Physical Bath Selection and Neutrino Mixing in VERSF

The common action construction of the neutral spectrum and its bath response

Keith Taylor

VERSF Standard-Model Derivation Programme

PBSN1 • 25 September 2026

General Reader Summary

Neutrino masses and neutrino mixing are different questions. Masses say how the neutrino modes travel; mixing says how the states produced with an electron, muon or tau overlap those travelling modes. A theory must fix both the neutrino equations and the charged-lepton reference frame, and this paper gathers the parts of VERSF that already do so into one construction.

The environment, called the bath, can help choose the state on which the calculation is done and can change the equations. But one particular set of coordinates used to describe bath fluctuations does not need to be excited before the direct construction can work. One existing kernel gives definite mixing numbers under stated assumptions; the mass matrix that was tested is a different construction, so those angles cannot simply be attached to its masses.

The tested candidate fails. Its three mostly-active neutrino states weigh about 0.05, 11 and 2300 times its own electron, and a common clock cannot change those ratios. The audit finds a single cause inside this candidate: the one weighting it uses to separate its generations is applied to the neutrinos exactly as to the charged leptons, and that both makes the neutrino masses far too widely separated and pins each neutrino flavour to one mass state. Take that weighting away and two of the three generations become exactly alike, so one required mass difference is absent. Moving the weighting to other places within the candidate does not help either. Within this construction, what is missing is a gentler rule that tells the neutrino generations apart; whether another part of the theory supplies one is the open question.

The last part follows the most promising source of such a rule. A folding process in the theory does generate two distinct generation-dependent patterns, enough in principle for three masses and a rotation between them, but the present equations do not choose the fold's angle and the simplest guess for it is wrong. Tracing the problem to its root identifies the first quantity the theory must return before the rest can be computed, so the next paper can compute that one object first, with several further steps still to follow, rather than scan another set of neutrino numbers.

Abstract

The VERSF corpus supplies a direct route from completion geometry to the neutral mass operator and its charged-current mixing frame; physical excitation of an auxiliary current chart is not a necessary intermediate step. We assemble the same-action neutral bilinears, kinetic normalization, closure Hamiltonians and weak-current bridge, incorporate the native record clock and stationary traffic of RCHS-25–27, the frequency-dependent normalization of the compatible scalar source and the full mixed-tensor completion descent, and state the compatibility conditions under which a

closure eigenframe becomes the physical PMNS frame. An independent evaluation of the published three-generation kernel reproduces its conditional angles, 33.401601°, 8.593903° and 48.369145°, with $\delta_{CP} = 227.235343°$, conditional on the kernel, charged frame, weak bridge and mass-column identification. An independent-admission extension of the stopped clock gives a mean waiting count proportional to $1/q$ and a mass-density scale proportional to q under fixed action per completed cycle, without computing mode coupling.

On a single shared charged and six-state neutral candidate, the three most active neutrino-to-electron mass ratios are (0.0544079021, 11.0846740, 2342.09765) at level 4. These are failed candidate outputs. An audit of that candidate shows why: with the address-depth factor D stripped, its Dirac and sterile cores are circulant to machine precision, so within this implementation the whole generation structure of the neutral sector is the same factor D that gives its charged block the ratios $1{:}\,204{:}\,42392$ (not the successful RRHF ratios $1{:}\,207{:}\,3576$), and that factor forces both the 14^{-4} gap ratio and the electron pole at once. Eight reassignments of D within the candidate do not repair it. The nonlinear single-fold source yields two noncommuting half- and quarter-sequence neutral responses, sufficient for three masses and frame rotation, but the equal-winding shortcut is falsified, the retained fold action is angularly flat because $\parallel Q(x) \parallel^2 = (x_1^2 + x_2^2)^2$, and faithful D_6 typing excludes the reduced scalar J–H shortcut. The representation-resolved current-history problem requires a rank-six $E_2 \otimes (C, W, Y)$ score; a positive history law with that geometry exists whose local Fisher metric is fixed up to one common gain g_{cyc} once the factor Gram of RCHS-1 is supplied, and the retained HCMR/RRHF source returns rank zero on that leg. The identified upstream bottleneck is the background-contracted graded source jet; the fold amplitude, orientation, restoring coefficient, physical addresses, neutral bilinear, kinetic normalization and current remain downstream of it. No neutrino mass, splitting or mixing angle is used to choose any coefficient.

Table of Contents

1 The prediction problem and the existing construction

The target is a definite neutral spectrum and a definite charged-current mixing matrix derived from VERSF inputs. The relevant construction is already distributed across the Master Substrate Action, history and branch closure, completion-channel Yukawa, weak-commitment, PMNS closure and neutral-sector papers [1–8]. Their results must be composed with their assumptions intact. An absence statement made for one finite implementation is not a theorem that the corresponding object is absent from the theory.

Three meanings of bath need to be distinguished. The Bath Criterion concerns shared conserved weight and the possibility of continuous redistribution [9,10]. A bath-resolved action describes physical unresolved modes and their state. An executable current chart describes selected coordinates of a history law. An identification between these objects is a mathematical assertion about embeddings and operators; terminology alone does not establish it.

This distinction changes the organization of the calculation. NSTC-1 explicitly allows the neutral sector to be obtained either through a source-derived neutral selector or through direct active–singlet mixed derivatives of the marked action [8, section 7]. On the second route, another selector is unnecessary. The neutral bilinear is computed on its typed carrier and then canonically normalized.

Bath dynamics enter when they select the background, contribute to those derivatives, or generate a self-energy. They need not act through an arbitrarily chosen set of current-score coordinates.

We therefore take the direct same-action route as the central construction. Sections 2–4 identify its mass and mixing operators and reproduce the existing conditional benchmark. Sections 5–7 incorporate the established temporal, source-normalization and completion results. Sections 8–9 retain the covariance and cycle-sector response calculations [24,25] within their proper domain. Section 10 specifies the observables and common-action assembly. Sections 11–12 evaluate admission, phase response and the common-completion candidate. Section 12.8 audits the shared candidate and locates the common origin of its two failures. Section 13 derives the nonlinear single-fold generation law, executes its first blind spectrum test, and then follows the failed shortcut upstream through the physical source-attachment problem. Sections 13.8–13.12 close the direct A1/current shortcuts, establish the fold-angle degeneracy at the retained order, identify the faithful rank-six current-history carrier and state the exact graded-jet release condition. Section 14 states the resulting physical target and conclusion.

The contribution is this composition and its verified consequences. The block reductions and spectral identities used below are standard mathematical operations. The VERSF content lies in the source operators, their physical identifications and the branch conditions under which those operations apply.

2 The neutral bilinear from the common action

2.1 Carrier and source level

Fix the physical branch, regulator, background and action level. The Master Substrate Action supplies amplitude, link, current, record, completion and matter variables; the history construction supplies the associated event law [1,2]. The neutral carrier must be the carrier admitted by the neutral branch and census results [7,8,11]. A convenient matrix written on an unrelated carrier is not a substitute.

Let $\hat{\Gamma}$ denote this common action, with all retained source insertions marked. For definiteness, write an all-left neutral multiplet $N = (\nu_L, N_R^c)^T$. Its quadratic mass term is represented by the complex symmetric matrix

$$\mathcal{M}_N = \begin{pmatrix} M_L & M_D \\ M_D^T & M_R \end{pmatrix}, \qquad \mathcal{M}_N^T = \mathcal{M}_N. \qquad (1)$$

The direct active–singlet mixed derivative supplies M_D; the singlet-completion bilinear supplies M_R; and the irreducible active bilinear supplies M_L. These are the three blocks requested by NSTC-1 [8, section 6]. A direct construction of these blocks already specifies the neutral operator. It does not require a second proof that a chosen auxiliary current coordinate is physically excited.

The word irreducible is essential. If sterile modes remain explicitly in (1), their exchange must not also be inserted into M_L. Where a controlled heavy-sterile expansion applies, eliminating them gives

$$M_{light} = M_L - M_D M_R^{-1} M_D^T + \cdots. \qquad (2)$$

Equation (2) is an approximation to a justified separated regime. The full matrix (1) remains the primary object when that separation has not been established. Setting $M_L = 0$ requires the applicable zero theorem or a declared approximation; it is not part of the definition of a neutral sector.

2.2 Relaxation and scalar differentiation

The primitive typed left-right bilinear and its completion insertion are provided by the action and Yukawa construction [1,3,12]. In a chosen carrier frame, denote the raw Dirac bilinear by B_{LR}. If the raw Higgs amplitude is h, the raw Yukawa insertion is its total derivative along the relaxed background,

$$Y_\nu^{raw} = \frac{dB_{LR}}{dh}, \qquad \frac{dB_{LR}}{dh} = \partial_h B_{LR} + \partial_a B_{LR}\,\frac{da_*}{dh}. \qquad (3)$$

For auxiliary coordinates satisfying $\hat{\Gamma}_a(a_*, h) = 0$, their derivative is

$$\frac{da_*}{dh} = -\hat{\Gamma}_{aa}^{-1}\hat{\Gamma}_{ah}. \qquad (4)$$

The inverse is taken on the constrained auxiliary domain. Thus the familiar Schur response is already available inside the direct mass construction. Equations (3)–(4) include completion and other relaxed fields without routing every contribution through a lower-dimensional probability chart. If a kinetic frame varies with h, differentiation of a canonically normalized vertex also includes that frame variation; the background value of a raw derivative alone does not include it.

The Majorana block is a bilinear on two same-chirality fermion legs. If E_N is its carrier embedding and $S_N = S_N^T$ its raw coefficient, then $M_R = E_N^T S_N E_N$. Symmetrizing an unrelated transport endomorphism does not derive this coefficient. This typing condition is already part of the mass and charge-conjugation construction [7].

2.3 Kinetic normalization and the current

The kinetic reconstruction and Higgs propagation papers supply finite kinetic operators and protected geometric reductions on their declared branches [13–15]. These results are inputs to the neutral calculation. They must not be replaced by an assumption that every kinetic residue is unknown, or by an identity matrix chosen for convenience.

For a positive neutral kinetic form Z_N on the admitted carrier, choose C_N with $C_N^\dagger Z_N C_N = I$. With $N = C_N N_c$, the canonical symmetric mass matrix is

$$\mathcal{M}_N^c = C_N^T \mathcal{M}_N C_N. \qquad (5)$$

A Dirac Yukawa written in left-right notation instead transforms as $Y^c = C_L^\dagger Y^{raw} C_R Z_H^{-1/2}$ when h is unnormalized. Charge conjugation converts the right-handed frame when that block is inserted into the all-left matrix (1). The transpose in (5) and the adjoint in the Dirac expression serve different purposes and must not be interchanged.

Let $U^T \mathcal{M}_N^c U = D_m$ be a Takagi factorization with $D_m \geq 0$. Let U_{eL} diagonalize the canonical charged-lepton left mass square. The charged-current coefficient for all neutral eigenstates is

$$\mathcal{V} = U_{eL}^\dagger \mathcal{J}_c U, \qquad U^T \mathcal{M}_N^c U = diag(m_1, \dots, m_n). \qquad (6)$$

Here $\mathcal{J}_c$ is the canonical weak-current bridge, including the active embedding. A three-column PMNS matrix is justified only after the spectrum and current establish which three states form the observable light active sector. If the selected matrix is nonunitary, its singular values and probability deficits are physical information. Replacing it by its polar unitary factor would change the theory being tested.

3 When closure geometry determines physical mixing

3.1 The normalized closure operator

The history closure and weak-doublet papers define ordered compression onto the weakly participating carrier [2,5,16]. In the notation of the PMNS closure theorem,

$$C_W = P_Y P_R P_C, \quad G_W = C_W C_W^{\dagger}, \quad \hat{C}_W = G_W^{-1/2} C_W, \quad H_W = \hat{C}_W H_{cl} \hat{C}_W^{\dagger}. \qquad (7)$$

The inverse square root is restricted to the support of G_W. The order and normalization of the compression are part of the construction. From it arise the charged-lepton closure operator h_ℓ and the neutrino closure kernel K_ν [5].

The weak-commitment theorem describes a regime in which a neutrino operator has the form

$$T_\nu = D_0 I + \varepsilon K_\nu, \qquad \varepsilon > 0. \qquad (8)$$

In that regime, residual diagonal gaps and off-diagonal transport co-scale with ε [4]. The eigenframe is independent of the common positive scale. This explains how substantial mixing can coexist with a small neutrino scale. It does not identify the eigenvalues of K_ν with masses or determine the value of that scale.

Completion-channel misalignment gives a complementary statement. A Yukawa operator $Y = U \Lambda W^{\dagger}$ has left mass square $YY^{\dagger} = U \Lambda^2 U^{\dagger}$ [3,6]. Relative left frames determine mixing after the weak bridge is included. The diagonal values in Λ require the completion calculation; they are not fixed by the singular-value identity itself.

3.2 Compatibility with the mass operator

The PMNS closure theorem states the compatibility required to transfer a closure eigenframe to the physical mass problem [5]. In the three-active Majorana reduction, use the convention $U_\nu^T M_\nu^c U_\nu = D_\nu$. The Hermitian operator with eigenvectors U_ν is

$$H_\nu^T = M_\nu^{c\dagger} M_\nu^c, \qquad H_\ell^L = M_\ell^c M_\ell^{c\dagger}. \qquad (9)$$

The superscript T labels the Takagi lift; it is not an instruction to transpose the displayed product. Using $M_\nu^c M_\nu^{c\dagger}$ instead conjugates the eigenframe and can reverse a CP-odd conclusion.

For simple spectra, common eigenframes follow from

$$\left[h_\ell, H_\ell^L\right] = 0, \qquad [K_\nu, H_\nu^T] = 0. \qquad (10)$$

The physical mass ordering, the current bridge and active completeness must also agree. For rank-one charged and neutral projectors, the invariant overlap is

$$|U_{\alpha i}|^2 = Tr\left(P_\alpha^\ell \mathcal{J}_c P_i^\nu \mathcal{J}_c^\dagger\right). \qquad (11)$$

This formulation removes arbitrary eigenvector phases and makes the charged reference frame explicit. It does not remove physical Majorana phases, which require the complex mass bilinear rather than its Hermitian square.

Assembly proposition. On a common canonically normalized three-active carrier, suppose (10) holds, the spectra are simple, the mass-column association is fixed and the weak bridge is unitary.

The closure projectors then determine the physical mixing probabilities through (11). The positive Takagi values of the mass bilinear determine the masses independently of those projector overlaps.

Proof. Commuting Hermitian operators with simple spectra share rank-one spectral projectors. Applying this separately to the charged and neutral pairs identifies the two physical eigenframes up to phases. Inserting their projectors in the common current gives (11). Equation (6) fixes the positive mass values. Neither the commutators nor the overlaps assign those values. Degeneracies require projectors onto whole eigenspaces and a further resolution within them. Condition (10) is a transfer condition, sufficient for carrying a closure frame to the mass problem; it is not a requirement on the theory. A direct calculation that returns H_ν^T itself fixes the frame without any closure kernel.

This proposition locates the direct calculation precisely. The key test is compatibility of source-derived closure and mass operators on the same carrier. Exciting a supplementary current chart is one possible way to perturb those operators, not an additional hypothesis of the proposition.

4 The existing numerical mixing kernel

Section 21 of the PMNS closure paper supplies the Hermitian example [5]

$$K_\nu = \begin{pmatrix} 0 & 1 & 1+\beta e^{i\psi} \\ 1 & r & 6/5 \\ 1+\beta e^{-i\psi} & 6/5 & r \end{pmatrix}, \quad r = \sqrt{3/2} - 6/5, \quad \beta = \sqrt{3}/20, \quad \psi = 3\pi/4. \qquad (12)$$

An independent Hermitian diagonalization gives ascending eigenvalues $(-1.184544169, -0.880766558, 2.114800470)$. The numerical eigen-equation residual is 1.74×10^{-15}. Following the example's column convention, the smallest electron component labels column 3; the remaining columns are ordered by decreasing electron magnitude. Relative to ascending eigenvalues, this gives columns (2,3,1).

The resulting modulus matrix is

$$|U| = \begin{pmatrix} 0.825459 & 0.544323 & 0.149430 \\ 0.310056 & 0.598057 & 0.739048 \\ 0.471680 & 0.588251 & 0.656870 \end{pmatrix}. \qquad (13)$$

Table 1. Independently reproduced conditional benchmark from equation (12).

Quantity	Value
θ_{12}	33.401601°
θ_{13}	8.593903°
θ_{23}	48.369145°
J_{CP}	−0.0244756590
δ_{CP}	227.235343°

The phase is recovered from both sine and cosine invariants, not from an inverse sine alone. We use $J_{CP} = Im\left(U_{e1} U_{\mu 2} U_{e2}^* U_{\mu 1}^*\right)$. Complex conjugation or a single muon–tau row exchange reverses this invariant; performing both preserves it. The attachment difference $|K_{e\tau}|^2 - \left|K_{e\mu}\right|^2 = 2\beta\cos\psi + \beta^2 = -0.1149744871$ is negative despite the upper-octant angle. Its sign alone is therefore not an octant theorem for this matrix.

These are definite mathematical consequences of (12), and should be retained. Their status is equally definite. The example assumes the displayed kernel, $U_{eL} = I$, $\mathcal{J}_c = I$, closure–mass compatibility and a complete three-active sector. Its electron-component labeling is an audit convention rather than a derivation of the physical mass order. No fit to new experimental data is performed here.

To promote Table 1 to a physical prediction, the same-action return must reproduce these assumptions or replace them with its actual values. In particular, arbitrary positive numbers d_i can share the same Takagi frame: $M = U^* diag(d_i) U^\dagger$ satisfies $M^\dagger M = U diag(d_i^2) U^\dagger$. Consequently, the mixing frame alone cannot fix neutrino masses. This identity does not assert that VERSF leaves the masses arbitrary; it identifies the role of the completion bilinear in fixing them.

5 The native clock and current history are already available

5.1 Record stopping and stationary traffic

RCHS-25 provides a source-stopped record clock on the directed D36 event chain [17]. Starting from its stationary distribution, T is the first admissible return to the starting state after all seven states have been visited, with $T \geq 7$. The evaluated law gives

$$\mathbb{E}T = 21.277682950918, \quad sd(T) = 9.355754165581, \quad \Pr(T = 7) = 0.006694060598121. \qquad (14)$$

This is a native event-grain construction. An independently fitted exponential hazard is not needed for this branch. A conversion to a dimensionful propagation clock must still follow the corresponding physical attachment; the presence of a discrete stopping law does not itself assign seconds to a Fold.

RCHS-26 constructs the orientation-even stationary traffic operator [18]. For stationary weights $\Pi = diag(\pi)$ and transition matrix P,

$$P^\sharp = \Pi^{-1} P^T \Pi, \qquad P_+ = \frac{1}{2}\left(P + P^\sharp\right). \qquad (15)$$

On the declared local class, this is the source-selected symmetric traffic law. Its matter attachment uses the EPM premise that the same pre-record Fold acts on the matter tangent through this traffic. The per-Fold convention and internal lift are therefore available source data, with a named physical premise. They should not be replaced by a claim that no temporal generator exists.

The relevant discrete eigenvalues include $\mu_{E1} = 1/2$ and $\mu_{E2} = -0.09938822758608408$. A negative eigenvalue of a discrete stochastic operator is compatible with nonnegative transition probabilities. It does prevent identifying that mode, without further structure, with the exponential of a positive scalar heat generator. The discrete Fold and the continuous scalar response of section 6 are distinct constructions until an attaching map relates them.

5.2 Terminal response and occupation response

Write $G(q) = \mathbb{E}(q^T)$ for the generating function of the native clock. Under the declared single-excitation, no-reinjection evolution, a homogeneous mode with multiplier μ has terminal and occupation factors

$$\mathbb{E}(\mu^T) = G(\mu), \qquad \mathbb{E}\left(\sum_{k=0}^{T-1} \mu^k\right) = \frac{1 - G(\mu)}{1 - \mu}. \qquad (16)$$

These identities follow pathwise from the finite geometric sum. The corresponding squared factors use μ^2. RCHS-27 and its continuation evaluate them for the E_2 mode [19].

Table 2. Native clock factors on the declared homogeneous single-excitation branch.

Quantity	Value
μ^2	0.009878019782703
Signed terminal factor $G(\mu)$	$-5.154458386 \times 10^{-10}$
Squared terminal factor $G(\mu^2)$	$6.283870234 \times 10^{-17}$
Root mean squared terminal amplitude	$7.927086623 \times 10^{-9}$
Signed occupation factor	0.90959678794
Squared occupation factor	1.00997656852

The small terminal value does not imply small accumulated occupation. For this branch, the occupation response remains of order unity. The raw current-occupation coefficient is $\eta_J = 0.005706350289534$, and the colour, weak and abelian E_2 current Gram has rank six [19]. This is a concrete current-history result, not an unspecified future input.

Its type still matters. A path-score Gram is an information or occupation metric. It becomes an action stiffness only if the action–history identification proves that equality with its units and normalization. The stationary score exposure and the propagated occupation can differ because they measure different histories. These results supply usable ingredients for bath response while preserving the direct neutral route of section 2.

Section 11 extends this same stopped-clock construction with an explicit admission factor and calculates which endpoint and occupation factors change. Sections 11.7–11.8 distinguish that execution from the mode coupling law and the subsequent conditional static phase calculation. The ratio calculation in section 12 removes a common clock conversion while retaining the relative neutral and charged dynamics.

6 Source normalization and completion descent

6.1 A compatible source has a frequency dependent normalization

The gauge-conserving temporal attachment and compatible-source execution provide an explicit example of how a static source identity extends in time [20,21]. On the stated 334-variable action chart, the Euclidean quadratic operator is

$$K(\omega_E) = (J_D + i\omega_E E)^\dagger (J_D + i\omega_E E) + J_V^\dagger J_V. \qquad (17)$$

The second term includes the action's potential rows. For the selected scalar source, the static coefficient map satisfies $C_0 J_D = J_s$ and $C_0 E = J_s/7$. Requiring the same source identity at nonzero frequency gives

$$C(\omega_E) = \frac{7}{7 + i\omega_E} C_0, \qquad C(\omega_E)(J_D + i\omega_E E) = J_s. \qquad (18)$$

Thus the static source normalization cannot simply be reused at every frequency. In the chosen Fourier convention, (18) corresponds to $\dot{s} + 7s = 7f$, together with its homogeneous initial-data term. This coefficient identity is not by itself a physical damping law for neutrino propagation.

On the real scalar-density branch, relaxing the currents gives the evaluated density operator [21]

$$K_\rho(\omega_E) = 2I + (I + L)^{-1}(L^2 + \omega_E^2 I). \qquad (19)$$

The nonuniform Laplacian eigenvalues are (2,2,4,4,5,7); after canonical normalization, the squared frequencies are (10,10,26,26,37,65). In particular, the source aligned with the $\lambda = 7$ mode produces a rank-one bath correction

$$\Delta K_b(\omega_E) = -\frac{8}{65 + \omega_E^2}\, gg^T, \quad \Delta L = -\frac{8}{65}\, gg^T, \quad \Delta Z = \frac{8}{65^2}\, gg^T. \qquad (20)$$

Here g is the normalized retained source direction. Equation (20) is an explicit source-induced stiffness and kinetic correction on this branch. It gives more information than a static force or a path-score variance alone. The associated conditional Gaussian oscillator has $\Omega = \sqrt{65}$; a physical use requires the same preparation and action normalization as the calculation.

The local completion displacement vanishes in this scalar test. That statement concerns one source and one response. It does not assert that the full mixed neutral bilinear, or its response to other action variations, vanishes.

6.2 The full completion tensor descends on the declared source stack

The finite fermion-kernel reconstruction contains a stronger completion result than the earlier primitive-leg-only statement [13, Corollary 17.1, pp. 43–44]. On its RRHF/FCHI single-primitive-leg completion structure,

$$c_H = \frac{1}{2}(1,1,1,1), \quad \| c_H \| = 1, \quad T_f = c_H \otimes B_f. \qquad (21)$$

The tensor B_f carries representation, generation, chirality and branch dependence. The mode-1 quotient acts on the primitive leg. Multiplicativity of the Frobenius norm therefore gives, for every nonzero B_f,

$$\frac{\| (I - P_{full})c_H \otimes B_f \|_F}{\| c_H \otimes B_f \|_F} = \frac{\| (I - P_{full})c_H \|}{\| c_H \|} = 1.149 \times 10^{-15}. \qquad (22)$$

The defect-only quotient has residual exactly one. The potential rows are therefore essential to the executed full-action descent. Within the stated structure, an additional sector-by-sector tensor publication is not mathematically required to establish the normalized mode-1 residual. The earlier claim that only a primitive leg passes while full tensor descent remains pending is superseded by this corollary.

Equation (22) establishes compatibility of the entire declared mixed tensor with that quotient. It does not determine the entries of B_f by projecting c_H. Those entries retain their role in the mass calculation. Likewise, the 89-page reconstruction supplies further geometric kinetic results on its declared physical spatial branch [14]; it should be consumed at that strength while preserving its source-identification premises.

Together, (18) and (22) correct two different reductions. A source must be normalized in the temporal operator actually used, and a completion tensor must be tested against the full action. Neither correction creates an additional bath-excitation requirement for the direct mixed derivative.

7 Where physical bath response enters the direct construction

7.1 Joint reduction of the operator and forcing

For retained bath coordinates b and auxiliary coordinates a, write the linearized equations at the selected action level as $\mathcal{A}x = f$. Exact block elimination gives

$$\mathcal{A}_{eff} = \mathcal{A}_{bb} - \mathcal{A}_{ba}\mathcal{A}_{aa}^{-1}\mathcal{A}_{ab}, \qquad f_{eff} = f_b - \mathcal{A}_{ba}\mathcal{A}_{aa}^{-1}f_a. \qquad (23)$$

Both expressions follow by solving the auxiliary equation and substituting it into the retained equation. The inverse includes the declared constrained domain and, for real-time response, the causal prescription. Initial data contribute a homogeneous solution. A Euclidean inverse, including (17), is not automatically a retarded response.

A change of retained coordinates transforms the operator, forcing and matter coupling together. A numerical stiffness in one chart cannot be paired with a source measured in another chart without that transformation. Bath degrees of freedom already eliminated into the action must not be integrated a second time.

If the kinetic expansion is admitted, its canonical transformation also acts on the neutral insertion and on the charged current. This is why the action-level elimination and the field normalization in section 2 belong to one calculation.

7.2 The current chart is a tested reduction

Suppose a physical history law $p_b(\omega)$ is known. Its centered scores are $s_a = \partial_{b_a}\log p_b|_0$. Let S_α be the centered scores of an existing current chart, with $F_{\alpha\beta} = \mathbb{E}(S_\alpha S_\beta)$. The least-squares projection and residual are

$$L_{\alpha a} = (F^+)_{\alpha\beta}\mathbb{E}(S_\beta s_a), \qquad r_a = s_a - S_\alpha L_{\alpha a}. \qquad (24)$$

The pseudoinverse handles chart redundancy. For a normalized classical history law and a history observable O, differentiation gives $\partial_a \mathbb{E}O = \mathbb{E}(\partial_a O) + \mathbb{E}[(O - \mathbb{E}O)s_a]$. The residual in (24) can therefore be discarded only if its downstream contribution vanishes or is evaluated separately.

When the projection is exact for the required observable, the current chart efficiently carries the bath response into the previously computed history and neutral derivatives [22,23]. When it is not exact, the direct derivative of $\hat{\Gamma}$ remains available. A projection failure diagnoses that reduction; it does not establish the absence of the direct neutral operator.

7.3 Mass and mixing derivatives

For a nondegenerate positive Takagi value, a variation of the canonical mass matrix obeys

$$\delta m_i = Re\left(u_i^T \delta M_\nu^c u_i\right), \qquad \delta H_\nu^T = \delta M_\nu^{c\dagger} M_\nu^c + M_\nu^{c\dagger}\delta M_\nu^c. \qquad (25)$$

Kinetic-frame variations are included in δM_ν^c. The corresponding rank-one projector derivative is

$$\delta P_i = \sum_{j \neq i} \frac{P_j \delta H_\nu^T P_i + P_i \delta H_\nu^T P_j}{m_i^2 - m_j^2}. \qquad (26)$$

Differentiating (11) then includes the charged projector, the neutral projector and both occurrences of the weak bridge. Equations (25)–(26) connect the source-derived bath response directly to mass shifts and mixing probabilities. They avoid assigning physical meaning to a phase-dependent eigenvector derivative. Degenerate or zero-mass limits require the corresponding subspace treatment.

For slow classical fluctuations, the chosen quasistatic observable can be averaged over the returned bath state. For rapid fluctuations, the appropriate fermion two-point function and self-energy determine propagation, with coherent and dissipative terms where present. Diagonalizing an averaged matrix and averaging its mixing probabilities are different operations. The existing fluctuation calculation below evaluates a specified quasistatic readout.

8 The executed cycle sector is a bounded example

The cycle-sector source execution [25] retains 18 current observables on a 334-variable unit-metric action chart: three cochains, real and imaginary parts, and three gauge factors [25]. The static Hessian has rank 313. For the retained observation matrix R, the null-space overlap is below 3.0×10^{-15}. With its Moore–Penrose inverse H^{+}, define

$$G = R^T H^+ R, \qquad K_{cyc} = G^{-1}, \qquad T = H^+ R K_{cyc}. \qquad (27)$$

The lift satisfies $R^T T = I$ and the full auxiliary stationarity condition. In the tested continuity-only temporal extension $K(\omega_E) = H + \omega_E K_1 + \omega_E^2 K_2$, the evaluated coefficients satisfy $\| K_1 T \|_F < 9.9 \times 10^{-15}$ and $\| K_2 T \|_F < 5.6 \times 10^{-15}$. Consequently the retained kernel is frequency independent within numerical precision on this extension.

All real-cycle stiffnesses are one. The imaginary colour-cycle entries are 0.998970635285, 0.999513168784 and 0.999513168784; the other imaginary-cycle entries are one. The lift includes nonzero gauge-link and record-phase components. The compatible scalar source has reduced overlap below 6.0×10^{-17} with these cycles. Direct checks at $\omega_E = 0, 0.1, 1, 10, 100$ confirm the polynomial certificate; the largest checked stationarity residual is 5.6×10^{-11} at the highest frequency.

These results concern the selected cycles and temporal extension. Their zero scalar-source overlap does not invalidate the nonzero scalar-density response (20), the native traffic operator (15), or the direct neutral mixed derivative. Each acts on a different specified subspace.

If the flat cycle kernel is assigned a conditional Gaussian action with unit action normalization, its formal covariance is $K_{cyc}^{-1}\delta(\sigma - \sigma')$. A symmetric frequency cutoff gives an equal-time covariance $(\omega_c/\pi)K_{cyc}^{-1}$. This expression is cutoff dependent; it is not a selected finite physical variance. The earlier execution therefore supplies a concrete stiffness and decoupling result, not a universal verdict on bath selection.

9 The retained second order neutral response

9.1 Constrained current geometry

The covariance calculation [24] evaluated the inherited level-4 current-to-neutral map [22–24]. Write $t = (t_0, t_c, t_s)$ and $z = (t_c, t_s)$. On the fixed-one-bit history manifold, the local entropy constraint gives

$$t_0 = -\alpha |z|^2 + O(|z|^3), \qquad \alpha = 0.17519593762712266. \qquad (28)$$

For any smooth chart output O, denote its free derivatives at the origin by $O_0 = \partial_{t_0} O$ and $O_{ij} = \partial_{t_i} \partial_{t_j} O$. Its constrained Hessian is

$$Q_{ij} = O_{ij} - 2\alpha\delta_{ij}O_0. \qquad (29)$$

For a centered small distribution with covariance Σ, this gives

$$\mathbb{E}O - O(0) = \frac{1}{2} \sum_{i,j \in \{c,s\}} Q_{ij}\, \Sigma_{ij} + O(\mathbb{E} \parallel z \parallel^3). \qquad (30)$$

A two-coordinate covariance has three independent entries. With real vectorization of complex outputs, the coefficient matrix has columns $\frac{1}{2} rvQ_{cc}$, rvQ_{cs} and $\frac{1}{2} rvQ_{ss}$. Its numerical rank is three for each evaluated output. The physically allowed domain remains the positive-semidefinite covariance cone. Rank three in second moments is not a claim of three independent first-order modes.

9.2 Numerical coefficients and inherited scope

The five chart inputs are $\chi = (g_{AB}, g_{BC}, g_{CA}, w_H, \gamma_T)$: three generation conductances and two history/completion coefficients. For isotropic covariance $\Sigma = vI$,

$$\mathbb{E}t_0 = -0.350391875254\, v + O(\mathbb{E} \parallel z \parallel^3). \qquad (31)$$

The coefficients of $\big(\mathbb{E}\chi - \chi(0)\big)/v$ are -0.005344947620 for each conductance, -0.002178717562 for w_H, and -0.00004658106641 for γ_T [24]. Their propagation gives the following norms.

Table 3. Conditional isotropic response per unit coordinate variance.

Chart output	Norm of response coefficient
Five inputs χ	9.5107505×10^{-3}
Dirac chart Y_ν	1.3430129×10^{-3}
Sterile chart M_R	9.4836146×10^{-4}
Full neutral chart $\mathcal{M}_6$	9.5929883×10^{-4}
Schur matrix	8.3050375×10^{-6}
Charged-frame probability array	7.4806907×10^{-4}
Record kernel	1.1922353×10^{-6}

The three singular values of the full-neutral covariance response are $(9.6385455 \times 10^{-4}, 8.5299535 \times 10^{-5}, 3.7399456 \times 10^{-5})$. Thus an isotropic source explores one fixed combination of three independent covariance responses; it does not supply three adjustable physical amplitudes.

The evaluation retains the archived D36 kernel, completion prescription, magnetic phase, address geometry and charged frame. It also retains that archive's zero active Majorana block, symmetric sterile compression and frozen kinetic convention. Those choices define the differentiated function. They do not certify the typed same-action return in section 2. The earlier candidate's active weight 1.996961 for its three lightest states and relative source discrepancy approximately 0.44865 remain properties of that candidate [22]. They are not corpus-wide impossibility results.

Central differences with steps 0.032, 0.016 and 0.008, followed by Richardson extrapolation, give relative changes below 6.8×10^{-9} for the matrix and five-input responses. The smaller probability and record responses have changes 1.88×10^{-7} and 4.36×10^{-5}. A four-point entropy-constrained ensemble at radius 0.005 agrees with the matrix coefficients within 4.43×10^{-6} relative error. This verifies the local curvature calculation at its stated precision.

The result remains useful: balanced fluctuations can change neutral coefficients even when the mean first-order displacement vanishes. The bath state determines whether this particular manifold and its covariance are physically realized. The direct mass calculation does not depend on that realization.

10 From the assembled operators to testable predictions

10.1 A finite same source calculation

The calculation can now be stated without repeating work already supplied by the corpus. Begin with the chosen common-action branch and its admitted carrier. Use the finite kinetic reconstruction, full-action completion descent and source normalization at their established strength [13–15,20,21,26,27]. Evaluate the active–singlet, singlet–singlet and irreducible active mixed derivatives on that same background. Carry the charged-lepton operator and weak current through the same canonical transformation.

Next solve the full neutral Takagi problem (6). Record the current weight of every state before identifying a light sector. Where the closure route is used to fix the mixing frame, evaluate (10) and match the mass columns; the conditional kernel in section 4 is then a concrete candidate to reproduce or replace. A direct mass-and-current calculation can also determine the mixing without first passing through that candidate kernel.

The temporal construction used in a dynamical dressing must be the one attached to those operators. RCHS-25–27 already supply the stopped Fold and occupation data on their declared branch. Equations (18)–(20) supply another explicit source response with its own normalization. Combining these results requires their attaching map, rather than an arbitrary conversion of one mode's decay into another mode's frequency. This is a source-matching step, not a demand for a new generic bath model.

10.2 The observables

For a justified unitary three-active reduction, the angles and Dirac invariant are read directly from the returned current matrix:

$$s_{13}^2 = |U_{e3}|^2, \quad s_{12}^2 = \frac{|U_{e2}|^2}{1 - |U_{e3}|^2}, \quad s_{23}^2 = \frac{|U_{\mu 3}|^2}{1 - |U_{e3}|^2}. \qquad (32)$$

Mass splittings follow from the Takagi values, $\Delta m_{ij}^2 = m_i^2 - m_j^2$. The complete complex matrix also supplies the Majorana phases when applicable. In the light Majorana regime, the standard absolute-mass combinations are

$$\Sigma_m = \sum_{i=1}^{3} m_i\,, \qquad m_\beta^2 = \sum_{i=1}^{3} |U_{ei}|^2\, m_i^2, \qquad m_{\beta\beta} = \left| \sum_{i=1}^{3} U_{ei}^2\, m_i \right|. \qquad (33)$$

The last quantity describes light-neutrino exchange with the usual left-handed current. Additional neutral states or other exchange mechanisms require their own terms. If the light charged-current block is nonunitary, the physical matrix $\mathcal{V}$ and its normalization enter the relevant amplitudes directly rather than being forced into (32).

A dimensionless mass ratio is a physical prediction when the dimensionless canonical blocks and their relative normalization are source-selected. Section 12 uses the charged electron singular value from the same construction to form m_{ν_i}/m_e, eliminating the common scale. Converting a resulting ratio to electronvolts using the measured electron mass is an electron-calibrated prediction, not an independent derivation of the dimensionful electron scale. Similarly, if a measured neutrino splitting calibrates a scale, that splitting is an input and cannot also count as an output.

Table 4. Numerical claims and the source object that determines them.

Claimed result	Determining object
Mass ratios and ordering	Canonical neutral bilinear and physical state labels; electron-normalized ratios also require the same-construction charged matrix and relative normalization
Masses in physical units	Same bilinear with the common dimensionful normalization
Mixing probabilities and Dirac phase	Charged and neutral frames with the weak-current bridge
Majorana phases	Complex symmetric mass bilinear and its Takagi convention
Bath-induced mass or mixing shift	Matched action response and state on those same operators
Regulator-independent prediction	The declared refinement comparison of those observables

These are independently testable claims. Agreement of a closure kernel's angles does not establish its masses; a completed mass spectrum does not establish a unitary three-active current; and a successful finite source certificate does not establish regulator convergence. Each comparison has an identifiable mathematical object.

10.3 Status of the direct construction

The direct neutral route, closure–mass compatibility theorem, native clock, current occupation law, compatible temporal source and full mixed-tensor descent are existing constructive inputs. Their roles are now explicit in one manuscript. The numerical mixing benchmark is reproduced with its CP sign and column convention; the earlier covariance response is preserved with its domain; and the auxiliary cycle zero is confined to the sector actually evaluated.

Section 11 adds an executed admission and clock calculation, two conditional neutral spectra and a static phase-response continuation. Section 12 then evaluates electron-normalized ratios from one shared charged and neutral construction. It separates completion-leg strength from coherent overlap and tests distinct coupling and address changes. These advances do not return a source-selected physical set of M_L, M_D, M_R, kinetic residues and charged-current coefficients. The displayed mass ratios diagnose the finite candidate; they do not promote Table 1 to an unconditional PMNS prediction or constitute physical neutrino mass predictions.

11 Void coupling anchoring and the emergent commitment clock

11.1 Admission anchoring and the mass convention

The calculation below varies a generic admission probability on the native record clock and evaluates two finite six-state implementations [28–32]. It was previously described as an incorporation of weaker void coupling and anchoring. Its actual scope is narrower: the phase distribution, alignment probability and native mode response were not recalculated. The admission factor was interpreted as a structural gate change, while the anchoring closure law was held fixed. Section 11.7 restores the separate coupling dependence required by the source papers.

The source distinguishes void coupling from void anchoring [28, section 2; 29, sections 1 and 5]. Coupling supplies micro-event opportunities through the relative phase dynamics. Anchoring is the accumulation and cycle closure required for irreversible commitment, characterized by K_c. A further structural gate g_a selects the aligned configurations that are anchorable, giving $p_{eff} = g_a p_{align}$. A highly coherent mode can have a small structural gate, but this possibility does not establish the coherence or coupling strength of the neutrino modes. Here g_a is distinct from the source vector g in (20). Under the action-per-completion convention [29,30],

$$\langle N \rangle = \frac{K_c}{p_{eff}}, \qquad m = \frac{\eta \hbar}{c^2 \Delta t \, \langle N \rangle}. \qquad (34)$$

Equation (34) retains the source's fixed-action assumption and its conversion from opportunity count to physical duration. No numerical value of Δt is inserted. An alternative mathematical-core note defines a resistance scale proportional to ticks per commitment [31]; that convention is not combined with the inverse-mean rest-energy convention used here. The two quantities require a physical identification before they can be equated.

11.2 Independent admission and the stopped clock

Let T be the first complete seven-state cover and return defined in section 5, with stationary-start mean $\mathbb{E}T = 21.2776829509186$ and standard deviation 9.3557541655814. It is a random completion count, not a fixed integer K_c. Write $G(z) = \mathbb{E}[z^T]$.

Introduce a relative extra admission probability $0 < q \leq 1$. Each native progress step is admitted independently with probability q; failed opportunities leave that progress unchanged. This is a declared extension of the native law. The variable q is neither an evaluated absolute gate g_a nor an assertion that the ungated source has $p_{eff} = 1$.

With independent geometric waiting counts W_k supported on the positive integers, the gated stopping count and its generating function are

$$T_q = \sum_{k=1}^{T} W_k \,, \qquad G_q(z) = G\left(\frac{qz}{1-(1-q)z}\right). \qquad (35)$$

Conditioning on T gives the exact moment identities

$$\mathbb{E}T_q = \frac{\mathbb{E}T}{q}, \qquad VarT_q = \frac{VarT + (1-q)\mathbb{E}T}{q^2}. \qquad (36)$$

If the action per completed cycle and the opportunity-to-duration conversion remain fixed, (34) then gives

$$\frac{m(q)}{m(1)} = \frac{\mathbb{E}T}{\mathbb{E}T_q} = q. \qquad (37)$$

Weaker admission and longer waiting describe the same mass suppression in this construction. Multiplying by an additional factor of q merely because the waiting time increased would count that effect twice. This is a conditional relation in common opportunity units, not a derivation of a second species-specific physical time coordinate.

Table 5. Exact clock and mass-density scaling for illustrative admission probabilities.

Relative admission q	Mean opportunities per completion	Relative mass from (37)
1	21.277683	1
0.1	212.776830	0.1
0.01	2127.768295	0.01
0.001	21277.682951	0.001

The uniform gate preserves the stationary distribution. Entropy production becomes $q\log 2$ per opportunity and remains $\log 2$ per admitted native step. It therefore changes the per-opportunity convention of the fixed-one-bit chart in section 9. No retuning is made to restore one bit per opportunity. The archived completion coefficient $\epsilon = 0.0521051847514$ is not identified with q.

11.3 Correlated propagation and stopping

The attachment of the gate to matter evolution determines the stopped response. Retain the single-excitation, no-reinjection premise of section 5. For a homogeneous source mode with multiplier μ, define $H(\mu) = [1 - G(\mu)]/(1 - \mu)$.

If the same admission event advances both matter and record progress, exactly T matter updates occur before commitment, irrespective of the failed opportunities. Thus

$$\mathbb{E}\left[y_{T_q}\right] = G(\mu)y_0, \qquad \mathbb{E}\sum_{n<T_q} y_n = \frac{H(\mu)}{q}y_0. \qquad (38)$$

The terminal second moment is also unchanged, with multiplier $G(\mu^2)$. Occupation per completed cycle increases as $1/q$, while its renewal average per opportunity is independent of q. For the native E_2 source mode this average is $H(\mu)/\mathbb{E}T = 0.0427488646222$, per unit initial amplitude. This is an occupation response, not an action stiffness or a mass.

If record progress alone is gated and matter continues updating on every opportunity, the terminal factor is instead $G_q(\mu)$. Table 6 evaluates both attachments at $\mu = -0.0993882275860841$. This source harmonic has not been identified with a neutrino mass eigenstate.

Table 6. Terminal response coefficients for two explicitly different gate attachments.

q	Shared matter and record admission	Record admission only
1	-5.15446×10^{-10}	-5.15446×10^{-10}
0.1	-5.15446×10^{-10}	-3.44795×10^{-17}

q	Shared matter and record admission	Record admission only
0.01	-5.15446×10^{-10}	-3.31763×10^{-24}
0.001	-5.15446×10^{-10}	-3.30493×10^{-31}

The difference is physical once an attachment is selected. The gate strength alone does not select a column. In the shared-admission case, replacing matter evolution by its averaged step and then treating stopping as independent is incorrect: at $q = 0.1$ that procedure gives $+8.64424 \times 10^{-5}$ instead of -5.15446×10^{-10}. The clock and updates must be averaged jointly.

11.4 Temporal response with the source attached consistently

For shared admission and a deterministic source b_n injected per admitted update, let $u_n = \mathbb{E}[y_n]$ denote the fixed-opportunity mean. It obeys

$$u_{n+1} = (1 - q + q\mu)u_n + qb_n, \qquad \chi_q(z) = \frac{q}{z - (1 - q + q\mu)}. \qquad (39)$$

Consequently,

$$\chi_q(1) = \frac{1}{1-\mu}, \qquad \chi_q\left(e^{i\omega}\right)^{-1} = (1-\mu) + \frac{i\omega}{q} + O(\omega^2/q). \qquad (40)$$

The temporal coefficient changes while the static response to the consistently attached source stays fixed. This is an exact discrete history response, with frequency measured per opportunity. Its linear frequency coefficient is not yet a fermion kinetic residue. A physical neutrino dispersion law requires its attachment to the neutral bilinear and current of section 2. The independent scalar factor in (18) is not multiplied into (39) without a map identifying their sources, carriers and clocks.

11.5 Two implementations in the six-state candidate

The scalar mass law leaves a further question: how does admission enter the neutral bilinear? To isolate it, retain the archived address geometry, charged-lepton matrix, identity kinetic and current conventions, and $M_L = 0$ [23,32]. These are declared assumptions of this finite candidate. It is distinct from the three-generation closure benchmark of section 4; the numerical angles in Table 1 are not substituted into this calculation.

Consider the family

$$M_D(q) = q^a M_D(1), \qquad M_R(q) = q^b M_R(1). \qquad (41)$$

The zero-momentum Schur matrix scales exactly as

$$M_S(q) = -M_D(q)M_R(q)^{-1}M_D(q)^T = q^{2a-b}M_S(1). \qquad (42)$$

Matching a leading light mass proportional to q requires $2a - b = 1$. It does not select a and b. We evaluate two choices: a fixed sterile block, $(a, b) = (1/2, 0)$, and a uniformly scaled full neutral matrix, $(a, b) = (1,1)$. The square root in the first choice follows from this matching condition; calling q a probability does not independently derive an amplitude factor. Both choices act on neutral completion coefficients and preserve the archived charged matrix by construction. Neither is a neutrino-only kinetic rescaling of a shared weak doublet.

All six Takagi values and current weights are computed before selecting the three most active states. Let their indices, ordered by mass, be i_1, i_2, i_3. Define the gap ratio and the phase-independent disappearance bound by

$$r_\Delta = \frac{m_{i_2}^2 - m_{i_1}^2}{m_{i_3}^2 - m_{i_1}^2}, \qquad 1 - P_{ee} \leq 1 - [\max(2p_{\max} - 1, 0)]^2. \qquad (43)$$

Here $p_{\max} = \max_i |\mathcal{V}_{ei}|^2$ uses all six states and the inherited normalized electron row. The bound follows from the triangle inequality for the coherent survival amplitude; it is not a prediction that a particular baseline attains the bound.

Table 7. Full six-state results for fixed sterile completion at refinement level 4.

q	Selected states	Minimum active weight	r_Δ	Disappearance ceiling
1	1, 2, 4	0.9859387	2.23988×10^{-5}	1.24585%
0.1	1, 2, 4	0.9993203	2.25122×10^{-5}	0.20582%
0.01	1, 2, 4	0.9537009	2.20637×10^{-5}	0.10061%
0.001	1, 2, 3	0.9995146	2.26641×10^{-5}	0.09008%

The purity dip at $q = 0.01$ reflects an active–sterile avoided crossing. The full calculation therefore matters even when the Schur scaling is simple. At $q = 0.001$, the three light masses are $(1.4179563 \times 10^{-11}, 2.9025777 \times 10^{-9}, 6.0969073 \times 10^{-7})\Lambda$, with Λ still uncalibrated. The charged ratios remain 1: 207.15058: 3576.08368.

As q tends to zero in this fixed-sterile implementation, the Schur gap ratio is $2.26447777 \times 10^{-5}$ and its electron probabilities are $(0.9997776723, 0.0002205492, 0.00000177851)$. The retained flavour hierarchy therefore survives the mass suppression. The previously supplied comparison value for the gap ratio, approximately 0.0298, is a diagnostic comparator [32], not an input to the executable or a new experimental-data compilation.

For uniform scaling of the full neutral matrix, every mass scales exactly as q, while every current probability stays fixed. The gap ratio remains $2.23988330 \times 10^{-5}$ and the disappearance ceiling remains 1.24585%. These calculations establish that the tested common gates can lower the neutral scale while retaining the charged hierarchy; they do not repair this candidate's relative masses and mixing.

11.6 What still prevents a mass derivation

The remaining task is to calculate the physical phase-coupling response and anchoring structure of the admitted neutrino modes, then evaluate their neutral mass operator. The scalar admission test supplies neither the phase-alignment probabilities nor the relative completion coefficients and their placement in the active and sterile blocks. Section 11.5 shows that two distinct six-state matrices obey the same leading scalar mass law. Scaling a fixed Schur matrix as in (42) leaves its mass ratios and eigenvectors unchanged; this does not constrain a coupling calculation that changes the underlying mode-dependent coefficients.

The required calculation has three identifiable outputs. First, obtain the relative phase history and its alignment probabilities, together with the separately defined structural admission and anchoring closure quantities. Section 11.8 evaluates a conditional static phase branch but not the physical neutrino history. Second, carry the returned contributions into the typed active–singlet, singlet–singlet and irreducible active derivatives, including relaxation, kinetic normalization and the charged-current bridge. This must determine actual block entries rather than select the exponents in (41) by convenience. Third, determine the relative charged and neutral normalization from that same construction. A common time and action calibration is additionally needed only for an independent dimensionful mass prediction; it cancels from the electron-normalized target in section 12.

The first three outputs determine the mass ratios and flavour weights. For frequency-dependent dressing, the propagation poles and their current residues must be obtained from the matched neutral two-point operator; a static Takagi problem is sufficient only when the resulting local mass and kinetic approximation is justified. Neither the homogeneous source multiplier in Table 6 nor the inverse susceptibility in (40) is a substitute for that operator.

A nonuniform admission profile can change the retained flavour pattern in the restricted history model below. Its normalization cannot be ignored. For example, if a row-stochastic history chain has state-dependent admission $Q_q = diag(q_i)$ with $0 < q_i \leq 1$, then

$$\hat{P} = I - Q_a + Q_a P, \qquad \hat{\pi}_x = \frac{\pi_x / q_x}{\sum_y \pi_y / q_y}. \qquad (44)$$

Let P_+ be the additive reversibilization of P in its stationary measure, and set $S = \Pi^{1/2} P_+ \Pi^{-1/2}$. Using the new stationary measure for the gated chain gives

$$I - \hat{S} = Q_a^{1/2} (I - S) Q_a^{1/2}. \qquad (45)$$

Unlike a uniform gate, a nonuniform gate can change the eigenvectors of this symmetric operator. Equations (44)–(45) specify how to carry a returned admission profile through this restricted history model. They do not assign the profile, identify seven history states with three flavours, or prove that the neutrino gate has this classical form.

The mode-resolved coupling and anchoring calculation must therefore be composed with the neutral completion and kinetic operators on one branch. Section 12 makes its immediate observable the three neutrino-to-electron mass ratios, preserving the relative scalar normalization. It evaluates the existing shared candidate and locates its unsuppressed neutral overlap, without choosing a replacement matrix from the desired neutrino outcome. That failed finite return does not establish that the required physical ingredients are absent from the corpus.

11.7 The distinct void coupling calculation

The admission calculation does not evaluate the resonant coupling of the neutrino modes to the substrate. The source papers distinguish the two physical processes [28, section 2; 29, sections 1 and 5]. Coupling determines the probability that a micro-event opportunity occurs; anchoring determines the accumulation and closure needed to commit an irreversible event. In the refined coupling model there is also a structural gate, the conditional probability that an aligned configuration is anchorable. The effective successful-event probability combines coupling and structural admission; anchoring depth enters the stopping rule. An aggregate admission factor cannot identify the separate dynamical origins of a change in event probability.

For a mode j, let $f_j(\Delta)$ be the normalized distribution of its relative phase on $[-\pi, \pi)$. Denote the alignment probability by A_j to distinguish it from the admission parameter q used above. Then

$$R_j = abs\left(\int_{-\pi}^{\pi} e^{i\Delta} f_j(\Delta)\, d\Delta\right), \qquad A_j = \int_{-\varepsilon_j}^{\varepsilon_j} f_j(\Delta)\, d\Delta. \qquad (46)$$

The phase window satisfies $0 < \varepsilon_j < \pi$. It is not the completion coefficient ϵ quoted in section 11.2. The source treats the window as substrate-defined and universal on its minimal branch; a channel-dependent window requires its own derivation. The phase law, rather than a selected numerical value of R_j, is the physical input that must be obtained from the mode dynamics.

Let g_j be the probability of anchoring admissibility conditional on alignment. The factorization $p_{eff,j} = g_j A_j$ then follows from conditional probability; it does not require the two events to be independent. With a fixed commitment threshold $K_{c,j}$ and independent successful increments, the source's counting law and conditional action bridge give

$$\mathbb{E}[N_j] = \frac{K_{c,j}}{g_j A_j}, \qquad m_j = \frac{\eta_j \hbar}{c^2 \Delta t}\frac{g_j A_j}{K_{c,j}}. \qquad (47)$$

Here the first expression is a pre-temporal tick count. The second uses the common tick-to-duration calibration and action-to-energy bridge stated in section 11.1. Coupling and anchoring remain distinct even though both contribute to the same completed-event density. If the phase history is correlated, a one-time probability A_j does not establish the independent-increment waiting law; the joint first-passage calculation must be used instead. Nor is $K_{c,j}$ identified with the random seven-state cover-and-return count T merely because both describe completion.

On the source's centered small-fluctuation Gaussian branch, $R_j = \exp(-\sigma_j^2/2)$, giving

$$A_j = erf\left(\frac{\varepsilon_j}{\sqrt{-4\ln R_j}}\right), \qquad \sigma_j \ll \pi. \qquad (48)$$

Equation (48) is distribution-dependent. The magnitude R_j alone does not determine a phase-window probability for an arbitrary circular distribution or a distribution locked away from the center of the window. For a weak-coupling calculation, the periodic phase density in (46) should therefore be retained. In particular, a uniform phase density gives $R_j = 0$ but $A_j = \varepsilon_j/\pi$, not zero. Extrapolating the unwrapped Gaussian expression to $R_j \to 0$ would incorrectly remove this finite-window floor. Any additional dynamical requirement for a random alignment to generate a micro-event must be specified rather than supplied by that extrapolation.

For physical modes to which (47) applies, a common time calibration cancels from the mass ratio:

$$\frac{m_i}{m_j} = \frac{\eta_i}{\eta_j}\frac{g_i}{g_j}\frac{A_i}{A_j}\frac{K_{c,j}}{K_{c,i}}. \qquad (49)$$

The factor A_i/A_j is the coupling contribution omitted from the numerical evaluation. It can differ between modes even when their structural gates are equal. The source also proposes a conditional multiple-sector alignment law, $A_j = s^{N_{eff,j}}$, with a common single-sector probability s. Under that independence hypothesis, the coupling ratio is $s^{N_{eff,i}-N_{eff,j}}$ [28, section 9]. The effective sector counts and their correlations must be calculated from the mode structure. They are not provided by the four illustrative values of q in Tables 5–7.

This correction does not call for multiplying the previous result by an unexplained second suppression. At fixed completion law and under the same independent thinning attachment, the aggregate admission change relative to a reference mode calculation would be

$$q_{eff,j} = \frac{g_j A_j}{g_j^{(0)} A_j^{(0)}}. \qquad (50)$$

Equation (35) describes thinning only when $0 < q_{eff,j} \leq 1$ and the native progress law is unchanged. A separately computed change in A_j belongs inside this effective probability. The reciprocal waiting-time change is then already determined by it. If coupling also changes the native progress matrix, the response multiplier, the completion law or the source attachment, its effect requires recalculation of those objects and cannot be represented by (50) alone.

The physical mode evaluation must obtain f_j and the alignment probabilities or their full joint history law, then combine them with the separately evaluated structural gate and anchoring closure. The conditional static calculation in section 11.8 is progress toward this attachment, not its completion. For a mixed neutral sector, scalar mode formulas alone do not supply an off-diagonal Majorana bilinear or the weak-current frame. Their coefficients must enter the typed neutral derivatives, kinetic operator and current on the same branch. The illustrative Adler dynamics in [28, section 10] is not a source-selected neutrino update law.

The existing numerical result is consequently restricted: a common independent admission suppression does not repair the fixed six-state candidate. It does not establish that the distinct void-coupling calculation is ineffective, nor that coupling affects only the overall mass scale. No mode-dependent coupling probabilities or neutrino masses derived from them are claimed in Tables 5–7.

11.8 What the action-derived phase calculation adds

The subsequent phase calculation evaluates an action-based alignment potential on a declared finite branch [33]. On a rank-d occupied support, the unit-metric polar-bond defect gives $\frac{1}{2} \| e^{i\theta}P - e^{i\phi}P \|^2 = d[1 - \cos(\theta - \phi)]$. For the archived wheel, the fixed-field compact action is $\Gamma_d = d \sum_e (1 - \cos\phi_e) + \sum_f \left[1 - \cos(B\phi)_f\right]$, where B is the face-edge incidence map. This is an explicit finite-phase continuation of the defect norm, not a proof that the local rate expansion fixes the complete physical probability law.

The full 334-coordinate unit-metric Hessian has invariant weak and abelian phase blocks $S_W = 24I_{12} + B^T B$ and $S_Y = 45I_{12} + B^T B$. Under the declared periodic Haar-reference law proportional to $e^{-\Gamma_d}$ with other fields held fixed, the weak ring coherence is $R = 0.979706030$, while the alignment probability for an illustrative half-window of 0.1 rad is $A = 0.380232430$. The corresponding abelian values are 0.989055644 and 0.500730532. These evaluated functions of the window distinguish coherence from alignment probability; the window is not selected as a neutrino parameter.

Emergent time is checked separately. In the inherited continuity-only temporal extension, both frequency coefficients annihilate the inserted weak and abelian phase subspaces: $K_1 Q_c = K_2 Q_c = 0$. Consequently this calculation supplies static phase stiffness and a conditional spatial measure, but no phase relaxation time or physical neutral-mode phase history. The source's operational transfer and record-clock constructions remain available with their stated carrier and metric conditions [1,2,22]. Their attachment to the neutral completion carrier must be evaluated; a spatial integration kernel is not that time-transfer operator.

The calculation advances the coupling side of the problem without replacing it by an anchoring count. It does not yet identify these record phases with the three physical neutrino interface phases, derive their probability normalization and alignment window, or produce the mixed neutral bilinear. Section 12 therefore evaluates what the existing completion candidate actually returns, while retaining this distinction.

12 Common completion readouts and neutrino to electron mass ratios

12.1 The charged benchmark and the neutral hierarchy are distinct

The next target is the three neutrino-to-electron mass ratios, with both numerator and denominator calculated from one canonical construction [34–36]. This removes a common unit without importing an electron coefficient from another model. It also exposes a necessary correction to the comparison with the successful charged-lepton calculation.

The charged localization law uses $\kappa = 8/3$ and a conditional tau participation factor $1/12$ [37–39]. The archived neutral construction instead uses the address-depth matrix $D = diag(14^{-2}, 14^{-1}, 1)$. A bare two-sided use of this matrix gives the ladder $1{:}\,196{:}\,38416$, not the charged law. The full mixed-tensor descent in section 6.2 remains valid: a common direct-sum parent can contain both sector constructions. That certificate does not identify them as readouts of the same nine-dimensional completion response [13,34].

Table 8. Charged ratios from the distinct constructions. Each row is normalized by its own electron value.

Construction	Muon / electron	Tau / electron
Conditional localization and participation law	207.127249	3575.141436
Retained RRHF benchmark, level 4	207.150579	3576.083677
Shared nine-dimensional completion candidate, level 4	204.167081	42392.252903

The admission tests of section 11 hold the RRHF charged matrix fixed. Their statement that the charged ratios are preserved is correct for that composite implementation. It is not a demonstration that the shared completion candidate in the last row reproduces those ratios. A passive change of field coordinates, with mass, kinetic form and current transformed together, cannot change this comparison.

12.2 Localization and participation in one completion functional

The history construction supplies a quadratic completion problem [2, sections 33–35]. On its positive metric and nonnegative penalty domain,

$$\Gamma(x;y) = \frac{1}{2}(x-y)^{\dagger}W(x-y) + \frac{1}{2}x^{\dagger}\Delta x, \qquad x_{\star} = R_{\Omega}y, \quad R_{\Omega} = (W+\Delta)^{-1}W. \qquad (51)$$

Whitening gives $A = W^{-1/2}\Delta W^{-1/2} \geq 0$ and the positive contraction $S_{\Omega} = (I+A)^{-1}$, with $R_{\Omega} = W^{-1/2}S_{\Omega}W^{1/2}$. Let $\mathcal{R}_{\Omega} = V_H^{\dagger}R_{\Omega}$ include the completion orientation. With physically typed, canonically normalized addresses, the frozen-background Dirac readouts are

$$Y_e = B_L^{\dagger}\mathcal{R}_{\Omega}B_e, \qquad Y_{\nu} = B_L^{\dagger}\mathcal{R}_{\Omega}B_N. \qquad (52)$$

The common weak doublet has one left address B_L. Its charged and neutral readouts can differ through their right addresses and the common geometry. Equation (52) is the linear completion branch; a scalar-dependent background requires the relaxed derivative (3), including any scalar dependence of the addresses and response. The Majorana self-readout has its own charge-conjugation pairing, as in section 2. Neither an identity kinetic metric nor a zero active Majorana term follows from (52).

The conditional charged law supplies an exact algebraic test of this construction [35,37]. Set $G = diag(0,1,2)$, $P_\tau = diag(0,0,1)$ and $C_{ch} = I - (11/12)P_\tau$. Its positive canonical mass magnitude is

$$|M_e| = m_0 C_{ch} e^{2\kappa G}, \qquad \kappa = 8/3. \qquad (53)$$

Choosing the reference $m_\Omega = m_0 e^{4\kappa}$, the dimensionless response has the exact representation

$$S_e = C_{ch} e^{2\kappa(G-2I)} = (I + A_{loc} + A_{maint})^{-1}, \qquad (54)$$

$$A_{loc} = e^{2\kappa(2I-G)} - I, \qquad A_{maint} = 11P_\tau. \qquad (55)$$

Both penalties are nonnegative, and $A_{loc}P_\tau = 0$. Therefore $I + A_{loc} + A_{maint} = (I + A_{loc})(I + A_{maint})$, proving (54). The localization entries are approximately (42900.697233,206.127249,0); the maintenance entries are (0,0,11). Their single variational response is (0.00002330910114,0.004827949994,1/12), which gives the first row of Table 8.

This derives a compatible resolvent from the conditional charged law. It does not derive the penalties from the physical stationary history. In particular, the charged source retains a stiffness-identification premise for κ and explicitly leaves the one-channel numerator in the tau factor conditional [37, sections 11–12]. The value 11 in (55) is $(1/12)^{-1} - 1$, not an independently evaluated maintenance rank.

Even on an aligned same-address branch, sector reduction requires relaxation. If E selects the charged carrier, decompose $K = I + A$ into that carrier and its complement. Then

$$E^\dagger S_\Omega E = [K_{ee} - K_{e\perp}K_{\perp\perp}^{-1}K_{\perp e}]^{-1}. \qquad (56)$$

The source-selected Schur penalty, rather than an arbitrarily chosen diagonal block, must match (55). General left-right readouts additionally require their actual polar frames. Nor can a response be substituted for an effective-action Hessian: minimizing (51) gives the kernel $W - W(W + \Delta)^{-1}W$, whose whitened form is $I - S_\Omega$. The physical insertion and any direct term determine which quantity enters a fermion bilinear.

The participation distinction can be stated directly. For a whitened physical constraint map F, the allowed-space projector is

$$Q = I - F^\dagger\left(FF^\dagger\right)^+F, \qquad \lim_{a\to\infty}\left(I + aF^\dagger F\right)^{-1} = Q. \qquad (57)$$

Eleven independent constraints in twelve dimensions imply $TrQ/12 = 1/12$. They do not fix the coherent amplitude $u_L^\dagger Q u_R$ to that value. The trace average requires an isotropic averaging rule; the mass readout requires physical left and right states. This is why a channel count alone cannot supply the neutral coupling or complete the charged participation derivation.

As a diagnostic, inserting $11q_\tau q_\tau^\dagger$ into the existing nine-dimensional penalty gives a Sherman–Morrison suppression $\left[1 + 11q_\tau^\dagger S q_\tau\right]^{-1} = 0.297573254$ at level 4, rather than 1/12 [35]. Its charged ratios become $1{:}\,204.145678{:}\,12616.1360$, and its active gap ratio is 7.44099×10^{-5}. The suppression differs from 1/12 by a factor 3.57, so the failure is not marginal. This tested insertion fails; it is not selected as a replacement model.

12.3 Electron normalization of the full neutral return

The numerical candidate uses the positive response $R = (w_H I + \Delta)^{-1} w_H I$, its oriented return $O = VR$, and normalized inclusions $E_s = I_3 \otimes v_s$. Here V denotes the archive's nine-dimensional unitary orientation, distinct from the charged-current matrix of (6). Its prescribed blocks are [36]

$$\begin{aligned} Y_e &= DE_L^\dagger O E_e D, & Y_\nu &= DE_L^\dagger O E_N D, \\ \widehat{M}_R &= sym[DE_N^T sym(O) E_N D], & \epsilon &= \frac{\sqrt{TrR/9}}{7\sqrt{2}}. \end{aligned} \tag{58}$$

With the archive's identity kinetic pairing and $M_L = 0$, the shared mass convention is

$$\widehat{M}_e = \frac{M_e}{\Lambda} = \epsilon Y_e, \qquad \widehat{\mathcal{M}}_\nu = \frac{\mathcal{M}_\nu}{\Lambda} = \begin{pmatrix} 0 & \epsilon Y_\nu \\ \epsilon Y_\nu^T & \widehat{M}_R \end{pmatrix}. \tag{59}$$

Let $c_e = \sigma_{\min}(\widehat{M}_e)$ and let μ_j be all six positive Takagi values of $\widehat{\mathcal{M}}_\nu$. The requested observable is exactly

$$\mathcal{M}_\nu^{(e)} = \frac{\widehat{\mathcal{M}}_\nu}{c_e}, \qquad r_j = \frac{m_j}{m_e} = \frac{\mu_j}{c_e}. \tag{60}$$

The common Λ, including a common clock calibration, cancels. The relative Dirac/Majorana coefficient ϵ does not cancel: in a controlled seesaw limit $r_i \simeq \epsilon\, \sigma_i\big(Y_\nu \widehat{M}_R^{-1} Y_\nu^T\big)/\sigma_{\min}(Y_e)$. The primary results use the full matrix (59). Substituting the successful RRHF electron coefficient into this denominator would require a derived relative normalization between the two constructions.

Table 9. The three most active candidate states, ordered by mass, and their common electron normalization.

Quantity	Level 3	Level 4
$c_e = m_e/\Lambda$	$2.602956976 \times 10^{-7}$	$2.598579243 \times 10^{-7}$
m_{ν_a}/m_e	0.0544970028	0.0544079021
m_{ν_b}/m_e	11.1045223	11.0846740
m_{ν_c}/m_e	2347.27646	2342.09765
Active weight of ν_a	0.997089	0.997097
Active weight of ν_b	0.985818	0.985939
Active weight of ν_c	0.992819	0.992822
Squared-gap ratio r_Δ	$2.238003782 \times 10^{-5}$	$2.239883297 \times 10^{-5}$

At level 4, $\epsilon = 0.0521051847514$. The six Takagi values, in ascending order and in common Λ units, are

$$\begin{aligned} \mu_{1,2,3} &= (1.4138324512 \times 10^{-8},\ 2.8804403814 \times 10^{-6},\ 4.8930232231 \times 10^{-6}), \\ \mu_{4,5,6} &= (6.0861263418 \times 10^{-4},\ 9.8977225478 \times 10^{-4},\ 0.20309588564). \end{aligned} \tag{61}$$

The most active states are 1, 2 and 4; the third-lightest state is predominantly sterile. Labels a, b, c identify this candidate triplet and do not assert a derived association with the observed oscillation labels. The current uses the shared candidate's charged left frame; its largest electron probability is 0.997076537. This differs from the RRHF-frame value 0.996875622 used in section 11.5 even though the retained neutral spectrum is the same.

Table 9 is a failed candidate result, not a physical neutrino mass prediction. Electron normalization makes the failure explicit. Using the measured electron mass energy, 0.51099895069 MeV [40], solely for an illustrative conversion gives $(27.8024\ keV, 5.66426\ MeV, 1.19681\ GeV)$. These values use the candidate's stated mass convention without an additional pole/running matching derivation. No neutrino observation selects a matrix entry. The ratio changes between levels are approximately $-0.1635\%, -0.1787\%, -0.2206\%$; this checks two-level stability, not continuum convergence.

12.4 Completion strength and coherent coupling are different observables

The positive response permits an exact decomposition of the candidate. Set $\tilde{L} = V^{\dagger}E_L$. For $f = e, N$, define

$$\begin{aligned} G_L &= \tilde{L}^{\dagger}R\tilde{L}, \qquad G_f = E_f^{\dagger}RE_f, \\ C_f &= G_L^{-1/2}\tilde{L}^{\dagger}RE_f G_f^{-1/2}. \end{aligned} \tag{62}$$

The columns of $R^{1/2}\tilde{L}G_L^{-1/2}$ and $R^{1/2}E_f G_f^{-1/2}$ are orthonormal. Hence the singular values of C_f are principal overlap amplitudes in [0,1]. The exact reconstruction is

$$Y_f = D\, G_L^{1/2} C_f G_f^{1/2}\, D. \tag{63}$$

The Gram factors measure completion-leg strength; C_f measures coherent overlap. Neither a Gram eigenvalue nor an overlap amplitude is thereby identified with the microscopic anchoring count or phase-window probability in (46)–(47). Those identifications require the physical source and history map.

Table 10. Singular values of the coherent overlap, in decreasing order, at level 4.

Readout	First	Second	Third
Charged C_e	0.877622878	0.845482752	0.836788219
Neutral C_N	0.875019983	0.840943931	0.831619641

The reconstruction residual is below 1.5×10^{-15}. The right-Gram spectra are also comparable, and $\| Y_\nu \|_F / \| Y_e \|_F = 1.010349436$. Thus the finite candidate does not implement substantially weaker neutral coupling: its charged and neutral completion subspaces overlap almost equally strongly. This is a statement about the evaluated operator, not a conclusion that physical neutrinos must couple this way.

12.5 Interaction suppression is not a neutral-address rescaling

Two diagnostic changes distinguish the proposed weak coupling from an address-amplitude change. First, hold the charged matrix, sterile block and kinetic pairing fixed, and change only $M_D \mapsto tM_D$. With $M_L = 0$, its deep seesaw ratios are approximately $t^2(0.0545667587, 11.1699207, 2347.25827)$. This reduces the scale but retains the leading relative spectrum and active eigenframe. Its limiting gap ratio, $2.26447777 \times 10^{-5}$, remains far from the previously supplied comparator 0.0298.

Second, a fixed-kinetic change of the neutral address $B_N \mapsto tB_N$ changes both readouts: $M_D \mapsto tM_D$ and $M_R \mapsto t^2 M_R$. The leading seesaw term obeys

$$-(tM_D)(t^2M_R)^{-1}(tM_D)^T = -M_DM_R^{-1}M_D^T. \qquad (64)$$

Shrinking that address therefore does not shrink this leading active operator. It lowers the sterile scale and eventually invalidates the separated approximation. The full calculation resolves the resulting admixture.

Table 11. Diagnostic probes at level 4. No physical value of t is selected.

Change	t	Three most active mass ratios to electron	Minimum active weight
Dirac interaction only	0.01	$5.45667428 \times 10^{-6}$; 0.00111699151; 0.234716709	0.9999607
Neutral address in both blocks	0.1	0.0441205; 9.04387; 1893.99	0.8364
Neutral address in both blocks	0.01	0.00919495; 1.87948; 391.118	0.5453

The Dirac-only probe has $r_\Delta = 2.26465142 \times 10^{-5}$. Active purity is not monotone throughout this family: at $t = 0.1$ an avoided crossing lowers one active weight to about 0.954. The two address probes no longer provide a clean three-state active sector. A passive field-coordinate rescaling is a third operation; transforming the kinetic metric and current consistently preserves the physical spectrum. None of these diagnostics assigns a physical coupling from a desired mass.

12.6 Physical addresses and the freedom left by charged masses

The source already specifies how to construct physical addresses. MCHFY distinguishes sparse occupancy selectors from their address-to-closure realization [12, section 22]:

$$S_{fX}^\dagger S_{fX} = I_3, \qquad B_{fX} = K_{cl}^{1/2} J_\Omega S_{fX} K_{fX}^{-1/2}. \qquad (65)$$

Representation support alone does not numerically supply J_Ω. MASD gives the typed terminal-support, projection and polar-decomposition rules [1, sections 29–35 and Appendix D]. The required step is their physical numerical realization on one selected background, including the neutral-extension premise and pairing, rather than another freely chosen normalization.

The current archive instead solves $X = H^{-1}Q$ and forms $J = X\left(X^\dagger W_{FE} X\right)^{-1/2}$ before normalizing terminal rows j_s. Interpreting the unnormalized inclusion as $I_3 \otimes j_s$ gives the Gram $\| j_s \|^2\, I_3$. At level 4 the row norms are (0.578057319,0.658742527,0.689581964) for L, e, N. This particular raw-row interpretation has scalar polar radii; it does not derive the separately prescribed generation factor D. It is not a theorem that the full physical map (65) has no hierarchy, and it does not authorize deleting D by hand.

Nor do charged masses determine the neutral readout merely because a positive common completion exists. For the charged response S_e in (54), the family

$$S_\Omega(t,U) = \begin{pmatrix} S_e & t\sqrt{S_e}U\sqrt{S_e} \\ t\sqrt{S_e}U^\dagger\sqrt{S_e} & S_e \end{pmatrix}, \quad 0 < t < 1, \quad U^\dagger U = I, \qquad (66)$$

is positive and contractive, with upper norm bound $(1+t)/12$. Its penalty $S_\Omega^{-1} - I$ is nonnegative. Fix the charged left and right embeddings to the first block and the neutral right embedding to the second. The charged response and total trace stay fixed, while the cross-readout is $t\sqrt{S_e}U\sqrt{S_e}$. With the same symmetric sterile block S_e on the real-U subfamily, the neutral spectrum changes [35]. This is a mathematical nonuniqueness witness, not a source-selected physical history.

The executed source derivatives reach the same practical conclusion for the current candidate. If Z spans directions preserving all three charged masses in the retained 26-coordinate source family, form

$$\left(J_{\log r}\right)_{ik} = \partial_k \log\mu_i - \partial_k \log c_e. \qquad (67)$$

The singular values of $J_{\log r}Z$ are (1.40630328,0.04095943,0.00460401) at level 4 and (1.41113569,0.04092452,0.00470193) at level 3. Both have rank three at threshold 10^{-8}, with charged-null residuals below 10^{-15} [36]. Thus all three ratios can change locally while the charged masses remain fixed. This is a response rank in the declared candidate family, not a count of physical modes or proof that its directions are physically selectable.

12.7 The next source selected return

The next target is the neutral cross-readout relative to its sterile self-readout, evaluated with the charged operator from the same physical construction. Concretely, the terminal supports, J_Ω, kinetic metrics, completion orientation and selected history must determine both $B_L^\dagger \mathcal{R}_\Omega B_N$ and the charge-conjugation-paired singlet bilinear. The irreducible active term and any frequency dependence must be included at the same action level. This preserves the existing constructive source rules while identifying which numerical realization remains to be returned.

Where void coupling is expressed through phase alignment, its physical phase history must be attached to these readouts; the static phase calculation and conditional probabilities in section 11.8 cannot select their coefficients by themselves. Likewise the localization and participation response must emerge with its physical left and right states, rather than be imposed after a neutral matrix has been computed.

The resulting calculation must output the charged spectrum, all neutral poles and current weights, and the three ratios (60) before comparison with neutrino observations. A common clock is retained consistently and cancels from those ratios; mode-dependent dynamics and relative scalar normalization remain. At present Table 9 is the evaluated candidate return. A source-selected physical neutrino-to-electron ratio triple has not yet been obtained.

12.8 Audit of the shared candidate: one origin of both failures

The candidate of section 12.3 admits a structural diagnosis. Throughout this subsection the charged frame is the RRHF frame, $U_{eL} = I$ with the address bridge $J_{e\nu} = I$ and the active projector $P_A = (I_3 \;\; 0)$, so that the electron row is read directly from the first three rows of the Takagi matrix; in that frame the largest electron probability is 0.996875622 and the disappearance ceiling $4p(1-p)$ is 1.24585%. In the shared candidate's own charged left frame the corresponding values are 0.997076537 and 1.16597% (section 12.3). The two frames give the same six masses and the same gap ratio.

Strip the address-depth factor from both dimensionless level-4 blocks,

$$Y = D^{-1}(\epsilon Y_\nu)D^{-1}, \qquad Y_R = D^{-1}\widehat{M}_R D^{-1}, \qquad D = diag(14^{-2}, 14^{-1}, 1). \qquad (68)$$

Both stripped cores are circulant as complex matrices, to the precision of the printed inputs. Writing P for the cyclic shift, the residuals $\| Y - PYP^T \|_{\max}$ and $\| Y_R - PY_RP^T \|_{\max}$ are 7.0×10^{-18} and 1.1×10^{-16}, relative to largest entries 0.01113 and 0.2024. The three cyclic classes of entries are, for Y : diagonal $0.01097360 - 0.00183931i$, forward off-diagonal $0.00267273 + 0.00017647i$, backward off-diagonal $0.00257263 - 0.00066434i$; and for Y_R : diagonal $0.19235749 - 0.06308391i$, both off-diagonals $0.04198669 - 0.01167559i$. This is numerical agreement at machine precision, not an algebraic identity proved from the construction; the construction makes it expected, since the magnetic phase $2\pi/21$ enters only through the generation triangle, the address operator commutes with that triangle, and the embeddings preserve it. Within this implementation the only C_3-breaking ingredient in generation space is therefore D. The same D produces the candidate's charged block, whose ratios are $1: 204.167: 42392.253$ (Table 8); those are not the successful RRHF ratios $1: 207: 3576$, and nothing here says the candidate reproduces the charged spectrum.

In the seesaw regime the light operator is DKD with $K = YY_R^{-1}Y^T$ circulant and $K_{off}/K_{diag} = 0.261$. Both failures then follow by hand. The gap ratio is $14^{-4} = 2.60 \times 10^{-5}$ up to the circulant correction (the evaluated value is 2.24×10^{-5}), and the active-sector leakage of the electron out of the lightest pole is $(0.261/14)^2 = 3.5 \times 10^{-4}$, to be compared with the evaluated 2.2×10^{-4} into state 2. This estimate is a three-active seesaw statement and is not the total leakage of the full six-state calculation, which is $1 - 0.996875622 = 0.003124378$; of that, about 0.002900 sits in the predominantly sterile state 3 and is active–sterile admixture, a separate effect that the DKD reduction does not describe. In this candidate the two discrepancies are one fact: D acts on the active neutrino field exactly as on the charged lepton.

Reassigning D within the candidate does not repair it. For $M_D = D^{a_1}YD^{a_2}$ and $M_R = D^bY_RD^b$ with exponents in $\{0,1\}$, the light operator is

$$M_{light} = -D^{a_1}\left(YD^{a_2-b}Y_R^{-1}D^{a_2-b}Y^T\right)D^{a_1}. \qquad (69)$$

The outer factor D^{a_1} and the inner factors D^{a_2-b} both act; the inner ones cancel only when $a_2 = b$, and otherwise change both the hierarchy and, through the sandwich with Y, the frame. All eight placements were evaluated on the full six-state Takagi problem, selecting the three states of largest total active weight before any ratio is read; the inputs and outputs are tabulated in Appendix A. Every placement with $a_1 = 1$ keeps $p_{e1} \geq 0.995$ and a gap ratio of at most approximately 5.09×10^{-3}. Every placement with $a_1 = 0$ and $b = 1$ floods the light triplet with sterile weight (electron row summing to 0.52–0.55 over its most active triplet), because the diagonal entry of Y, of magnitude 0.0111, then exceeds the light sterile eigenvalue $14^{-4} \times 0.2$. The fully symmetric placement gives a degenerate pair, so one of the two required splittings is absent. In an electron-aligned basis of the degenerate subspace its electron row on the light triplet is $(0,\ 0.664761,\ 0.332276)$; columns inside a degenerate subspace depend on the basis chosen and do not establish physical mixing angles. What is basis-independent is that the non-degenerate state carries equal normalized flavour proportions, 1/3 each, while its physical current probabilities are 0.332276 per flavour, the shortfall from 1/3 being the small sterile admixture. Normalized flavour proportions and physical current probabilities are therefore kept distinct here and in Table A1. The per-generation light mass ratio required by the benchmark, 0.1726, would need an effective exponent $14^{-2/3} = 0.1722$; this is recorded as a numerical coincidence and not used.

The audit therefore sharpens section 12.6 for this implementation. In the tested candidate, and in its eight D-placements, the active neutrino cannot carry the address ladder that its charged block carries and still meet the data, and no other generation-space operator is present in the neutral sector to lift the degeneracy that removing the ladder leaves. These are conclusions about the specific construction and the placements tested; they do not exhaust the source-derived mass, kinetic and current structures that a different neutral attachment could return. Any such attachment must still respect the common weak-doublet construction, so the source question it raises is whether the accessibility weight attaches to the $SU(2)$ doublet as a whole or only to the charged attachment. If to the doublet, this finite realization cannot meet the data and must be replaced; if not, the neutral attachment that replaces it is what the remaining sections seek.

13 Nonlinear single-fold attachment and the generation law of the neutral mass response

13.1 The source-derived nonlinear fold response

The next calculation should not begin by assigning three neutral coefficients. The original nonlinear EPMD executable already supplies a more restrictive source. Retain the two real E_1 fold coordinates $x = (x_1, x_2)$. In the normalized diagonal source convention, the reduced cubic action contains the E_2 coupling [42]

$$\Gamma_{eff}^{(3)} = \frac{1}{2}[y_1(x_1^2 - x_2^2) + 2y_2x_1x_2], \qquad g_{source} = \frac{1}{2}. \qquad (70)$$

Define $Q(x) = (x_1^2 - x_2^2,\ 2x_1x_2)$. The action-gradient force on the E_2 pair is therefore

$$F_{E_2}(x) = -\frac{1}{2}Q(x). \qquad (71)$$

This coefficient is not fitted to a neutrino observable. The archived three-copy normalization has already been pulled back to the normalized diagonal, so dividing the coefficient by $\sqrt{3}$ again would double count the source normalization. The absolute physical fold amplitude and the identification of this finite source coordinate with the physical neutral coordinate remain separate questions [42].

The complementary field directions are relaxed rather than held affine. If V is the linearly relaxed lift, A spans the auxiliary complement and H is the native Hessian, the minimizing second attachment B is fixed by stationarity:

$$B_{ij} = -A(A^THA)^{-1}A^T\,\Gamma^{(3)}[\cdot, V_i, V_j]. \qquad (72)$$

The full second derivative of any field residual therefore contains both its direct curvature and the relaxed attachment term. On the finite source used in [42], the latter is numerically substantial and cannot be set to zero without changing the calculation.

After retaining only the E_1 pair and relaxing all 185 complementary physical directions, the induced E_2 displacement is

$$y_\star(x) = -\chi Q(x) + O(\| x \|^3), \qquad \chi = \frac{g_{source}}{\kappa_{E_2}}. \qquad (73)$$

For the unit residual metric $\chi = 0.096153846154$, so $\kappa_{E_2} = 26/5$ exactly; for the archived history-diagonal metric $\chi = 0.221883660298$. A typed-history candidate gives $\chi = 0.170990002238$ but is not selected here as the physical branch. The same second attachment also generates an A_1 singlet density displacement proportional to $\| x \|^2$. In the archived separable scalar action the local Higgs completion coordinates remain exactly unchanged; the nonzero response lies in density and real current. This is a result about that source action, not a theorem that the complete representation-resolved neutral completion must have zero Higgs response [42].

13.2 Generation refinement fixes the available scaling laws

The nonlinear source becomes predictive across generations only after a physical generation map is specified. Decompose the retained E_1 coordinate into a generation-even component u and a reversal-odd component v. The VERSF half-lazy construction gives an exact one-half eigenvalue on the odd readout when its binary-admissibility, cost-symmetry and log-access intertwining conditions hold. If the physical generation step descends to the same operator on this neutral fold coordinate, then [41]

$$x^{(g)} = \left(u,\ 2^{-(g-1)}v\right), \qquad g = 1,2,3. \qquad (74)$$

This is a conditional physical bridge, not an automatic consequence of writing a two-component fold. It can be rejected if the neutral attachment does not satisfy the required intertwining. Conditional on the bridge, substitution into the action-derived Q gives

$$Q\left(x^{(g)}\right) = \left(u^2 - 4^{-(g-1)}v^2,\ 2^{-(g-1)} \cdot 2uv\right), \qquad (75)$$

so the relaxed E_2 background obeys

$$y_\star^{(g)} = -\chi\left(u^2 - 4^{-(g-1)}v^2,\ 2^{-(g-1)} \cdot 2uv\right) + O(\| x \|^3). \qquad (76)$$

The generation-dependent part of the first E_2 component scales as 1, 1/4, 1/16 and the second component as 1, 1/2, 1/4. The A_1 singlet displacement carries $\| x^{(g)} \|^2 = u^2 + 4^{-(g-1)}v^2$, so its non-common generation dependence follows the same quarter-sequence. At this order the nonlinear fold supplies only a common component, a half-sequence and a quarter-sequence. It does not license three independent generation weights. The restriction is on the form, not yet on the values: the three sequences $(1,1,1)$, $(1, 1/2, 1/4)$ and $(1, 1/4, 1/16)$ are linearly independent and therefore span every triple of generation values when their coefficients are unrestricted. Predictive content arises only when the coefficient matrices are independently determined by the pullback of section 13.3.

13.3 Pullback to the canonically normalized neutral operator

The preceding statement is a result in the physical fold and background coordinates. It becomes a mass statement only after the neutral bilinear and kinetic functional are pulled back along the same relaxed attachment. For a genuine common attachment $\Xi(x)$, the required second derivatives are [42]

$$\partial_i \partial_j (N_\nu \circ \Xi) = D^2 N_\nu[V_i, V_j] + DN_\nu[B_{ij}], \qquad \partial_i \partial_j (Z_\nu \circ \Xi) = D^2 Z_\nu[V_i, V_j] + DZ_\nu[B_{ij}]. \qquad (77)$$

The first term is the direct curvature of the neutral readout; the second is the response of that readout to the relaxed second attachment. Both are required. Replacing either derivative by an identity map would introduce a new assumption and would incorrectly identify the finite source E_2 coordinate with a neutral matrix coordinate.

Section 7 gives the canonical fixed-background mass and kinetic vertices. In a local Hermitian whitening frame, define the canonical response to a physical background mode a by

$$C_a = G_a - \frac{1}{2}(J_a^T M_N + M_N J_a). \qquad (78)$$

Once the representation-resolved pullback assigns the fold's half-sequence and quarter-sequence to their canonical neutral response matrices A and B, and all generation-independent second-order pieces, including the u^2 parts, are collected into a common matrix $\overline{M}_N$, the neutral matrix has the constrained form

$$M_N^{(g)} = \overline{M}_N + 2^{-(g-1)}A + 4^{-(g-1)}B + O(\| x \|^3), \qquad (79)$$

that is, $M_N^{(1)} = \overline{M}_N + A + B$, $M_N^{(2)} = \overline{M}_N + 1/2\,A + 1/4\,B$ and $M_N^{(3)} = \overline{M}_N + 1/4\,A + 1/16\,B$. The matrices A and B are common response matrices returned by one physical neutral pullback; they are not three separately adjustable generation coefficients. Equation (79) indexes a response by generation; the assembly of the three $M_N^{(g)}$ into the single physical neutral mass operator on the three-generation carrier (whether as diagonal blocks, as a generation-labelled source entering one resolvent, or otherwise) is part of the representation-resolved pullback and is not fixed by (79) itself. Section 13.5 uses one particular diagnostic assembly for that reason. The question that follows is whether a concrete existing neutral response can realize this structure without using neutrino data to choose its coefficients.

13.4 What is now derived and what remains open

The primitive-source side has advanced materially. The nonlinear E_1^2-to-E_2 driving coefficient is explicitly returned; the full second-order relaxed attachment has been calculated rather than suppressed; the singlet density response has been retained; and, conditional on the generation-intertwining premise, the generation algebra admits only the half- and quarter-sequences above. These results use no neutrino masses, splittings or PMNS angles as inputs [41,42].

The physical mass certificate remains open for a narrower reason. The calculation still requires the representation-resolved derivatives of the neutral bilinear and kinetic functional, including the physical sector addresses, support and polar readout, positive completion factors and their derivatives. The physical fold amplitude, the source-selected value of κ_{E_2}, the canonical neutral residue and the branch selection must travel with that return. The local source result $g_{source} = 1/2$ cannot simply be divided by an unrelated neutral restoring coefficient and called a mass coupling.

Before that return exists, a falsifiable diagnostic is possible: use the simplest same-wheel endpoint identification already present in the finite neutral model, propagate the derived half- and quarter-sequence sources through its existing completion resolvent, and test the resulting spectrum blind. Passing it would not prove the physical pullback; failing it can reject the shortcut.

13.5 First direct contraction into the existing neutral response

To test whether the generation law (79) has enough structure to generate a nontrivial spectrum, take the simplest direct same-wheel endpoint identification as a diagnostic bridge. This is not declared to be the physical representation-resolved pullback. Let the two generation-dependent source shapes be

$$S_1 = diag(1,\ 1/2\,,\ 1/4), \qquad S_2 = diag(1,\ 1/4\,,\ 1/16). \qquad (80)$$

Applying the diagnostic seven-address completion response of [42] to these two sources gives the dimensionless candidate response matrices below; Appendix B specifies the inherited inputs, the full response map and its independent stationary-solution check:

$$A_0 \approx \begin{pmatrix} 0.620280 & 0.090793 & 0.079326 \\ 0.090793 & 0.321134 & 0.056391 \\ 0.079326 & 0.056391 & 0.171562 \end{pmatrix}, \qquad B_0 \approx \begin{pmatrix} 0.614782 & 0.073828 & 0.065228 \\ 0.073828 & 0.166065 & 0.030826 \\ 0.065228 & 0.030826 & 0.053885 \end{pmatrix}. \tag{81}$$

Both matrices are full rank and their commutator is nonzero, $\|[A_0, B_0]\|_F = 0.0367$, so the two fold-generated responses are not simultaneously diagonalizable. Even this deliberately minimal bridge has enough algebraic structure to generate three distinct masses and a nontrivial neutral-frame rotation; it is not merely a common scalar rescaling.

13.6 Blind equal-winding spectrum test

The selected self-closing wheel mode is a complex first harmonic. The cleanest symmetry benchmark therefore assigns equal norm to its reversal-even and reversal-odd real components, $u = v = 1/\sqrt{2}$. No neutrino mass, splitting or mixing angle is used in this choice. In the direct diagnostic bridge the generation source is

$$s_g = \left(u + i\,2^{-(g-1)}v\right)^2, \qquad g = 1,2,3, \tag{82}$$

mapped into the complex neutral candidate by the explicit contraction (B10) of Appendix B, which includes the generation-independent part of this diagnostic source and prints the complete complex matrix. The resulting singular values, up to the common unknown dimensional scale, are

$$m_1 : m_2 : m_3 = 0.286635 : 0.382294 : 0.687699 = 1 : 1.33373 : 2.39922, \tag{83}$$

with dimensionless mass-squared-gap ratio $(m_2^2 - m_1^2)/(m_3^2 - m_1^2) = 0.1638$. The oscillation benchmark of section 11.5 is 0.0298 for normal ordering. The equal-winding E_2-only shortcut is therefore quantitatively wrong by a factor of five and a half. This is a useful falsification: the nonlinear fold response is nontrivial, but the simplest direct endpoint identification does not reproduce the physical hierarchy.

13.7 Orientation diagnostic, mixing check and no-fit firewall

As a diagnostic only, vary the orientation $\varphi = \arctan(v/u)$ while keeping the same derived response matrices. A value near $\varphi \approx 13.28^\circ$ gives a relative spectrum 1: 1.0320: 1.7889 and a mass-squared-gap ratio 0.02956. This shows that the fold-response family contains enough structure to reach the observed hierarchy.

That value is not a prediction. The nonlinear source calculation explicitly leaves the absolute fold direction unselected; choosing φ because it reproduces the oscillation ratio would move the fit into the fold orientation and is forbidden as a derivation step. The same diagnostic orientation also fails to reproduce a PMNS-like frame: the best column assignment gives approximately $\theta_{12} \approx 37.4^\circ$, $\theta_{23} \approx 41.7^\circ$ and $\theta_{13} \approx 3.5^\circ$. Matching the gap ratio alone is not enough. The quoted diagnostic angles use an identity charged-lepton frame and current bridge, with columns (2,3,1) relative to ascending singular values; this permutation compares frame shapes only and does not establish a physical mass-column identification. Appendix B reproduces the calculation.

The result is therefore two-sided. The nonlinear one-fold mechanism has sufficient rank and noncommuting structure to generate three distinct masses and flavour rotation, and the simplest E_2-only equal-winding realization is rejected while an angle chosen after looking at the data is inadmissible. The remainder of this section follows the missing common term, kinetic normalization and fold selection back into the source instead of introducing another spectrum parameter.

13.8 Source audit of the common/background response

The nonzero A_1 singlet attachment is a real density response, but density is not itself a neutral mass baseline. In the archived separable finite action the complete second attachment gives zero response in the local Higgs/completion coordinates h and X_H. The same fold also generates a real current, but the projection of that current onto the three established divergence-free neutral cycle modes is zero to numerical precision: the relative projection is below 1.8×10^{-15}. Thus the two most immediate routes, A_1 to local completion and fold current to the known neutral cycle response, do not generate the missing $\overline{M}_N$ on the executed branch [42,47].

The primitive terminal address maps do not rescue this result. In the Master-Action construction the raw representation-typed addresses are frozen before response is evaluated; their primitive density derivatives therefore vanish by construction. A derived scalar-compatible support could in principle move, but the explicit mixed density–Higgs attachment evaluated on the archived branch also vanishes. These statements localize the missing coupling; they are not a theorem that the full representation-resolved neutral readout is zero.

13.9 The retained fold action does not select an orientation

For $x = (x_1, x_2)$ in the retained E_1 fold plane the nonlinear source is $Q(x) = (x_1^2 - x_2^2,\ 2x_1x_2)$, with invariant norm

$$\parallel Q(x) \parallel^2 = (x_1^2 - x_2^2)^2 + 4x_1^2x_2^2 = (x_1^2 + x_2^2)^2. \qquad (84)$$

After the E_2 response is relaxed, the induced quartic contribution therefore depends on the fold radius but not on the angle $\varphi = \arctan(x_2/x_1)$. The A_1 singlet term is also proportional to $\parallel x \parallel^2$. Hence, at the order retained in the present source action,

$$\frac{\partial \Gamma_{eff}}{\partial \varphi} = 0. \qquad (85)$$

The 13.28° orientation of section 13.7 is consequently a reachability diagnostic only. It cannot be derived from the current cubic/quartic fold action. Physical orientation selection must enter through a higher-order anisotropy or another source-coupled sector. This closes the possibility of promoting the diagnostic angle by a stationary-point argument that is not present at this order. The same flatness is the fold-plane counterpart of the isotropy noted for the address ladder in section 12.8: an invariant quadratic on a rotation doublet carries no angle.

13.10 Faithful current-history typing excludes the reduced scalar shortcut

A one-bit reduced history metric contains a nonzero scalar J–H correlation. That reduced coefficient is useful diagnostically, but faithful representation typing changes the conclusion. The generation-sensitive current-cycle doublet transforms as E_2 under the wheel D_6 symmetry, whereas the local H-completion carrier is spatially trivial. Therefore

$$Hom_{D_6}\left(E_2,\ A_1^{\oplus 3}\right) = 0, \qquad (86)$$

so the reduced scalar J–H overlap cannot be lifted to the physical E_2 fold mode. The only surviving linear E_2 cross-defect portal in the present faithful census is the scalar current to primitive Abelian-closure line. The direct common neutral term cannot therefore be obtained by promoting the reduced W_{HJ} number [45,46].

The same typing identifies the correct size of the missing history response. The colour/weak/hypercharge factor quotient is positive rank three and the current has two real E_2 quadratures, so the physical current-history score must have rank six. A single scalar E_2 bath pair has rank at most two and cannot carry the full response.

13.11 Positive current-history realization and the one remaining source gain

The absence of the retained source leg is not a positivity obstruction. RCCSJ-1 constructs an explicit positive continuous-time history generator by exponentially tilting the canonical wheel jump rates with the representation-resolved current-cycle score. With one common real gain g_{cyc} its Fisher-rate metric is

$$W_{full}\left(g_{cyc}\right) = g_{cyc}^2\,\frac{2}{31}\,H_{CWY} \otimes I_2, \qquad W_{m=2/4}\left(g_{cyc}\right) = g_{cyc}^2\,\frac{1}{62}\,H_{CWY} \otimes I_2, \qquad H_{CWY} = \frac{1}{49}\begin{pmatrix} 36 & 18 & 36 \\ 18 & 26 & 26 \\ 36 & 26 & 47 \end{pmatrix}. \qquad (87)$$

The factor Gram H_{CWY} is supplied by the source-generated current stiffness of RCHS-1 [48]; it is positive definite, with eigenvalues 0.0787, 0.2527 and 1.8931. Positivity and D_6 covariance alone do not fix the relative colour, weak and hypercharge weights; given H_{CWY} from [48] and the terminal representation census, they fix the local Fisher metric up to the one common strength g_{cyc}. This is uniqueness of the local metric, not of the complete history law, whose higher-order structure is not addressed here. In the canonical current-cycle normalization it is related to the direct current coefficient by

$$\lambda_J = 0.303446843055259\ g_{cyc}^2. \qquad (88)$$

The completed retained-order source does not determine that gain [45]. The explicit HCMR/RRHF current-cycle execution returns a $0_{6\times 6}$ pullback where the required target has rank six. Its later XI-3 composition remains radially blind because the current enters through the normalized ray $\rho_J = |J\rangle\langle J|/\langle J|J\rangle$. This is a source-level no-return, not $\lambda_J = 0$. No further differentiation or Schur reduction of the same retained blocks can create the missing magnitude-bearing leg [44,46].

13.12 Exact release condition and paper boundary

The residual-level current coordinate is already normalized: $P_{cyc}\delta D_J = P_{cyc}\delta J$. What remains is the map from that fixed physical current tangent into the physical history law. Because the primitive gauge source is graded, the correct object is not a new six-column linear map at the flat source. Colour and weak first appear as quadratic Casimir jets, while hypercharge, record, completion and CAR/Fock occupy different differential orders. At a nonzero current background J_0 those higher jets acquire ordinary first variations by background contraction [45].

The release object for the next calculation is therefore

$$T_{cyc}^{(n)} = P_{E_2}^{hist}\, D\Phi_{gr}^{(n)}(J_0)\, P_{E_2\otimes CWY}^{(D_J)}, \qquad (89)$$

or, equivalently on a completed common metric/history construction, the corresponding representation-resolved projection of $D_J W_{prim}(J_0)$. A physical pass requires rank six, D_6 covariance, regulator stability, positive history evolution and a source-fixed coefficient before any neutrino or Standard-Model target is inspected. Once the raw transfer derivative is returned, the Perron/Doob normalization and terminal score follow uniquely; there is no second current-score scale to choose [43,45,47]. This is one identified upstream bottleneck, not the last unknown. Returning $T_{cyc}^{(n)}$ fixes g_{cyc} and hence the history response; it does not by itself determine the physical fold amplitude, the fold orientation, the source-selected κ_{E_2}, the physical sector addresses, the neutral bilinear, the kinetic normalization or the current. Each of those remains a downstream calculation that must be carried out on the same background once the bottleneck is cleared.

This is the natural stopping point of the present paper. It has derived the nonlinear generation law under the stated intertwining assumption, shown that the resulting response has sufficient rank for three masses and frame rotation, falsified the simplest equal-winding realization, proved that the retained fold action does not select the diagnostic angle, and identified one explicitly typed Master-Action source derivative as the upstream bottleneck of the remaining physical normalization. A further neutrino spectrum scan before that derivative is returned would only recycle an unselected coefficient.

14 Conclusion

The VERSF neutrino programme has a direct operator route from completion geometry to the neutral bilinear, its canonical spectrum and its charged-current frame. The present paper has now taken that route far enough to separate the derived neutral mechanism from the remaining source-normalisation problem. Physical bath selection contributes through the selected background and matched response; excitation of an arbitrary auxiliary current chart is not a universal prerequisite.

The temporal, completion and charged-reference analyses delimit what may be compared consistently, and the audit of section 12.8 shows that the shared candidate's two failures have one origin, within that implementation, in its address ladder acting on the active neutrino. They supply a native stopped clock, compatible source normalization, finite kinetic/completion operators and explicit failed common-completion candidates. None of those earlier conditional constructions is promoted by the new fold calculation; rather, they provide the operator environment in which the fold-generated neutral response must ultimately be evaluated.

The nonlinear single-fold continuation materially advances the mass problem. Under the stated generation-intertwining assumption, it derives the half- and quarter-sequence dependence of the second-order response, and propagates them to two explicit full-rank neutral response matrices with a nonzero commutator. The blind equal-winding benchmark therefore constitutes a genuine test rather than a fit: it produces three masses but the wrong mass-gap hierarchy and is rejected. The alternative 13.28° orientation demonstrates reachability only; the retained action is exactly angularly flat at this order and does not select it.

The source audit follows the failed shortcut through the source rather than introducing another neutral coefficient. The direct A1-density/local-Higgs route, the known divergence-free current-cycle route and the primitive address-motion route do not provide the missing common response on the executed branch. Faithful D6 typing also excludes promotion of the reduced scalar J-H correlation. The correct current-history carrier is rank six, $E_2 \otimes (C, W, Y)$. A positive history law with the required factor geometry exists, and its local Fisher metric is fixed up to one common gain once the RCHS-1 factor Gram is supplied, but the retained HCMR/RRHF source returns rank zero on that physical leg. The identified upstream bottleneck is the background-contracted graded source jet T_{cyc}, equivalently the representation-resolved current derivative of the completed operational metric/history generator; the fold amplitude, orientation, restoring coefficient, addresses, neutral bilinear, kinetic normalization and current remain downstream of it. This is a sufficiently narrow and source-native target to close the present paper. The paper therefore advances the neutrino-mass derivation from an unspecified coefficient problem to a generation law derived under the stated intertwining assumption, a falsified blind benchmark and one explicit Master-Action attachment boundary; it does not yet issue a physical neutrino mass triple or PMNS matrix.

Appendix A Numerical verification and source scope

Closure kernel and Takagi convention. The verification script independently diagonalizes (12), reconstructs the three mixing angles, and obtains the Dirac phase from both its sine and cosine invariants. It checks the Takagi convention through $M = U^* D U^\dagger$ and verifies that $M^\dagger M$ shares the kernel's eigenframe. The auxiliary positive values used for that algebraic check are not proposed neutrino masses. The same script recomputes the occupation identities in Table 2 and the coefficients in (20) from the cited source values.

Covariance and cycle-sector calculations. The covariance tensors and their convergence and finite-ensemble checks are inherited from the executed covariance calculation [24]. The 334-variable cycle coefficients and residuals are inherited from its separately executed source continuation [25]. Neither is described here as a new computation, and no physical covariance is inferred from their unit-metric normalization.

Source editions. The source comparison distinguishes the 71-page and 89-page fermion-kernel editions bearing the suffix (4). Corollary 17.1 belongs to the 71-page text; its full mixed-tensor result supersedes the earlier primitive-leg-only account for that certificate. Source premises are carried forward at the level at which they were stated, including the EPM matter attachment and RRHF/FCHI completion structure.

Corpus discovery screens. For the admission and clock calculation, all 841 available catalogue entries were screened for anchoring gates, record-clock laws and neutral temporal response; the searches matched 73, 81 and 128 entries, with editions and duplicates included. For the common-readout and ratio calculations, the same 841 materialized entries, comprising 790 distinct text hashes with no missing text files, were screened for raw addresses, interface typing, relative scale, physical source selection and maintenance differentials; the retained source ledger contains 22 full texts [36]. These are complete discovery screens of the available corpus followed by contextual reading of controlling sources, not proof audits of every stored edition. Exact identities, hashes and matches accompany the working archives.

Admission executable. The admission executable [32] checks six immutable input hashes and reuses the native enumeration of 1,479 transient clock states. Direct absorbing-chain inversions verify the gated moment identities with maximum relative error below 1.8×10^{-15}. Both neutral implementations are evaluated at refinement levels 3 and 4; Takagi residuals are below 7×10^{-16}

and current-completeness residuals below 2×10^{-15}. The probabilities $q = 1,0.1,0.01,0.001$ illustrate independent admission identities and were not selected by observations. The package contains the admission derivation, executable, source ledger, discovery screen, numerical inputs, matrices and current coefficients.

Phase continuation. The phase calculation [33] checks its invariant blocks against the full 334-coordinate action, with maximum block residual 2.47×10^{-14} and covariance residual below 3.7×10^{-16}. Periodic quadratures with 96 and 192 points agree to better than 1.3×10^{-15} for the reported coherence and alignment quantities. The fixed-field measure, illustrative alignment window and continuity-only temporal scope are explicit.

Charged comparison and common readout. The charged/neutral comparison and common-completion derivation [34,35] retain the source premises behind the successful charged law and test the diagnostic tau insertion without selecting it as a physical model. For [36], the primary neutral Takagi values are verified with a 60-decimal-digit real symmetric dilation of the complex six-state matrix; this checks arithmetic on retained double-precision inputs, not 60-digit physical accuracy. Agreement with double-precision singular values is better than 2.1×10^{-13} relative. Charged and neutral reconstruction, current completeness and the coupling factorization are checked. Rescaling both mass matrices by 10^{-6} or 10^{6} changes the ratios by less than 1.1×10^{-14} relative. The archive includes both regulator levels, all six masses and current coefficients, canonical candidate matrices, the electron-normalized neutral matrix, address diagnostics, coupling probes and source-response arrays. Candidate results are stored separately from the still-unreturned physical prediction.

Structural audit of the shared candidate (section 12.8). Inputs: the printed level-4 blocks ϵY_ν and $\widehat{M}_R$ of [36], zero active block, identity kinetic pairing, RRHF charged frame $U_{eL} = I$, bridge $J_{e\nu} = I$, active projector $P_A = (I_3 \;\; 0)$; all six states kept; the light triplet is the three states of largest total active weight $\sum_\alpha |N_{\alpha i}|^2$. From these inputs alone the six Takagi values, the full 3×6 probability table, the RRHF-frame ceiling 0.0124584638, the gap ratio 2.2399×10^{-5} and the identity $\prod \mu_i = \left|\det \widehat{M}_D\right|^2$ were reproduced to all printed digits. The cyclic-shift residuals of the D-stripped cores are 7.0×10^{-18} and 1.1×10^{-16}. The eight placement tests use $M_D = D^{a_1} Y D^{a_2}$, $M_R = D^b Y_R D^b$ and the same frame, current and selection rule:

Table A1. Placement tests of the address factor at level 4, RRHF charged frame. The last three columns are the physical electron current probabilities on the light triplet, in ascending mass order; the column heads abbreviate p_{e1} and p_{e,i_k}.

(a_1, a_2, b)	Triplet	Gap ratio	p_e1	p_e,i_1	p_e,i_2	p_e,i_3
(1, 1, 1)	1, 2, 4	2.2399×10^{-5}	0.99688	0.99688	2.2×10^{-4}	1.6×10^{-6}
(1, 0, 1)	1, 2, 5	5.0872×10^{-3}	0.99514	0.99514	3.1×10^{-3}	7.7×10^{-5}
(0, 1, 1)	1, 3, 4	0.24871	0.48169	0.48169	0.00837	0.06134
(0, 0, 1)	1, 2, 4	0.34383	0.26356	0.03836	0.22709	0.25346
(1, 1, 0)	1, 2, 3	5.1265×10^{-10}	0.99982	0.99982	1.8×10^{-4}	1.4×10^{-6}
(1, 0, 0)	1, 2, 3	2.2779×10^{-5}	0.99978	0.99978	2.2×10^{-4}	1.8×10^{-6}
(0, 1, 0)	1, 2, 3	1.5149×10^{-5}	0.92799	0.92799	0.02171	0.05010
(0, 0, 0)	1, 2, 3	0	0.66476	below 10^{-6}	0.66476	0.33228

Here p_{e1} is the largest electron probability over all six states. The row (0,0,0) is the C_3-symmetric limit; its first two masses coincide to 10^{-15} relative, so the gap ratio is zero and the electron entries on the degenerate pair are basis-dependent (the values shown use the electron-aligned basis); the non-degenerate entry 0.33228 is a physical current probability, below the normalized flavour proportion 1/3 by the sterile admixture. The executable reproducing this table from the printed blocks accompanies the reproduction archive of [36].

Nonlinear fold source. The source executable reproduces $g_{source} = 1/2$, evaluates the complete second-order attachment in the original 226 field coordinates and relaxes the 185-dimensional complement of the retained E_1 pair. Fifty-five independent finite-difference and stationarity checks pass, with largest finite-difference residual 4.48×10^{-11} and algebraic residuals below 8×10^{-15}. The result files retain the full lifts, second attachments and pulled-back source residual derivatives. These checks validate the declared finite-source differentiation; they do not substitute for the still-missing representation-resolved neutral readout or kinetic residue.

Blind same-wheel diagnostic. Appendix B makes the fold-spectrum calculation self-contained: it prints the three inherited response inputs, derives the compressed map from the seven-address Sylvester equation, specifies the complex source contraction and supplies a short reproduction program. An independent evaluation from the printed inputs alone reproduces both matrices in (81), the singular values in (83), the 13.28° diagnostic spectrum and its quoted mixing angles, and the closed response map agrees with the full seven-address stationary solve to an absolute Frobenius residual below 10^{-15} for the two generation shapes and the identity source. The response machinery is the same-wheel scalar-completion model in [42]; it is distinct from the six-state charged/neutral candidate of [36]. No neutrino observable selects the equal-winding benchmark. The 13.28° orientation remains a diagnostic choice, and numerical reproducibility does not establish a physical neutral pullback, kinetic residue, selected fold orientation or mass scale.

Source-attachment audit. The archived nonlinear fold calculation retains the nonzero A_1 density and current responses while returning zero local Higgs/completion second attachment on the separable branch. The induced fold current is numerically orthogonal to the established divergence-free neutral cycle subspace at relative level below 1.8×10^{-15}. The fold angular invariant (84) is checked algebraically. The faithful current-history audit distinguishes the reduced scalar history metric from the D_6-typed E_2 current carrier, for which the executed retained HCMR/RRHF pullback has rank zero against a rank-six target. RCCSJ-1 independently constructs a positive rank-six history realization with Fisher metric $g_{cyc}^2(2/31)H_{CWY} \otimes I_2$ and leaves only its common source gain unselected. These checks localize the remaining physical input to the background-contracted graded source jet; they do not set that gain from the neutrino data.

Appendix B Explicit reproduction of the fold spectrum

B1 Inputs and source provenance

This appendix reproduces the dimensionless diagnostic of sections 13.5–13.7. It uses the finite same-wheel response of [42], with identity fermion kinetic pairing, charged-lepton frame and current bridge. These are diagnostic premises. The physical sector attachment and common dimensional mass scale remain unselected.

The three numerical inputs are the scalar response weight w, the hub-to-rim conductance c_s and the neighbouring-rim conductance c_r. They are inherited from the driven seven-state history calculation in [42], before any comparison with neutrino data:

$$w = 0.15246711893053608, \qquad c_s = 0.027060532986541, \qquad c_r = 0.04048748943797676. \qquad (B1)$$

The retained archive is Electron_Rules_Neutral_Extension_Reproduction.zip, with SHA-256 digest e47b9aa2b0c84ec975b14a0c16f9f4aeda0b7b36d54c6a0f8c24645ae64127c9. Within its top-level folder of the same name, neutral_baseline_interference_results.json records these inputs under response (w, spoke, rim). The formulas below occur in neutral_action_selection_test.py and neutral_baseline_interference_test.py. The archive's fourteenth calculation, neutral_second_attachment_test.py, supplies the separate nonlinear attachment results; it does not by itself select a physical neutral readout.

B2 The response map

Let P_0 be the projector onto the uniform vector of the three-address carrier and P_2 its orthogonal complement, with $\mathbf{1}$ the three-component all-ones vector. For any complex 3×3 source S, extend the real scalar-completion response complex linearly:

$$P_0 = \frac{1}{3}\mathbf{1}\mathbf{1}^T, \qquad P_2 = I_3 - P_0, \qquad (B2)$$

$$R(S) = r_0 P_0 S P_0 + r_A P_2 S P_2 + r_B (P_0 S P_2 + P_2 S P_0), \qquad (B3)$$

$$r_A = \frac{w}{w + c_s + 3c_r} = 0.5065519054695844, \qquad (B4)$$

$$r_B = 2w\left(\frac{6/7}{2w + c_s + 3c_r} + \frac{1/7}{2w + 8c_s + 3c_r}\right) = 0.6441594422558468, \qquad (B5)$$

$$r_0 = \frac{36}{49} + \frac{w}{49(w + 7c_s)} + \frac{24w}{49(2w + 7c_s)} = 0.8948550682202270. \qquad (B6)$$

B3 Full carrier check and complex contraction

The full-carrier realization fixes the map without a fitted propagation matrix. Number the wheel hub 0 and the rim vertices 1 to 6. Give every hub–rim edge conductance c_s and each neighbouring rim edge conductance c_r, with indices taken cyclically. If C is this symmetric conductance matrix, define $\Delta = diag(C\mathbf{1}) - C$ and $H = wI_7 + \Delta$. The 7×3 embedding E has $E_{j+1,j} = E_{j+4,j} = 1/\sqrt{2}$ for $j = 0,1,2$, with all other entries zero. Then

$$HX + XH = 2wESE^T, \qquad R(S) = E^T XE. \qquad (B7)$$

Because H is positive, the Sylvester equation has a unique solution. Direct evaluation of (B7) agrees with (B3) for S_1, S_2 and I_3 to an absolute Frobenius residual below 10^{-15} with a standard solver. The saved H in neutral_baseline_interference_arrays.npz agrees with the wheel construction above to numerical precision. The response matrices of section 13.5 are

$$A_0 = R(S_1), \qquad B_0 = R(S_2), \qquad C_0 = R(I_3), \qquad (B8)$$

of which the first two reproduce (81). The identity response needed for the complete complex source is

$$C_0 = \begin{pmatrix} 0.635986293 & 0.129434388 & 0.129434388 \\ 0.129434388 & 0.635986293 & 0.129434388 \\ 0.129434388 & 0.129434388 & 0.635986293 \end{pmatrix}. \qquad (B9)$$

For real u and v, the source (82) has real part $u^2 I_3 - v^2 S_2$ and imaginary part $2uvS_1$, so by complex linearity of R its complete diagnostic contraction is

$$M(u, v) = u^2 C_0 - v^2 B_0 + 2iuvA_0. \qquad (B10)$$

The $u^2 C_0$ term is fixed by this diagnostic assembly; it is not an added fit parameter or an evaluation of the physical baseline sought in section 13.8. $M(u, v)$ is complex symmetric, so its singular values are its Takagi values. For equal winding, $u = v = 1/\sqrt{2}$, write $M_{45} = M^R + iM^I$. In the dimensionless normalization used in (83),

$$M^R = \begin{pmatrix} 0.010602008 & 0.027802950 & 0.032103186 \\ 0.027802950 & 0.234960857 & 0.049304128 \\ 0.032103186 & 0.049304128 & 0.291050569 \end{pmatrix}, \qquad (B11)$$

$$M^I = \begin{pmatrix} 0.620279614 & 0.090793119 & 0.079325824 \\ 0.090793119 & 0.321134482 & 0.056391235 \\ 0.079325824 & 0.056391235 & 0.171561917 \end{pmatrix}. \qquad (B12)$$

B4 Numerical outputs and a minimal reproduction

Sorting singular values in ascending order gives the following dimensionless outputs. Their common scale is omitted, and neither row is a physical mass prediction.

Table B1. Fold-spectrum diagnostic outputs from (B10).

Orientation φ	Singular values	Gap ratio
45°	0.286634587805, 0.382294398409, 0.687698786635	0.163752544401
13.28°	0.483962239128, 0.499451377771, 0.865746312518	0.029560084950

The commutator norm is $\|[A_0, B_0]\|_F = 0.036700241543$. At 13.28°, using columns (2,3,1) of the ascending-singular-value left frame gives $\theta_{12} = 37.379918°$, $\theta_{13} = 3.523535°$ and $\theta_{23} = 41.747550°$. The gap ratio is always computed in ascending mass order. The separate column permutation checks only the frame shape and cannot supply a consistent physical mass assignment by itself.

The following program requires only NumPy and uses the printed inputs directly. It prints A_0, B_0, C_0, the full complex matrix for each orientation, the ordered singular values, their gap ratio and the diagnostic angles. It uses no observed neutrino input to construct the matrices; the 13.28° angle is the stated orientation diagnostic, not an action-selected prediction.

```
import numpy as np
w = 0.15246711893053608
cs, cr = 0.027060532986541, 0.04048748943797676
P = np.ones((3, 3))/3; Q = np.eye(3) - P
a = w/(w + cs + 3*cr)
b = 2*w*((6/7)/(2*w+cs+3*cr) + (1/7)/(2*w+8*cs+3*cr))
t = 36/49 + w/(49*(w+7*cs)) + 24*w/(49*(2*w+7*cs))
def R(S):
    return t*P@S@P + a*Q@S@Q + b*(P@S@Q + Q@S@P)
A = R(np.diag([1, .5, .25]))
B = R(np.diag([1, .25, .0625])); C0 = R(np.eye(3))
print("A0, B0, C0:", A, B, C0, sep="\n")
print("commutator:", np.linalg.norm(A@B - B@A))
for phi in (45., 13.28):
    u, v = np.cos(np.deg2rad(phi)), np.sin(np.deg2rad(phi))
    M = u*u*C0 - v*v*B + 2j*u*v*A
    U, m, _ = np.linalg.svd(M)
    order = np.argsort(m); m = m[order]; U = U[:, order]
    gap = (m[1]**2-m[0]**2)/(m[2]**2-m[0]**2)
    p = abs(U[:, [1, 2, 0]])**2
    angles = np.rad2deg([
        np.arctan2(np.sqrt(p[0,1]), np.sqrt(p[0,0])),
        np.arcsin(np.sqrt(p[0,2])),
        np.arctan2(np.sqrt(p[1,2]), np.sqrt(p[2,2]))])
    print(phi, M, m, gap, angles, sep="\n")
```

This reproduction closes the numerical specification of the diagnostic. It leaves the physical source gain, neutral attachment, kinetic normalization, mass scale and fold selection at the open boundary stated in section 13.12.

References

All VERSF works below are by K. Taylor. Parenthesized edition labels identify the supplied corpus versions, rather than journal volume numbers. Working calculations are identified as such.

[1] *The VERSF Master Substrate Action and Constitutive-Descent Theorem*. Edition (11). Common action, six-defect parent and typed primitive bilinears.

[2] *The VERSF History Measure Generation Geometry and Branch Closure Theorem*. Edition (11), especially section 56A. Common history generator and its distinct bath, kinetic and information compressions.

[3] *The Yukawa Operator from Completion Channel Misalignment in VERSF*. Edition (15)(1). Completion insertion and relative flavour frames.

[4] *The Weak Commitment Neutrino Operator in VERSF*. Edition (15), sections 2–3. Weak-commitment co-scaling and the neutrino eigenframe.

[5] *The PMNS Closure Hamiltonian Theorem in VERSF*. Edition (13), compatibility theorem and section 21. Canonical mass lifts, current bridge and numerical benchmark.

[6] *Completion Density and Flavour Mixing from Refinement Geometry*. Edition (4)(1), sections 6–7. Completion geometry and relative mixing frames.

[7] *The Neutrino Mass and Charge Conjugation Completion Theorem in VERSF*. Edition (13). Typed neutral bilinears and Majorana completion.

[8] *NSTC-1 The Same Branch Neutral Sector Return*. Edition (10), sections 6–7. Full neutral return and the direct mixed-derivative alternative to an additional selector.

[9] *The Bath Criterion*. Edition (9)(1). Conserved shared weight and continuous redistribution.

[10] *Bath or Ledger*. Edition (9)(1). Physical interpretation of the shared transport structure.

[11] *The Neutral Fermion Branch and Carrier Census Completion Theorem in VERSF*. Edition (8). Neutral carrier and branch census.

[12] *The Microscopic Closure Hessian and Full Yukawa Evaluation Theorem in VERSF*. Edition (13). Same-action mixed derivatives and Yukawa evaluation.

[13] *The Finite Master Action Fermion Kernel Reconstruction and Source Identity Gate in VERSF*. Edition (4), 71-page source, Corollary 17.1, pp. 43–44. Full mixed-tensor completion descent on the declared RRHF/FCHI structure.

[14] *The Finite Master Action Fermion Kernel Reconstruction and Source Identity Gate in VERSF*. Edition (4), 89-page source. Finite kinetic reconstruction and geometric principal-symbol continuation.

[15] *The Refinement Geometry Higgs Spatial Propagation and Kinetic Residue Reduction Theorem in VERSF*. Edition (4). Protected scalar propagation and kinetic reductions.

[16] *The Weak Doublet Projection Theorem in VERSF*. Edition (9)(1). Ordered carrier projection. See also *The Projected Weak Doublet Closure Hamiltonian in VERSF*, edition (18)(1).

[17] *RCHS-25 Record Clock Event Grain Separation and Source Stopped Endpoint Theorem*. Working calculation. Native cover-and-return stopping law.

[18] *RCHS-26 D36 Stationary Traffic Matter Endpoint Generator*. Working calculation. Orientation-even local traffic and EPM matter attachment.

[19] *RCHS-27 Source Stopped Pre Record Current Occupation and E2 Gram*. Working calculation, with *Corpus Checked Fold to Neutrino Response* continuation. Stopped occupation factors and current Gram.

[20] *VERSF Gauge Conserving Temporal Attachment*. Working calculation. Continuity attachment on the full finite action chart.

[21] *Compatible Field Source History Response*. Working calculation. Frequency-dependent source identity and relaxed scalar response.

[22] *RCHS-35 Configuration Dynamics Refinement Geometry and Record Conditioned Curvature Response in VERSF*. 22 September 2026, Parts XIII–XV and consolidated claim ledger.

[23] *Current Cycle Attachment to Generation History and the Neutral Matrices*. 22 September 2026, RCHS-35 archive A69, with *Current_Cycle_Attachment_Reproduction.zip*. See also the raw-transfer derivative execution of 2 September 2026.

[24] *Physical Bath Selection and Neutrino Mixing in VERSF*. PBSN1 covariance calculation, 23 September 2026, section 8 and numerical reproduction. Conditional constrained covariance response retained in section 9.

[25] *PBSN1 Section 10 Source Execution*. Corrected working calculation and reproduction, 23 September 2026. The 334-variable cycle-sector execution and its scoped temporal conclusion.

[26] *The Multi-Regulator Full-Faithful Source Assembly, Exact Six-Mode Bath Identification and HM4 Provenance Execution-Boundary Theorem in VERSF*. Edition (12). Common source assembly and regulator programme.

[27] *The Primitive Full-Carrier Provenance, Recurrent Fisher Pullback, Source-Projected Bath Factorisation and Conditional Doob–Kraus Descent Theorem in VERSF*. Edition (15). Primitive carrier and source provenance.

[28] *Void Coupling as Phase Coherence*. Edition (5)(2), sections 2, 5–7, 9–10 and 12. Distinct phase coupling, alignment probability, structural admission and anchoring depth; conditional bridge formulas and the eigenmode computation programme.

[29] *Void Anchoring, Action Quantization, and the Mass Scale*. Edition (1)(4), sections 5–9. Conditional action-per-completion law and the inverse-mean rest-energy convention.

[30] *Yukawa Hierarchy*. Edition (7)(1), section 5c. Reconciliation of the anchoring probability and completion-count mass coordinate.

[31] *VERSF Void Anchoring and Void Coupling Mathematical Core*. Version 2, section 8. Alternative resistance-scale convention, distinguished from the rest-energy convention in section 11.1.

[32] *Incorporating Void Coupling and Emergent Time into the Neutrino Calculation*. Working calculation, 23 September 2026. Derivation, full-catalogue discovery ledger and executable six-state results; accompanying archive *VERSF_Void_Coupling_and_Emergent_Time_Reproduction.zip*.

[33] From the VERSF Action to Phase Coupling. Working calculation, 23 September 2026. Polar-bond phase potential, full-action phase blocks, conditional periodic alignment functions and temporal-scope check. Report: VERSF_Action_to_Phase_Coupling_Derivation.md; accompanying reproduction archive.

[34] Charged and Neutral Completion Comparison in VERSF. Working calculation, 23 September 2026. Distinct RRHF and neutral address hierarchies, common-parent scope, canonical-rescaling audit and shared-operator comparison. Report: VERSF_Charged_Neutral_Completion_Comparison.md; accompanying reproduction archive.

[35] Charged Localization, Participation and Neutral Masses as Common Completion Readouts. Working calculation, 23 September 2026. Conditional charged resolvent realization, Schur reduction, projector-versus-amplitude distinction, diagnostic insertion and positive common-completion nonuniqueness family. Report: VERSF_Common_Completion_Sector_Readouts.md; accompanying reproduction archive.

[36] Three Neutrino-to-Electron Mass Ratios from the Common Completion Candidate. Working calculation, 23 September 2026. Same-construction electron normalization, full six-state results, coherent-overlap factorization, coupling probes, raw-address audit and charged-preserving ratio derivatives. Report: VERSF_Neutrino_to_Electron_Mass_Ratios.md. Reproduction: VERSF_Neutrino_Electron_Ratios_Reproduction.zip.

[37] The Charged-Lepton Absolute Scale Theorem in VERSF. Edition (8), sections 11–12. Canonical localization–closure magnitude, stiffness-identification premise and conditional tau participation factor.

[38] A Charged-Lepton Ratio Law from Localization and Saturated Participation. Edition (3)(1). Conditional charged hierarchy and ratio construction.

[39] Charged-Lepton Participation Gating in Saturated Closure. Edition (3)(1). Charged participation and gating construction; used with the conditionality stated in [37].

[40] NIST. 2022 CODATA Recommended Values of the Fundamental Physical Constants. Electron mass energy equivalent in MeV: 0.51099895069(16). https://physics.nist.gov/cuu/Constants/Table/allascii.txt (accessed 23 September 2026). Used only for the illustrative electron-calibrated conversion in section 12.3.

[41] Uniform Binary Admissibility and the Half-Lazy Closure Operator. Working manuscript, sections 7–10. Generation-independent log-access intertwining, odd-mode one-half eigenvalue and the conditional χ-halving theorem.

[42] Applying the Electron Rules to a Neutral Mass Readout. Working calculation, 24–25 September 2026. Reproduction archive: Electron_Rules_Neutral_Extension_Reproduction.zip (SHA-256 in Appendix B1). The seven-address response used in sections 13.5–13.7 is specified by neutral_baseline_interference_test.py and neutral_action_selection_test.py; its frozen parameters are in neutral_baseline_interference_results.json. The nonlinear EPMD differentiation and complete second-order relaxed attachment are in neutral_second_attachment_test.py and its accompanying inputs and results. Appendix B gives a self-contained reproduction of the fold-spectrum diagnostic.

[43] RCHS-8 Primitive-Quadratic-Polar-Lift and MBCI-Homogeneity-Separation. September 2026. Unique positive source-equivariant quadratic primitive current intensity up to one common scale, exact C/W/Y scores and one-Fold Fisher decomposition.

[44] PCRS-1R Primitive Current-to-Record Kraus Descent and Gauge-Normalisation Boundary in VERSF. 2 September 2026. Forward marked-stack audit, exact projective radial blindness and the current-cycle Kraus no-return theorem.

[45] RCCSJ-1 The Representation-Resolved Current-Cycle Source Jet and Positive-History Realisation in VERSF. 3 September 2026. Graded-source typing, D6 E2 rigidity, rank-six requirement and positive current-history realization with one common gain.

[46] Actual HCMR/RRHF Current-Cycle Source Execution in VERSF. Executable source audit. Exact retained-order rank-0 versus rank-6 current-cycle no-return and exhaustion of the radially blind HCMR/RRHF/XI-3/XI-4 route.

[47] CRTF-HMGBC Current-Cycle Raw-Transfer Derivative Execution and HM12 Current-Score Normalisation/Doob-Descent Audit. 2 September 2026. Exact raw-transfer/Doob score formulas, canonical branch zero and conditional irreversible current-cycle history scores.

[48] RCHS-1 Twelve-Sector Current Rigidity, Source-Generated Rank-Six Current Stiffness and Current-Link Schur Reciprocity, updated through RCHS-4. September 2026. Exact H_CWY factor geometry, source current stiffness and current-link reciprocity with physical normalization boundary.

[49] Can the Master Action Select the Neutrino History-Source Variation? Working source audit and calculation, 21 September 2026. Coupled Master-Action stationarity, source-shape test and physical current-to-history derivative boundary.