

The First Source-Fixed Neutrino Mass-Ratio Prediction Attempt in VERSF

Fold indistinguishability, the projection hypothesis, and a neutrino hierarchy from
universal completion geometry

Keith Taylor, AIDA Institute
September 2026

General Reader Summary

This paper asks a deliberately testable question: can one specified VERSF neutrino construction produce a definite numerical mass-spacing prediction without tuning it to the measured neutrino hierarchy? The projected-source construction tested here gives a frozen mass-spacing ratio of 0.02822, and the archived calculation has been independently rerun to reproduce that value. Because the projection hypothesis was introduced after earlier constructions had failed, the paper does not claim that this is a fully blind or uniquely selected prediction of all VERSF; it is the frozen prediction of the declared projected-source construction.

That number can be tested by experiment. Neutrino oscillation measurements determine the same relative spacing between the three neutrino masses, and future measurements should make that spacing increasingly precise. If sufficiently precise data settle on a value inconsistent with 0.02822, this projected-source construction would be falsified under its stated assumptions. Agreement would support the construction, although it would not by itself prove the whole VERSF framework.

The paper also isolates what remains theoretically open. Across the tested range, the predicted spacing changes almost directly with the size of the input coming from the Fold construction, so the present difference from the experimental comparator amounts to a five-to-six-percent question about how strongly that source should enter the later history calculation. That is a post-prediction diagnostic, not an invitation to tune the answer: the frozen value stays on record unless VERSF independently derives a different conversion from its own physics. One obvious shortcut—simply setting the source-action scale equal to the history scale—is rejected because the two quantities being equated are not the same kind of object.

Abstract

This paper is a deliberately predictive test of a specified VERSF neutrino construction: the aim is to carry the framework to a definite mass-spacing ratio, freeze that result, and expose it to increasingly precise experiment rather than tune a coefficient to the measured hierarchy. The same-wheel and D-free single-orientation constructions are rejected at retained order. Under Hypothesis P, projection to the reflection-fixed Fold source gives the active-dominant shape $m_1:m_2:m_3 = 1:1.0468:2.0967$ and $R_v = 0.02822$. The complete projected-source numerical chain has now been independently reconstructed from the archived execution packages. A finite-amplitude scan shows that the returned hierarchy is approximately proportional to the source amplitude to within about 1% across the tested factor-three range, $R_v \approx 0.1696|q|$, with an effective finite-amplitude slope of about 0.17. The stated experimental comparator corresponds to $|q| \approx 0.1756$, about 5.6% above the archived amplitude 0.1664; this is a post-freeze diagnostic, not a retuning prescription. The canonical source susceptibility $g_{ax}/\kappa E^2 = (1/2)/3 = 1/6$ and the type-correct same-history value $\kappa E^2 \approx 2.92415$ are independently derived in the Primitive-Matter Selection programme [4], not selected from neutrino data. Direct

identification of the source-action stiffness with the history/Fisher stiffness remains inadmissible because the two are differently typed; its numerical value $R_v \approx 0.00991$ is diagnostic only. Thus $R_v = 0.02822$ is the frozen prediction of the declared projected-source construction. The remaining normalisation and projection/generation provenance determine whether it can be promoted to a fully first-principles VERSF prediction, not whether the present construction makes a testable numerical claim.

Contents

1. Objective, claim level and comparison protocol
2. Source audit and inherited neutral structure
3. Derivation of the half/quarter generation law
4. Quartic flatness and the sixth-order anisotropy
5. Admissible Fold orientations and the two readings of indistinguishability
6. Same-wheel diagnostic and its rejection
7. Direct Dirac, sterile and kinetic responses
8. Mixed fourth-order Fold-fermion vertex
9. CFB-2 capacity test
10. Universal history-derived completion geometry
11. D-free representation-resolved neutral compression
12. Retained-order corrections and the sensitivity of the ratio
13. Absolute-scale consequence and cosmological comparison
14. Status ledger, reproducibility package and open gates
15. Conclusion
 - Appendix A: Sixth-order orientation-selection derivation
 - Appendix B: Why CFB-2 gives a scale-free ratio
 - Appendix C: Claim ledger
 - Appendix D: Arithmetic identities used in the text
 - Appendix E: Projection, reflection invariance and the generation average
 - Appendix F: Non-extension of finite-set zeros to the projected source
 - Appendix G: Schur cancellation versus active-sterile mixing
 - Appendix H: First-order variation of the ratio under a mass and kinetic correction
 - Appendix I: Consistent-basis test of representative equivalence
 - References

1 Objective, claim level and comparison protocol

1.1 Objective

The objective is to determine how far the present VERSF source stack can be pushed toward a neutrino mass-ratio prediction without inserting measured neutrino masses, mass splittings or PMNS data into the source selection. Every branch choice is frozen before observational comparison. A numerical agreement obtained only after choosing a Fold direction, completion graph or neutral coefficient from the measured hierarchy is classified as reachability, not prediction.

The principal observable is the scale-free mass-squared gap ratio

$$R_\nu = \frac{m_2^2 - m_1^2}{m_3^2 - m_1^2}.$$

An unknown common mass scale cancels from R_ν , so the ratio can be tested before the absolute scale in electronvolts is derived. Section 13 shows that the converse also holds: once R_ν and the two remaining ratios are fixed, one measured splitting fixes the absolute scale.

1.2 Outcome classes

Outcome	Criterion	Interpretation
PASS	One common source branch fixes the neutral operator and the frozen R_ν is statistically compatible with the comparator.	A source-fixed mass-ratio prediction candidate.
FALSIFIED AT THIS ORDER	The source closes uniquely but the frozen R_ν is incompatible with the comparator.	The tested branch is rejected.
UNDERDETERMINED	A continuous source coefficient, Fold direction or neutral readout remains free.	No numerical prediction is licensed.
CAPACITY ONLY	A declared candidate geometry can produce a good ratio but is not source-derived.	Mechanism test, not a prediction.
CONDITIONAL CANDIDATE	The source closes to a frozen R_ν only under an assumption not yet derived from the source.	Candidate, with the assumption named as the decisive open object; its discrepancy from the comparator is reported separately.

1.3 Non-insertion rule and pre-declared comparator

The following quantities are comparison outputs only and may not be used to select a source coefficient before the calculation is frozen: the measured neutrino mass splittings, the PMNS matrix, the lightest neutrino mass and the reference ratio. Post-failure inverse calculations are permitted only when labelled as diagnostics of the tensor structure a missing source term would have to generate.

The comparator is fixed in advance as follows. Ordering: normal, $m_1 < m_2 < m_3$; the ratio as defined is positive only for normal ordering, and no inverted-ordering comparison is made after the fact. Dataset: NuFIT 6.0, normal ordering, IceCube 2024 with Super-Kamiokande atmospheric data [20], with $\Delta m_{21}^2 = (7.49 \pm 0.19) \times 10^{-5} \text{ eV}^2$ and $\Delta m_{31}^2 = (2.513_{-0.019}^{+0.021}) \times 10^{-3} \text{ eV}^2$, with central-value ratio

$$R_\nu^{ref} = 0.0298050.$$

A local approximation with $\sigma_S = 0.19 \times 10^{-5} \text{ eV}^2$, $\sigma_D = 0.020 \times 10^{-3} \text{ eV}^2$ and zero covariance gives $\sigma_R \approx 7.92 \times 10^{-4}$ (Appendix H); this symmetrises the asymmetric intervals and neglects their correlation, and is not a replacement for the published likelihood. The PDG 2025 tables [22] quote Δm_{32}^2 rather than Δm_{31}^2 ; the denominator used here requires $\Delta m_{31}^2 = \Delta m_{32}^2 + \Delta m_{21}^2$, and the corresponding PDG fit gives 0.02981, consistent with the comparator. The comparison statistic is the signed local residual $(R_\nu - R_\nu^{ref})/\sigma_R$ (theory minus observation), used as a description of the discrepancy and not as an exclusion level: an exclusion level would require a specified threshold, the experimental likelihood and a justified treatment of theoretical uncertainty, and no theoretical error is assigned merely to enlarge the acceptance interval. Use of a standard three-neutrino fit is itself conditional on the completed neutral and current model reproducing that effective oscillation description.

1.4 Chronology of assumptions

The order in which the constructions of this paper were fixed is recorded so that the claim level of each result can be read off directly.

1. The same-wheel response matrices A_0, B_0 were frozen before the sixth-order angle selection existed (Section 6).
2. The sixth-order orientation selection was derived from the nonlinear source (Section 4), with no neutrino input.
3. The universal completion geometry and the D-free compression were executed on the single-orientation source images, giving $R_\nu \approx 0.0536$, and this was compared with the comparator (Sections 10 and 11.2).
4. The No-Subordinate-Distinguishability principle was then articulated for this calculation, the projected source Q_{phys} was defined, and the projection calculation was executed and frozen before its output was compared (Sections 5 and 11.3 to 11.5).

Consequently the projection result is not fitted to neutrino data: no neutrino datum enters any number in the chain. It is not, however, the output of a blinded design, because the decision to project post-dates knowledge that the single-orientation result was too high. Throughout the paper the phrase "not fitted to neutrino data" is used for what the record supports, and "frozen before comparison" for the numerical protocol; the word "blind" is reserved for the same-wheel diagnostic, whose construction genuinely preceded every selection.

2 Source audit and inherited neutral structure

2.1 Audit scope

The calculation reconciles the controlling neutral, generation, source-provenance, history, completion and CAR/Fock papers with the September execution notes and machine-readable returns [1] to [13], [17] to [19]. It consumes the corpus-wide catalogues and re-reads the source documents that control the mass-ratio chain. This is a source reconciliation and targeted proof audit of the relevant chain, not a line-by-line re-proof of every historical document.

2.2 Neutral parent and whitening convention

In the all-left-handed convention the six-state neutral quadratic coefficient is

$$B_N = (B_L \ B_D ; B_D^T \ B_R),$$

with $B_N = B_N^T$ the complex symmetric neutral mass coefficient and $K_N = K_N^\dagger > 0$ its Hermitian kinetic metric. Let $Z = K_N^{-1/2}$ denote the principal positive Hermitian inverse square root. With $\psi = Z\chi$ the canonical mass matrix and its Takagi decomposition are

$$\mathcal{M}_N = Z^T B_N Z, U^T \mathcal{M}_N U = \text{diag}(m_1, \dots, m_6), m_i \geq 0.$$

The kinetic square root is real symmetric only when the kinetic matrix is real; the transpose in the mass congruence and the Hermitian conjugate in the kinetic transformation are distinguished throughout. A block-diagonal kinetic metric $K_N = \text{diag}(K_L, K_R)$ in the typed active/sterile basis is an assumption of the declared typed parent; if a returned kinetic correction has active-sterile cross terms it is whitened as part of the full matrix. On a type-I-dominant branch the light operator

$$S_c = K_L^{-T/2} (B_L - B_D B_R^{-1} B_D^T) K_L^{-1/2}$$

is used as a reduction tool only; the six-state Takagi problem remains primary wherever the sterile hierarchy is not controlled, and every headline number in Sections 11 and 13 is a six-state result.

2.3 Frozen-address warning

The raw representation addresses are fixed within the primitive source chart. A physical derivative of the neutral readout may come from the completion operator, support/polar motion, the physical chart, the record map or the fermion kinetic metric, but it cannot be manufactured by differentiating a frozen raw address as if it were a free flavour matrix.

3 Derivation of the half/quarter generation law

3.1 The retained Fold coordinate

Retain the real E_1 Fold coordinate $x = (u, v)$, where u is reversal-even and v is reversal-odd. Under the half-lazy generation refinement, conditional on the neutral Fold intertwining with the same generation operator, the odd component is halved at each generation step:

$$x^{(g)} = (u, 2^{-(g-1)} v), g = 1, 2, 3.$$

One common u and one initial odd amplitude v generate all three layers; this is stronger than assigning three independent generation coefficients.

3.2 Quadratic $E_1^2 \rightarrow E_2$ source map

The primitive nonlinear source returns the unique D_6 -covariant quadratic map

$$Q(u, v) = (u^2 - v^2, 2uv).$$

Substituting the half-lazy refinement gives

$$Q(x^{(g)}) = (u^2 - 4^{-(g-1)} v^2, 2^{-(g-1)} 2uv),$$

so the only non-common generation sequences available at this order are $1, 1/2, 1/4$ and $1, 1/4, 1/16$. After one common representation-resolved neutral pullback the neutral response has the constrained form

$$M_N^{(g)} = M_N + 2^{-(g-1)} A + 4^{-(g-1)} B + O(|x|^3),$$

with A and B common response matrices that are not adjustable generation by generation.

3.3 Order of projection and generation refinement

The half-lazy map acts on the Fold coordinate before the quadratic source is formed. With $\lambda_g = 2^{-(g-1)}$ and $T_g = \text{diag}(1, \lambda_g)$,

$$Q(T_g x) = (u^2 - \lambda_g^2 v^2, 2\lambda_g uv).$$

This map commutes with the residual sign changes of u and v . For $g > 1$ it does not commute with a general rotation of the E_1 plane, so full D_6 covariance of the refinement requires a simultaneous transformation of the refinement frame; the residual reflection is sufficient for the comparison made in this paper.

Under the orbit-invariant reading of Section 5.4 the generation law is applied to any one representative. Under the projection hypothesis of Section 5.5 the order of the two operations matters. Averaging after applying the generation map to the four unit-amplitude admissible representatives gives (Appendix E)

$$\bar{q}_g = \frac{1}{4} \sum_{\varphi \in \mathcal{A}_F} Q(T_g x(\varphi)) = \left(\frac{1}{4} - \frac{3}{4} 4^{-(g-1)}, 0\right), \bar{q}_1 = \left(-\frac{1}{2}, 0\right), \bar{q}_2 = \left(\frac{1}{16}, 0\right), \bar{q}_3 = \left(\frac{13}{64}, 0\right).$$

These retain a common contribution and a quarter-law contribution, and the odd half-law contribution cancels; they are not obtained by multiplying $(-1/2, 0)$ by one decaying coefficient. Projecting first and refining the bisector preimage $x_* = (0, 1/\sqrt{2})$ instead gives $Q(T_g x_*) = (-4^{-(g-1)}/2, 0)$, that is $(-1/2, 0)$, $(-1/8, 0)$, $(-1/32, 0)$ (Appendix E). The two orders therefore define two explicit candidate sequences s_g for the first component, with the second component zero in both:

g	Refine then project, $s_g = \frac{1}{4} - \frac{3}{4} 4^{-(g-1)}$	Project then refine on x_* , $s_g = -\frac{1}{2} 4^{-(g-1)}$
1	$-1/2$	$-1/2$
2	$1/16$	$-1/8$
3	$13/64$	$-1/32$

These sequences illustrate two inequivalent prescriptions; they coincide at $g = 1$ and differ in sign and magnitude at $g = 2, 3$. The archived rerun recovers the numerical projected-source insertion used by the downstream history solve, but it does not reconstruct the upstream generation ordering that produced that insertion. In particular it does not decide between refine-then-project and project-then-refine, nor whether the scalar invariant r is retained. Those provenance choices remain part of Hypothesis P and must be derived or recovered explicitly; the reported spectrum is conditional on them and is not described as following uniquely from averaging the refined representatives.

4 Quartic flatness and the sixth-order anisotropy

4.1 Flatness at quartic order

At cubic and quartic order the induced E_2 norm satisfies $|Q(x)|^2 = (u^2 + v^2)^2$, so the reduced action is independent of the polar angle φ in the E_1 plane at that order. A fitted angle (such as the value 13.28° that reproduces the measured hierarchy in the quartic bridge) demonstrates reachability but cannot be promoted to a physical source orientation.

4.2 Sixth-order reduction

The full nonlinear EPMD source is reduced one order further. The real edge current is eliminated exactly, the five non-retained density modes are relaxed order by order, and the effective action is expanded through $O(R^6)$. On the unit residual metric,

$$\Gamma_{eff} = \frac{5}{3} R^2 - \frac{179}{1560} R^4 + \left(\frac{1635523}{702336960} - \frac{4987}{35116848} \cos 6\varphi \right) R^6 + O(R^8),$$

with angular coefficient

$$c_6 = -\frac{4987}{35116848} = -1.4201161 \times 10^{-4}.$$

An independent execution with the archived history-diagonal metric gives

$$\Gamma_{eff}^{hist} = 1.053014485 R^2 - 0.156158297 R^4 + (-0.006695457 - 0.007855703 \cos 6\varphi) R^6 + \dots$$

The sign of the angular coefficient is the same on both metrics. Its relative size is not: on the unit metric $|c_6|$ is about 6% of the isotropic R^6 coefficient, on the history metric about 117%, and the isotropic R^6 coefficient itself changes sign. The robustness claim is therefore restricted to the sign of c_6 , which is all that the orientation selection uses.

4.3 Origin of the sixfold term

The sixth-order term is generated by hidden-mode relaxation rather than inserted by symmetry. The E_1 Fold excites the alternating $m = 3$ wheel mode at cubic order. Eliminating that mode lowers the effective action by a contribution proportional to $R^6 \cos^2 3\varphi$, which contains the $\cos 6\varphi$ anisotropy. D_6 symmetry determines the functional form; the nonlinear source calculation determines the sign.

4.4 Amplitude and truncation

A negative $\cos 6\varphi$ coefficient selects angular minima at fixed nonzero amplitude R within the validity of the expansion. It does not by itself fix R , nor establish a stable nonzero radial solution: on the unit metric the truncated radial action has a negative quartic and a positive sextic term and the radial stationary point lies where the R^6 term is no longer small, while on the history metric the sextic coefficient is negative along directions that include the selected axes. A finite-order asymptotic polynomial is not the full action at large amplitude, and its radial behaviour establishes nothing about the physical Fold amplitude. Two consequences are drawn. First, the orientation selection is used only as an angular statement at fixed R , and no claim is made that the truncated Γ_{eff} determines the physical Fold amplitude. Second, the amplitude that enters the downstream history solve is not taken from Γ_{eff} at all; it is the separately returned E_2 history response of Section 10.3, whose normalisation is stated there. Control of the omitted $O(R^8)$ terms at the physical amplitude is an open obligation (Section 14.3).

5 Admissible Fold orientations and the two readings of indistinguishability

5.1 Sixth-order minima and the half-lazy gate

Because $c_6 < 0$, the local angular minima satisfy $\varphi = n\pi/3$: the six candidates are $0^\circ, 60^\circ, 120^\circ, 180^\circ, 240^\circ, 300^\circ$. The half-lazy generation law independently requires a non-degenerate reversal-odd carrier; the 0° and 180° extrema have $v = 0$ and erase the odd generation component. Conditional on the neutral generation intertwining of Section 3, the admissible set is

$$\mathcal{A}_F = \{60^\circ, 120^\circ, 240^\circ, 300^\circ\}.$$

5.2 The No-Subordinate-Distinguishability principle

VERSF treats the Fold as the minimal physical unit of distinguishability: its internal four-state architecture supports one irreducible committed binary outcome, not several independently readable outcomes [16]. The corresponding quotient rule is stated as follows.

Principle NSD. A primitive Fold carries exactly one irreducible committed distinction. A variable that distinguishes two mathematical realisations of the same Fold, but is not exported into a distinct record or defined relationally against an external physical frame, is not an additional physical degree of freedom of that Fold.

The principle does not deny the pre/post-commitment architecture of a Fold; it denies an additional hidden committed label inside that architecture. An orientation can become physical if another record stores a relative orientation; it is not automatically a primitive property of an isolated Fold.

Consistency with orientation-type quantities elsewhere in the programme. The charged-flavour curvature closure fixes a definite phase and axis (for example $\arg \zeta = 5\pi/6$ in the same-carrier bridge). NSD does not quotient those, because they are defined relationally: the phase is the angle between two carriers (the curvature image and the Wilson return) that are both present in the same record, so the relative orientation is a recorded fact. The Fold angle of the present section is the orientation of a single primitive coordinate with respect to a symmetry axis of its own carrier, with no second carrier in the record. This is the criterion applied uniformly: an orientation is physical when it is relative to a second recorded object, and unphysical when it is absolute with respect to an unrecorded frame. Any future construction that supplies such a second object for the Fold reopens Section 5.4.

5.3 Residual representative symmetry

The admissible set is preserved by $\varphi \rightarrow \varphi + \pi$, $\varphi \rightarrow -\varphi$ and $\varphi \rightarrow \pi - \varphi$. These exchange the four representatives without creating a record label. Under $Q(\cos\varphi, \sin\varphi) = (\cos 2\varphi, \sin 2\varphi)$,

$$Q(60^\circ) = Q(240^\circ) = q_+ = \left(-\frac{1}{2}, +\frac{\sqrt{3}}{2}\right), Q(120^\circ) = Q(300^\circ) = q_- = \left(-\frac{1}{2}, -\frac{\sqrt{3}}{2}\right).$$

Q is already invariant under $\varphi \rightarrow \varphi + \pi$; the remaining residual action on the source is the reflection $q_\pm \rightarrow q_\mp$, which flips the sign of the second E_2 component.

5.4 Orbit-invariant reading

Under NSD the physical state is the orbit $\{q_+, q_-\}$, and a physical observable is a function on the quotient, that is, a reflection-invariant function of q . Such a function may depend on q_y^2 without distinguishing the sign of q_y . On this reading the prediction rule is: evaluate the downstream calculation on either representative and require the result to be reflection-invariant. Section 11.2 returns $R_v = 0.05357926$ and 0.05361611 on q_+ and q_- ; the 7×10^{-4} relative difference between them is the measured departure of the implemented chain from exact reflection equivariance and is discussed in Section 11.2. The orbit-invariant value is $R_v \approx 0.0536$.

5.5 Projection hypothesis

A second reading takes the physical source to be the component of q that is invariant under the residual action, obtained by projection onto the reflection-fixed subspace:

$$Q_{phys} = \frac{1}{4} \sum_{\varphi \in \mathcal{A}_F} Q(\varphi) = \frac{q_+ + q_-}{2} = \left(-\frac{1}{2}, 0\right).$$

Write the residual reflection on the quadratic source as $\sigma(s, t) = (s, -t)$ and the fixed-subspace projection as $P(s, t) = (s, 0)$. NSD requires physical observables to be independent of the sign of the unrecorded coordinate; it does not require independence of its squared magnitude. A reflection-invariant scalar observable may therefore have the form $F(s, t) = f(s, t^2)$, and evaluating it at the projected source replaces $f(s, t^2)$ by $f(s, 0)$. That replacement is an additional physical prescription. For each observable to which the projection is applied, the identity required is

$$F(q) = F(Pq),$$

a factorisation of the response through P . Reflection invariance alone establishes $F(q) = F(\sigma q)$ and not this stronger identity. Two responses must be kept apart. The unprojected construction of Section 11.2 defines a trial response F_{trial} , evaluated on the representatives $q_{\pm}$ and returning $R_v \simeq 0.0536$. Hypothesis P does not assert that F_{trial} factors through P . The two reported constructions give different outputs, 0.0536 and 0.02822; that difference does not by itself establish or refute factorisation of a single unchanged response, since it would have to be shown that the generation rule and every other input were held fixed between the two evaluations, which has not yet been established. Hypothesis P asserts that the physical response is a different object, $F_{phys} = F \circ P$, obtained by evaluating the same downstream chain on the projected source, and the outstanding task is to derive from the source and record map why the physical construction uses F_{phys} rather than F_{trial} . For the full neutral construction a sufficient condition would be that its effective mass and kinetic data, with the physical current included, depend on the source only through Pq up to a consistent change of field basis; equality of a single mass ratio at one source value would be a weaker result and would not establish the factorisation (Appendix E).

Three further facts are stated so that the status of the operation is not overstated.

1. It changes an invariant magnitude. $|q_{\pm}|^2 = 1$ whereas $|Pq|^2 = 1/4$: the polynomial $F(s, t) = s^2 + t^2$ is reflection-invariant and takes the value 1 on both representatives but $1/4$ at their projection. Equality on the orbit does not imply equality at the projected mean. Averaging the observable gives $[F(q_+) + F(q_-)]/2 = F(q_+)$; averaging its argument gives $F(Pq_+)$. These are different operations.
2. It is the quadratic image of a bisector. Since $|Q(u, v)| = u^2 + v^2$, the vector $(-1/2, 0)$ can be represented as $Q(0, 1/\sqrt{2})$, a preimage at $\varphi = 90^\circ$ with $R^2 = 1/2$ lying on a bisector rather than a selected angular minimum. This is a warning about changing the source; it is not an instruction to substitute that Fold configuration into channels defined on the original Fold coordinate.
3. It interacts with the generation law. Projecting after refinement and projecting before it give different source sequences (Section 3.3), so the order used is part of the hypothesis.

The projection is accordingly recorded as an additional modelling assumption, Hypothesis P, beyond NSD:

Hypothesis P. The nonlinear downstream response of the neutral sector to an unrecorded-orientation Fold source factors through the reflection-fixed projection of the quadratic source, evaluated at the projected magnitude, with a generation prescription whose refine/project ordering and treatment of the scalar invariant r remain to be derived or recovered. The archived rerun fixes the numerical source vector actually inserted into the downstream history calculation, but does not by itself establish that upstream provenance (Section 3.3).

The obligation that would promote Hypothesis P from assumption to theorem is a derivation, from the reduced action or from the record map, of the factorisation $F(q) = F(Pq)$ for the neutral mass, kinetic and current data: that is, a proof that every coefficient of t^{2k} , $k \geq 1$, in the reflection-invariant expansion of the response vanishes, not merely that their sum cancels at one orbit. Until

that derivation exists, results obtained under Hypothesis P are conditional candidates in the sense of Section 1.2.

5.6 Prediction rule under the two readings

Under the orbit-invariant reading the retained-order prediction is the single-orientation D-free result of Section 11.2. Under Hypothesis P it is the projected-source result of Section 11.4. The paper reports both, grades both in Appendix C, and identifies the choice between them, which is the quotient operation itself, as the decisive open object of the neutrino sector at this order. For clarity, the projected value is placed on record as the frozen prediction of the declared projected-source construction, while its claim grade remains conditional; this does not assert that VERSF uniquely selects Hypothesis P over the orbit-invariant reading.

6 Same-wheel diagnostic and its rejection

6.1 Frozen diagnostic bridge

A deliberately simple same-wheel endpoint bridge tests whether the source-derived half/quarter law has enough algebraic structure to generate three masses. Its response matrices A_0 and B_0 are full rank and non-commuting and were frozen before the sixth-order angle selection was derived.

6.2 Spectrum on the selected orientations

Evaluating the diagnostic on any of the four admissible minima gives the Takagi values, up to a common scale,

$$(m_1, m_2, m_3) = (0.16735103, 0.27233946, 0.64858693),$$

$$m_1 : m_2 : m_3 = 1 : 1.62735454 : 3.87560756, R_\nu = 0.11756373.$$

The diagnostic has the same value on all four admissible orientations, so it is invariant on that discrete set. This establishes equality between those representative evaluations only. Its value on the projected source is not established by that equality. The archived rerun now evaluates the declared projected-source bridge numerically, but that calculation remains conditional on the projected-source and generation prescription whose upstream provenance is discussed in Sections 3.3, 5.5 and 14.2.

6.3 Verdict

The same-wheel value 0.11756 is strongly incompatible with the comparator (a factor of about four). The failure is that the first two masses remain too widely separated relative to the third. The correct inference is not that the sixth-order Fold selection failed; it is that the tested construction, under its stated assumptions, is rejected. The same-wheel pullback is the assumption most directly implicated, because it is the only ingredient of that construction that is not shared with the later chain, but the rejection does not by itself prove that no other assumption of the construction is responsible.

Test	Frozen output	Status
Equal-winding diagnostic	$R_\nu \approx 0.1638$	Failed; angle unselected at that stage.
Same-wheel on selected orientations	$R_\nu = 0.117564$	Failed; construction rejected.
D-free single orientation ($q_\pm$)	$R_\nu = 0.053579 / 0.053616$	Orbit-invariant reading; strongly incompatible (about 80% high).

Test	Frozen output	Status
D-free projected source (Hypothesis P)	$R_\nu = 0.0282214$	Conditional candidate; about 2σ low (provisional, Section 11.5).
Comparator (NuFIT 6.0, NO)	0.029805 ± 0.00079	Post-freeze comparison only.

7 Direct Dirac, sterile and kinetic responses

7.1 Canonical variation of the light operator

For $S = B_L - B_D B_R^{-1} B_D^T$ the first variation is

$$\delta S = \delta B_L - \delta B_D B_R^{-1} B_D^T - B_D B_R^{-1} \delta B_D^T + B_D B_R^{-1} \delta B_R B_R^{-1} B_D^T.$$

Kinetic whitening adds the corresponding δK_L terms. On the retained branch $B_L = 0$ (Section 12.1). A common rescaling $\delta B_D = a B_D$, $\delta B_R = b B_R$ and $\delta K_L = z K_L$ (the rescaling of the kinetic matrix itself, not of the identity, unless the basis is already canonical) then changes S_c only by the overall factor $(1 + 2a - b - z)$ at first order and cannot alter R_ν . Were B_L nonzero it would have to scale by the same factor $2a - b$ for the statement to hold. The needed response must therefore contain a non-scalar, generation-resolved component.

7.2 Direct completion route

On the 60° representative, both the linear lift and the complete relaxed second attachment have exactly zero local Higgs/completion component on the archived separable branch. The Fold is nonzero, but its response lies in density and real current rather than the local completion coordinate. The completion-mediated contributions to the returned B_D and B_R therefore vanish on this branch at the declared order. This zero is a statement about the selected configurations and the implemented response maps. A linear response in the Fold coordinate, a quadratic response in that coordinate, and a response linear in its E_2 image are distinct maps. A zero on a spanning set proves that a linear map vanishes on its whole domain, but the corresponding inference fails for a quadratic map evaluated on a finite set: vanishing of a second-order attachment on the four selected orientations constrains its symmetric tensor A to $A_{uv} = 0$ and $A_{uu} + 3A_{vv} = 0$, which does not force $A = 0$ (Appendix F). The full second-order attachment tensor, including any scalar quadratic channel, must therefore be evaluated on the projected source or shown to factor through the projection. If the retained response is known independently to be linear in the E_2 image, $C_2(x) = \mathcal{L}(Q(x))$, then vanishing on q_+ and q_- , which span E_2 , does give $\mathcal{L} = 0$ and hence $\mathcal{L}(Pq) = 0$; this sufficient condition requires the pure- E_2 factorisation to be established from the actual source tensor first.

7.3 Current-cycle route

The Fold also generates a nonzero real current, but the selected Fold current is orthogonal, to numerical roundoff, to the divergence-free neutral current-cycle carrier. The faithful current-history audit independently returns a rank-zero source pullback where the physical current-history target has rank six. The retained current-cycle channel therefore does not supply the required Fold-to-neutral correction. As in Section 7.2, the statement is established on the returned configurations; its extension to the projected source requires either the factorisation or a direct evaluation.

7.4 Protected kinetic and active-Majorana routes

The evaluated defect/history/bath/commitment kinetic corrections on the protected common line are zero. A direct generation-dependent fermion residue remains logically possible but has not been returned. If the written fermion action is truncated to its Dirac and sterile quadratic terms, eliminating

bosonic auxiliaries at tree level does not generate a new active quadratic Majorana coefficient; an irreducible two-interface active term remains a distinct unreturned object (Section 12.1).

7.5 Declared-order verdict

The retained source terms give $Q_D^{returned} = 0$ and $Q_R^{returned} = 0$ on the configurations for which they have been explicitly returned, through the implemented completion and current routes, with the truncated active block and protected kinetic corrections also returning zero. Their extension to Hypothesis P is conditional on the factorisation or direct evaluation of the relevant completion and current maps. The absence of an independent active Majorana term in the explicitly truncated quadratic action (Section 12.1) is a separate retained-order statement and does not rely on this extension. The same-wheel hierarchy cannot be repaired by further differentiation of these returned scalar channels.

8 Mixed fourth-order Fold-fermion vertex

8.1 Differential order

The primitive MASD matter term is linear in the completion coordinate X and bilinear in fermions, so its first nonzero completion-matter interaction is a mixed third derivative $T_{X\psi\psi}$ with one bosonic leg and two fermion legs. A direct Fold repair of the form $\psi\psi x^2$ requires a mixed fourth-order source tensor with two Fold legs and two fermion legs.

The certified CAR/Fock gauge and completion jets each carry one primitive source leg. Because the selected EPMD Fold has zero second-order projection into both the relevant link/bond and local-completion source legs on the returned branch, the composed two-Fold/two-fermion vertex is zero by the chain rule on the declared stack. A nonzero fourth-order term would have to be a new or previously unreturned primitive source dependence, not a hidden consequence of the returned third-order leg. As in Section 7, this zero is established on the returned Fold configurations and its extension to the projected source carries the same qualification.

8.2 Symmetry-allowed generation tensors

For the three-generation permutation carrier $P = 1 \oplus E_2$, the Hermitian generation-matrix space contains three E_2 copies: diagonal traceless splitting, symmetric off-diagonal mixing, and antisymmetric off-diagonal current/CP structure. For a symmetric Majorana mass correction the first two are the mass-producing copies.

8.3 Post-failure inverse diagnostic

Using the measured hierarchy only after the same-wheel failure, an inverse diagnostic asks what missing matrix correction the same-wheel spectrum would need. This is a tensor census for a future source term, not a prediction.

Allowed correction class	Best or required R_v behaviour	Interpretation
Diagonal E_2 only	minimum $R_v \approx 0.09239$	Insufficient.
Symmetric off-diagonal E_2 only	minimum $R_v \approx 0.08571$	Insufficient.
Both real E_2 copies	reaches ≈ 0.0298	Requires both mass-shaping channels.
A_1 + both E_2 copies	reaches ≈ 0.0298 with smaller norm	Common terms assist but do not replace E_2 structure.

The efficient correction needs both diagonal generation splitting and symmetric off-diagonal mixing. Section 10.5 shows that the universal completion response excites both automatically.

9 CFB-2 capacity test

9.1 Why CFB-2 is relevant but not physical

The Master-Action corpus contains a candidate generation-resolved scalar-stretch law, MA9'. Its CFB-2 execution freezes a positive completion resolvent $\mathcal{R}_f = (I + L_f^2)^{-1}$ with integer-weight three-generation path Laplacians; the neutrino candidate uses path weights $(2, 3)$, declared before its spectrum was inspected. This makes CFB-2 a legitimate capacity test. It does not make the graph source-derived: HMGBC requires one universal history generation object rather than a separately chosen neutrino graph.

9.2 CFB-2 neutrino ratio

For the neutrino graph the completion operator has eigenvalues

$$\text{spec}(\mathcal{R}_f^{(v)}) = (1, 0.1528470774, 0.01681873236).$$

On the CFB-2 neutral branch the Dirac and sterile blocks are proportional to the same Y_ν , so each generation shares the same two-by-two active/sterile factor and the three active-dominant light masses are proportional to the three eigenvalues (Appendix B). The frozen branch returns

$$(m_1, m_2, m_3) \propto (0.00377992384, 0.03435159675, 0.22474487139), R_\nu = 0.02308589,$$

about 22.5% below the comparator. Because the $(2, 3)$ graph is a constitutive choice rather than a history-derived object, the result is classified as capacity, not prediction.

9.3 Half-lazy shape diagnostic

For a three-node weighted path $L(a, b)$, requiring the centred half-lazy profile $(1, 1/2, 1/4)$ to be an eigenmode of the path Laplacian gives the edge ratio $b/a = 8/5 = 1.6$ (Appendix D). CFB-2 used $b/a = 3/2$. The proximity is noted, but the identification of the half-lazy depth profile with an MA9' path-Laplacian eigenmode is not a theorem of the physical source and is not promoted.

10 Universal history-derived completion geometry

10.1 The universal HMGBC object

Stationary history flow defines one symmetric conductance geometry, pushed to the terminal-generation carrier and converted into the universal covariant Laplacian Δ_Ω . HMGBC does not select independent quark, charged-lepton and neutrino graphs. Completion is determined by

$$\mathcal{R}_\Omega = (W_H + \Delta_\Omega)^{-1} W_H,$$

and this is the object on which the Fold source acts.

10.2 Same-history background conductance

Using the irreversible one-bit history rather than the older reversible benchmark, the antipodal three-generation quotient returns equal background conductances

$$g_{AB} = g_{BC} = g_{CA} = 0.0990153342003,$$

about 15.1% above the reversible value 8/93. The background is circulant, so relative generation splitting must come from a non-uniform source response, not from a declared neutrino graph.

10.3 Source-linked E_2 amplitude and its normalisation

An independent irreversible-history third-jet execution returns a nonzero $E_1^2 \rightarrow E_2$ history response. In the declared history normalisation of [7] (the inner product on the history carrier against which the E_2 harmonic is normalised; its explicit definition is contained in the recovered basis and normalisation record of Section 14.2), the relaxed response magnitude is

$$|q_{E_2}| = 0.1663661412.$$

The source vector inserted into the history solve is $q_{E_2} Q$, where Q is the quadratic source image of Section 5.3 or 5.5 in the E_2 harmonic basis. For a single representative $|Q| = 1$ and the inserted magnitude is 0.1663661412; for the projected source $|Q| = 1/2$ and the inserted magnitude is 0.0831830706. The accompanying carrier audit establishes norm-preserving identifications of the relevant harmonics on the declared history carrier. The same-action nonlinear physical certificate for the amplitude remains a separate source-provenance obligation (Section 14.3), and the finite-amplitude response is evaluated in Section 12.6, which finds $Rv/|qE_2| \approx 0.1696$ across the tested range. The remaining issue is the physical provenance of the amplitude, not its numerical sensitivity.

10.4 Source images of the four admissible representatives

Because Q doubles the Fold angle, the four admissible representatives map into two E_2 directions,

$$60^\circ, 240^\circ \rightarrow 120^\circ \text{ in } E_2; 120^\circ, 300^\circ \rightarrow 240^\circ \text{ in } E_2.$$

A reversed E_2 vector (overall minus sign) is a useful robustness diagnostic but is not the quadratic image of any sixth-order minimum and is not treated as a source-selected branch.

10.5 Exact nonlinear universal completion response

The source amplitude is propagated through the full exponential history family at finite amplitude. The stationary distribution, terminal conductances, W_H and universal completion resolvent are recomputed rather than truncated at first order. For the $120^\circ E_2$ representative this gives

$$(g_{AB}, g_{BC}, g_{CA}) = (0.09897769, 0.11042009, 0.08286056).$$

The universal response excites both symmetric E_2 mass-shaping channels of Section 8.2, diagonal traceless splitting and symmetric off-diagonal mixing, in relative magnitude about 1.46 to 1.47 on the finite source response, with no neutrino datum used.

10.6 Archived full-neutral branch remains rejected

The same source-linked history deformation inserted into the older MHPE/FCHI six-state construction leaves its active-dominant gap ratio of order 2×10^{-5} . The conductance response does not rescue that architecture, which motivates the D-free calculation below.

11 D-free representation-resolved neutral compression

11.1 Cancellation of the sterile address radius: scope

Write the joint neutral geometry as

$$M_D = a X H_N, M_R = b H_N^T P_N H_N,$$

where H_N is the positive right-neutral address radius. For invertible H_N and P_N ,

$$M_D M_R^{-1} M_D^T = \frac{a^2}{b} X P_N^{-1} X^T.$$

Every invertible generation-dependent factor carried purely on the sterile right address cancels from the leading zero-momentum Schur complement. This is why the earlier depth matrix $\text{diag}(14^{-2}, 14^{-1}, 1)$ is not retained as an independent neutrino input ("D-free" throughout means free of that matrix). A derived active-side factor would survive through X or K_L and would have to be included.

The identity is exact for the zero-momentum Schur term and proves cancellation from the leading light operator only. Passing to the sterile coordinate $n = H_N^{-1} n'$ removes H_N from the mass blocks but transforms the sterile kinetic metric to

$$K_R' = H_N^{-\dagger} K_R H_N^{-1},$$

and keeping this transformed metric is necessary for an equivalent physical parent; the identity does not justify setting it to the identity. The exact six-state spectrum reported below applies to the archived parent and its declared kinetic metric, which the reproducibility package must state explicitly. Agreement between the Schur and Takagi results for that parent (Section 11.4) confirms the consistency of that one parent. It does not establish independence from other physically admissible sterile parents: the light Schur mass can be unchanged while the exact light mass and active weight change materially (Appendix G). The relevant remaining condition is source uniqueness of the parent, or a demonstrated insensitivity of the observable across the allowed family (Section 14.3).

11.2 Single-orientation D-free results and the equivalence test

Using the exact nonlinear history response and the typed active/sterile subspaces, with no depth matrix, the two quadratic images give

$$q_+ = Q(120^\circ_{E_2}): m_1 : m_2 : m_3 = 1 : 1.09041985 : 2.12785637, R_\nu = 0.05357926,$$

$$q_- = Q(240^\circ_{E_2}): m_1 : m_2 : m_3 = 1 : 1.09039471 : 2.12704615, R_\nu = 0.05361611.$$

Each result occurs twice among the four Fold representatives. The active weights of the three light states are approximately (0.99715, 0.99711, 0.99682), where the active weight of a mass state is defined as $w_i = \|\Pi_{act} U e_i\|^2$, with U the Takagi unitary of the whitened six-state operator and Π_{act} the projector onto the typed active subspace; this is a basis-typed quantity and becomes the physical weak-current weight only once the current map of Section 14.3 is returned. In a controlled seesaw regime the leading departure of w_i from unity is quadratic in the dimensionless active-sterile mixing amplitude $T = M_R^{-1} M_D^T$ formed from the canonically normalised blocks (Appendix G), and a change of source can change T linearly. The leading Schur value and the exact six-state value agree at the 10^{-5} level in R_ν on both images.

If q_+ and q_- describe the same physical state, the two spectra should coincide once the entire construction (history geometry, representation embeddings, sterile pairing, kinetic and current maps) is transformed together by the reflection. They differ at 7×10^{-4} relative. That difference has one identified contributor, the E_2 response in the γ_T metric row whose microscopic origin is unresolved (Section 14.1), but holding γ_T at its symmetric value does not remove it, so the residual is either a broken equivariance in the implementation, an approximation, or a genuine relational frame. Which of these it is has not been determined. The required test is the consistent-basis test of Appendix I: the reflection must be realised as an invertible matrix G on the typed carrier, supplied by the actual reflection action and not fitted, under which the mass coefficient, kinetic metric and current map transform together; equivalent data then have identical Takagi values and identical canonical current amplitudes, and the dimensionless congruence residuals of Appendix I measure any failure. A small change in the final ratio after freezing one metric coefficient is not a substitute for these

transformation checks. Under the orbit-invariant reading of Section 5.4 the retained-order prediction is the common value $R_\nu \approx 0.0536$, and this equivalence test is part of its certificate.

11.3 Projected source through the nonlinear history

Under Hypothesis P the projected Fold source is $Q_{\text{phys}} = (-1/2, 0)$. The recovered execution record shows the numerical insertion actually used by the downstream run: after whitening the entropy-constrained, Schur-relaxed even-history metric to (radial, E2-cos, E2-sin), the source vector is $0.16637(0, -1/2, 0)$, of magnitude 0.08318. This closes the computational ambiguity about the inserted vector. It does not, however, determine the upstream refine/project ordering or whether the scalar invariant r was retained; those remain provenance obligations of Hypothesis P. The nonlinear history is then recomputed from the archived irreversible one-bit kernel rather than inferred from the printed output.

$$(g_{AB}, g_{BC}, g_{CA}) = (0.09242420, 0.10680687, 0.09673587), w_H = 0.1523387853, \gamma_T = 0.5371166018.$$

No neutrino observable enters this projection or the history solve. The generation splitting of the conductances is roughly half that of Section 10.5, as expected from the halved source magnitude; the change of direction (from the $120^\circ E_2$ axis to the fixed axis) changes which generation pair is split most.

11.4 Projected-source spectrum

The exact D-free six-state neutral operator returns the three active-dominant masses, up to the unresolved common scale, in the ratio

$$m_1 : m_2 : m_3 = 1 : 1.0468 : 2.0967, R_\nu = 0.02822,$$

with active weights $(0.99714849, 0.99712739, 0.99682513)$. The leading Schur value is $R_\nu = 0.02821131$; the difference of 1.012×10^{-5} is absolute, about 0.036% relative, so the exact six-state result is stable against finite sterile mixing on this branch. Headline statements in this paper use $R_\nu \approx 0.02822$; full precision is retained in the tables for reproducibility, without any implication of comparable theoretical accuracy.

The similarity of the active weights across Sections 11.2 and 11.4 is a numerical observation for the declared parent. Their leading departure from unity is governed by the dimensionless active-sterile mixing matrix formed from the canonically normalised Dirac and sterile blocks, not by the light-mass scale a^2/b alone (Appendix G). No statement about the order at which the weights depend on the source deformation is inferred from the closeness of the printed values; that would require an explicit symmetry or derivative calculation.

11.5 Post-freeze statistical comparison

Only after the projected result is frozen is the comparator opened. Against $R_\nu^{\text{ref}} = 0.0298050$ with the local $\sigma_R \approx 7.92 \times 10^{-4}$ of Section 1.3, the projected result $R_\nu = 0.02822143$ has a signed residual of -1.998 local standard deviations on the propagated ratio (a 5.3% deficit); profiling the common scale in a two-observable Gaussian comparison gives -2.008 (Appendix H). Both are approximately two standard deviations. This is a description of the residual under the stated approximation, not a declaration of statistical acceptance or a full oscillation fit.

The single-orientation value 0.0536 and the same-wheel value 0.1176 are far from the benchmark and are strongly incompatible with it within their retained-order constructions. Local uncertainties are not extrapolated into calibrated tail probabilities for them. These results reject those implementations under their assumptions; they do not establish a no-go theorem for every VERSF neutrino construction. For the projected result the statement made is this and no more: it has an approximately

two-standard-deviation local discrepancy (signed residual -1.998 propagated, -2.008 profiled, two-sided $p \approx 0.045$ under the Gaussian approximation), and no formal acceptance or exclusion claim is made without the full likelihood and the completed physical model. What distinguishes it from the reachability scans and fitted inverse diagnostics is only that no continuous coefficient was adjusted to obtain it.

11.6 Claim level

Under the orbit-invariant reading the single-orientation construction is **FALSIFIED AT THIS ORDER**: its retained-order value 0.0536 is incompatible with the selected benchmark, with its source equivariance still unresolved. Under Hypothesis P the projected calculation is a **CONDITIONAL CANDIDATE** with an approximately two-sigma local residual, dependent on Hypothesis P, its generation prescription and the declared source and parent normalisation. The choice of projection was made after the earlier disagreement was known; the calculation contains no explicit fit of a continuous coefficient to the neutrino target, but it was not designed under blinding. A later derivation of an external record that physically distinguishes the Fold orientation would reopen Section 5.4 and must be treated upstream, not selected after seeing neutrino data. The projected value is nevertheless placed on record as the frozen prediction of that declared construction for future experimental testing; its conditional grade is unchanged because the projection, generation provenance and source normalisation are not yet uniquely derived from VERSF.

12 Retained-order corrections and the sensitivity of the ratio

12.1 Irreducible active Majorana block

The primitive active Majorana block is zero at bare single-interface order. The allowed local two-interface coefficient is

$$\mathcal{L}_5 = -\frac{1}{2} \kappa_{ij}^{irr} (L_i H)(L_j H) + h.c., B_L^{irr} = \frac{v^2}{2} \kappa^{irr}.$$

In the explicitly retained Dirac-plus-sterile quadratic action there is no independent all-fields $1\text{PI } LLHH$ coefficient. Bosonic stationary relaxation begins at $O(\psi^2)$ and changes the action first at $O(\psi^4)$; it cannot manufacture a new quadratic active mass. Therefore $B_L^{irr} = 0$ at the retained tree order. The sterile-exchange term $-B_D B_R^{-1} B_D^T$ is already included and is not counted again.

12.2 Active kinetic residue

The primitive shared-link CAR/Fock kinetic operator uses the same faithful physical link and carries no independent matter-link normalisation. The three generation copies of one representation therefore receive no independent tree-level left kinetic coefficient from that term. Known defect/bath residue corrections in the audited channels are zero. A generation-resolved physical self-energy remains a higher-order return, not a retained-order free matrix.

12.3 Dimension-six matching is contained in the exact parent

Integrating out heavy neutral states produces the dimension-six active kinetic operator of an effective description. It is not an additional correction: the exact six-state Takagi diagonalisation already contains the full active-sterile mixing that generates it, which is what the Schur/Takagi agreement of Section 11.4 records.

12.4 Sensitivity of R_ν to small mass shifts

Let $A = m_2^2 - m_1^2$, $B = m_3^2 - m_1^2$ and $R_\nu = A/B$. For fractional mass shifts $\varepsilon_i = \delta m_i/m_i$ the first-order variation is

$$\frac{\delta R_\nu}{R_\nu} = c_1 \varepsilon_1 + c_2 \varepsilon_2 + c_3 \varepsilon_3, c_1 = -\frac{2m_1^2}{A} + \frac{2m_1^2}{B}, c_2 = \frac{2m_2^2}{A}, c_3 = -\frac{2m_3^2}{B}.$$

On the projected spectrum

$$(c_1, c_2, c_3) = (-20.27869, 22.86761, -2.58891).$$

They sum to zero, as they must, because a common mass rescaling leaves R_ν unchanged. Differential corrections to the nearly degenerate lower pair can nevertheless produce an amplified change in the ratio: increasing m_2 alone by 0.17% changes R_ν from 0.02822143 to 0.02931947, a 3.89% increase, and with the other two masses fixed the increase needed to reach the comparator is 0.24508%. These are diagnostics of sensitivity, not source-selected corrections.

12.5 Ordinary radiative corrections are neither excluded nor implicated by size

The D-free branch has $\|Y_\nu\|_2 \simeq 0.314$ in the normalisation in which $M_D = \nu Y_\nu/\sqrt{2}$ with the common heavy scale set by b ; the basic loop factor is $\|Y_\nu\|_2^2/(16\pi^2) \simeq 6.24 \times 10^{-4}$, and with the declared heavy/electroweak logarithm the estimated two-sided flavour deformation of the mass matrix is of order 1.7×10^{-3} per order-one coefficient. A mass correction of that order is not excluded as relevant to the residual. Its effect on the ratio depends entirely on its matrix direction: it can increase, decrease or leave the ratio unchanged, and for a correction proportional to the mass matrix itself the ratio shift is exactly zero. In the illustrative pattern in which only m_2 shifts, the displacement needed to reach the comparator is about 1.44 times 1.7×10^{-3} ; this number is not a derived loop coefficient.

The retained-order conclusion is therefore that the residual cannot be assigned to, or ruled out as arising from, ordinary loops on a size estimate alone. A source-derived mass correction δB and kinetic correction δK must be canonically normalised and propagated through the six-state spectrum, with the light states identified through their currents; Appendix H gives the first-order formula $\delta m_i = \text{Re}(u_i^T \delta M u_i)$ with $\delta M = \delta B - \frac{1}{2} \delta K^T M - \frac{1}{2} M \delta K$, which tests any returned loop matrix without fitting its direction to the residual. The physical Yukawa normalisation, sterile thresholds and scale entering the logarithm must be specified on the same branch; inferring a light absolute scale from an observed splitting (Section 13) does not by itself determine those heavy-sector quantities. The source-amplitude sensitivity is now evaluated directly in Section 12.6: the scan is approximately proportional, $R_\nu \approx 0.1696|qE2|$ to within about 1% over the tested range, with an effective finite-amplitude slope of about 0.17. The remaining open question is the physical derivation of $|qE2|$.

12.6 Final normalisation attack: amplitude response and the action/Fisher boundary

The projected-source chain was independently rerun from the archived Generation History to Neutrino Response and MASD/HMGBC source packages. The full machine-precision reproduction values are recorded in Section 14.2. At ordinary reporting precision the rerun gives $m_1:m_2:m_3 = 1:1.0468:2.0967$ and $R_\nu = 0.02822$. The printed hierarchy is therefore computationally reproducible from the archived inputs rather than being only a reported downstream number.

The source-amplitude scan reveals a more useful result than “robustness.” For the four tested amplitudes, $R_\nu/|q| = 0.16963, 0.16964, 0.16970$ and 0.1684 . Thus, over nearly a factor of three in $|q|$, the frozen calculation is approximately proportional to the source amplitude to within about 1%: $R_\nu \approx 0.1696|q|$. The corresponding finite-amplitude secant slopes are close to 0.17, so the previously listed source-amplitude-sensitivity obligation is numerically discharged at the accuracy of

this scan. After the prediction is frozen, the stated oscillation comparator corresponds to $|q| \approx 0.1756$, about 5.6% above the archived amplitude 0.1664. The residual is therefore principally an amplitude-normalisation discrepancy: the open action-to-history conversion is quantitatively a 5–6% question on this branch.

The two nearby source-side amplitudes used in the scan have independent provenance. In the Primitive-Matter Selection programme, the primitive nonlinear third-jet calculation gives $g_{ax} = 1/2$ (Theorem 32.62 in [4]). PMS-EX97 returns the canonical master-Hessian E2 stiffness $\kappa E2 = 3$, and PMS-EX98–EX99 then give the clock-invariant susceptibility $g_{ax}/\kappa E2 = 1/6$. A later type-correct lift of the same irreversible history law returns $\kappa E2 = 2.92415$, giving $g_{ax}/\kappa E2 = 0.17099$ (PMS-EX106). These are source/history diagnostics, not values chosen from the neutrino comparator.

The final obvious shortcut is rejected on typing grounds independently of its numerical effect. The primitive source programme derives a positive rank-six current action and an exact current-link Schur reciprocity, but the action Hessian and the history/Fisher metric live on differently typed objects; no canonical map has been derived that identifies their absolute stiffnesses. Directly equating them is therefore not an admissible normalisation rule. If that inadmissible identification is nevertheless evaluated post-freeze, it gives $|q| \approx 0.05885$ and $R_v \approx 0.00991$. The numerical disagreement illustrates the consequence of the typing error; it is not the basis for rejecting the shortcut. The outstanding quantity is one absolute source-to-history/current normalisation, not a freely adjustable neutrino coefficient.

13 Absolute-scale consequence and cosmological comparison

13.1 One measured splitting fixes the scale

Both R_v and the full mass ratios are scale free. Write $m_i = s r_i$, where the dimensionless shape r_i is fixed by a particular construction. Admitting the measured atmospheric splitting after fixing that shape gives

$$s = \sqrt{\frac{\Delta m_{31}^2}{r_3^2 - r_1^2}}, \sum_i m_i = \sqrt{\Delta m_{31}^2} \frac{r_1 + r_2 + r_3}{\sqrt{r_3^2 - r_1^2}}, \Delta m_{21}^2 = R_v \Delta m_{31}^2.$$

These are consequences of a predicted shape calibrated by one measured splitting, a post-freeze operation permitted by Section 1.3; they are not absolute masses derived from the VERSF action alone, and R_v by itself does not determine the sum, since the full shape is needed. With $\Delta m_{31}^2 = 2.513 \times 10^{-3} \text{ eV}^2$ the constructions of this paper give the following. Both single-orientation branches are listed, since they differ at the displayed precision. The CFB-2 row is a capacity diagnostic.

Construction	m_1 (eV)	m_2 (eV)	m_3 (eV)	$\sum m_\nu$ (eV)	Solar splitting (10^{-5} eV^2)
Same-wheel	0.013388	0.021787	0.051887	0.087062	29.544
D-free single orientation, q_+	0.026690	0.029103	0.056792	0.112585	13.464
D-free single orientation, q_-	0.026703	0.029117	0.056798	0.112618	13.474
Projected source (Hypothesis P)	0.027202	0.028476	0.057035	0.112713	7.092
CFB-2 capacity diagnostic	0.000843	0.007663	0.050137	0.058643	5.801

Construction	m_1 (eV)	m_2 (eV)	m_3 (eV)	$\sum m_\nu$ (eV)	Solar splitting (10^{-5} eV ²)
Measured solar benchmark (NuFIT 6.0)					7.49

The projected shape has a close lower pair and a third mass about 2.1 times the lightest; describing all three as quasi-degenerate would be imprecise. Neither the same-wheel nor the CFB-2 shape has a sum near 0.11 eV. The table compares distinct constructions and is not a list of interchangeable physical predictions.

13.2 Cosmological comparison

The DESI DR2 BAO plus CMB analysis reports $\sum m_\nu < 0.0642$ eV at 95% under its baseline Λ CDM model with three degenerate neutrino states, and a lightest-mass bound of 0.023 eV when oscillation constraints are included; its baseline evolving-dark-energy (w_0 w_a CDM) analysis instead gives $\sum m_\nu < 0.163$ eV [21]. These limits are conditional on the analysis model, data combination and priors.

The projected and single-orientation sums near 0.113 eV (both branches) exceed the former benchmark, by a factor of about 1.75, and lie below the latter. The CFB-2 diagnostic sum of 0.0586 eV does not exceed 0.0642 eV, although its solar splitting disagrees with the benchmark. A statement that every construction is excluded by the same sum bound would therefore be incorrect, and a quantitative exclusion claim for any of them requires a likelihood analysis for the actual spectrum rather than a comparison of central values with a headline bound.

13.3 What the comparison does and does not constrain

At fixed shape and fixed atmospheric splitting, the inferred sum cannot be lowered by changing only the common scale. A cosmological discrepancy must instead be addressed through a derived change of shape, a justified change in cosmological assumptions, or rejection of the tested construction. The present sum is not immune to future source-derived corrections to the shape.

Applying a three-neutrino cosmological constraint to the complete six-state parent also requires the sterile masses, production history and abundances, whose cosmological effects cannot be inferred from the three active weights alone. The present comparison is a benchmark test under standard three-light-neutrino abundance assumptions, not a cosmological fit of VERSF.

A beta-decay prediction requires the physical electron current. In the standard unitary three-light-state limit $m_\beta^2 = \sum_{i=1}^3 |U_{ei}|^2 m_i^2$, and mass ratios alone do not determine the weights $|U_{ei}|^2$. Until the charged-lepton frame and current are returned, no value of m_β is claimed as a VERSF prediction; a value obtained by inserting measured mixing angles would be a further post-calibration consequence and would be labelled as such. A future measurement of $\sum m_\nu$ near the inferred sum would support this particular shape without uniquely validating its mechanism.

14 Status ledger, reproducibility package and open gates

14.1 Completion-metric provenance audit

The generation-history source audit contains one unresolved equivariance issue: the γ_T metric row has an E_2 response whose microscopic origin is not identified, whereas the w_H scalar row respects the expected symmetry. Numerically, holding γ_T at its symmetric-background value while retaining the projected conductances changes the result only from $R_\nu = 0.02822143$ to 0.02822217; the

hierarchy is not produced by that anomaly, but the anomaly is one candidate contributor to the q_+/q_- non-equivalence of Section 11.2 and must be closed. Freezing this coefficient is not a substitute for the consistent-basis test of Appendix I.

14.2 Reproducibility package and independent rerun

The central projected-source calculation has now been independently reconstructed from the archived execution packages. The recovered files provide the irreversible history kernel, stationary measure, defect generators, constrained relaxation and whitening, the numerical projected-source insertion actually used by the run, the universal completion data, the typed neutral embeddings and the six-state parent. The rerun reproduces the quoted conductances, completion coefficients, active weights, mass ratios and R_v to the displayed precision. Numerical reproduction of this declared parent is therefore closed. The archive does not, by itself, settle the upstream refine/project ordering or the treatment of the scalar invariant r ; those remain physical/provenance obligations of Hypothesis P.

1. Recovered projected-source input: the downstream run inserts $0.1663661412(0, -1/2, 0)$ in the whitened even-history coordinates (radial, E_2 -cos, E_2 -sin). No separate hidden generation coefficient is introduced at this insertion stage. This fixes what was executed, not which upstream refine/project prescription or scalar-invariant rule uniquely derives it.
2. Recovered history evolution: the irreversible one-bit transition family, stationary-distribution solve, entropy-constrained defect relaxation and terminal three-generation conductance push-forward are executable from the archive.
3. Recovered source normalisation and basis: the E_2 harmonic carrier and whitening are explicit; the lift satisfies an identity-normalisation residual of 3.96×10^{-16} .
4. Recovered nonlinear completion output for Q_{phys} : $(g_{AB}, g_{BC}, g_{CA}) = (0.0924241994, 0.1068068739, 0.0967358714)$, $w_H = 0.152338785322$ and $\gamma_T = 0.537116601784$.
5. Recovered typed active/sterile embeddings and archived parent: the level-3 and level-4 address embeddings are isometric in the executed package. Their physical uniqueness and full current/kinetic admission remain open.
6. Recovered six-state calculation: at level 4 the exact parent gives active weights $(0.99714849, 0.99712739, 0.99682513)$, mass ratios $1:1.0468248785:2.0966838931$ and $R_v = 0.02822142937$.
7. Recovered Schur/Takagi consistency: the leading light ratio and exact six-state ratio differ only at about 10^{-5} in R_v on this parent, so the retained hierarchy is not generated by finite sterile mixing.
8. Still open: the upstream projected-generation provenance (refine/project order and treatment of the scalar invariant r), the consistent-basis reflection test for the two single-orientation representatives, and the source-unique physical sterile pairing/current map.
9. Amplitude-response execution: the tested points give $R_v/|q|$ near 0.1696 and are consistent with proportionality to within about 1% across nearly a factor-three amplitude range; the effective finite-amplitude slope is about 0.17. The comparator corresponds to $|q| \approx 0.1756$, about 5.6% above the archived amplitude. The source-amplitude sensitivity obligation is therefore numerically closed at this accuracy, while the absolute normalisation remains open.

Accordingly, numerical reproduction is no longer an open gate for the projected-source branch. The remaining gates concern physical source admission, the projection hypothesis, the action-to-history normalisation, the full kinetic/current map and the absolute neutrino mass scale.

Two variations are distinguished. A coordinate transformation of one parent, with its kinetic metric and currents transformed, must preserve physical results. A change to a different physical sterile parent need not preserve them; source uniqueness of the parent, or a demonstrated insensitivity of

the observable across the allowed family, is the relevant remaining condition, and independence from every arbitrary sterile parent is not a requirement.

14.3 Open gates

Object	Current status	Consequence
Half/quarter generation law	Derived conditional on neutral generation intertwining	No three independent generation coefficients.
Sixth-order Fold minima	Derived by $\cos 6\phi$ anisotropy; four survive the half-lazy gate	Continuous angle scan closed; amplitude and $O(R^8)$ control open.
NSD principle	Foundational; identifies the four representatives as unrecorded realisations of one orbit	Factorisation through the projection is a separate hypothesis.
Quotient operation (orbit invariant versus Hypothesis P)	Not derived; decisive open object	Decides between 0.0536 (incompatible with the benchmark) and 0.0282 (conditional candidate).
Projected generation rule	Downstream numerical insertion recovered; upstream refine/project ordering and treatment of scalar invariant r not yet derived or uniquely reconstructed.	Hypothesis P remains conditional on that provenance; numerical reproduction does not settle it.
q_+ / q_- equivalence	7×10^{-4} relative non-equivalence unexplained; consistent-basis test not yet run	Source equivariance unresolved.
Direct zero responses	Established only on explicitly returned configurations	Extension to the projected source needs a factorisation theorem or direct evaluation.
Source amplitude $ q_{E_2} $	Declared history normalisation; finite-amplitude response measured as $R_v/ q \approx 0.1696$; same-action physical certificate open.	Comparator corresponds to $ q \approx 0.1756$, about 5.6% above the archived amplitude; sensitivity is closed, normalisation provenance is not.
Same-wheel spectrum	$R_v = 0.117564$, about four times the comparator	Construction rejected under its assumptions; not a no-go for the sector.
Sterile parent	Radius cancels from the leading Schur term; equivalent full parents require transformed kinetics and currents	Source uniqueness or family insensitivity to be shown.
Retained B_L^{irr}	Zero in explicit tree-level truncation	No extra active Majorana repair at this order.
Generation-dependent K_L	Absent from primitive retained kinetic term	Higher self-energy return only.
Radiative effects	Potentially relevant in size; sign and magnitude of the ratio shift require the correction matrix	Must be computed from source before the residual is interpreted.

Object	Current status	Consequence
Mass sum	About 0.1127 eV for the projected shape after calibration by the atmospheric splitting	Above the Λ CDM benchmark; within the w_0 w_a CDM benchmark; model dependent.
Reproduction package	Central projected-source parent independently reproduced from the archived execution packages.	Numerical reproduction closed; remaining obligations are physical/provenance items listed in Sections 14.2–14.3.
Typed pairing, residues, current map, PMNS	Open	Needed for the full neutrino-sector certificate.

15 Conclusion

The nonlinear source removes the continuous Fold-angle freedom and leaves four admissible sixth-order orientations after the half-lazy non-degeneracy gate. The same-wheel bridge, frozen before any selection, gives $R_\nu = 0.1176$ and is rejected under its assumptions. The universal history-derived completion geometry with the sterile depth matrix removed gives $R_\nu \simeq 0.0536$ on either quadratic source image, incompatible with the selected oscillation benchmark at retained order. The projected-source construction gives

$$m_1 : m_2 : m_3 = 1 : 1.0468 : 2.0967, R_\nu \simeq 0.02822,$$

For the declared projected-source construction, the frozen numerical prediction is $R_\nu = 0.02822$, with an approximately two-sigma local residual against the stated NuFIT 6.0 normal-ordering comparator. Because Hypothesis P was introduced after earlier constructions were known to fail, this is not claimed as a fully blind prediction; it is the fixed prediction of the construction actually tested in this paper.

The unresolved issue is whether the physical source and record map justify the projection and its generation provenance. Indistinguishability of the source representatives alone does not provide that derivation: the No-Subordinate-Distinguishability principle makes the four orientations one orbit, but averaging the source vectors before a nonlinear response is a further step that halves an invariant magnitude, lands on the image of a bisector, and interacts with the order of generation refinement. The archived rerun fixes the numerical vector used downstream, but it does not decide the refine/project ordering or whether the scalar invariant r is retained. Those questions must be settled upstream.

The algebra assembled here specifies how a returned mass and kinetic correction changes the ratio, how the reflection equivalence of the two source images is to be tested, and how a fixed mass shape implies an absolute scale once one measured splitting is supplied, giving $\sum m_\nu \simeq 0.1127$ eV for the projected shape. These are tests of stated constructions. Their success would strengthen the corresponding source chain; their failure would reject that chain under its assumptions without establishing a no-go theorem for the VERSF neutrino sector as a whole.

The final normalisation attack sharpens the remaining boundary. The projected-source chain is independently reproducible, and the source-amplitude scan is approximately proportional to within about 1% over nearly a factor of three in amplitude, with an effective slope of about 0.17. The stated comparator is reached at $|q| \approx 0.1756$, only about 5.6% above the archived amplitude 0.1664. Thus the residual is quantitatively an upstream normalisation question, not an unexplained downstream neutrino effect.

The primitive current programme also derives a positive source current action and an exact current-link relative normalisation. Direct identification of the source-action stiffness with the history/Fisher

stiffness is not admissible because the two are differently typed and no canonical identification has been derived. Its post-freeze value $R_v \approx 0.00991$ is therefore diagnostic only; the typing mismatch is the decisive objection.

The paper consequently stops at a precise frontier. The frozen projected-source construction makes a quantitative prediction now: $R_v = 0.02822$. The source-amplitude response is quantified, and the difference from the stated comparator is traceable to a 5–6% source-normalisation question on this branch. The unresolved normalisation and projection/generation provenance limit the claim that the value is derived from VERSF at full first-principles closure; they do not remove its status as the prediction of the declared construction. Unless an independent VERSF derivation changes the upstream normalisation, $R_v = 0.02822$ remains the number to be tested. Increasingly precise neutrino experiments can therefore confirm or rule out this construction under its stated assumptions.

Appendix A: Sixth-order orientation-selection derivation

Let the retained Fold doublet be $x = R(\cos\varphi e_c + \sin\varphi e_s)$ on the seven-node wheel carrier (hub plus six rim nodes). The density space decomposes under D_6 as $2A_1 \oplus B \oplus E_1 \oplus E_2$, of dimension $1 + 1 + 1 + 2 + 2 = 7$: two A_1 copies (hub and rim mean), the alternating $m = 3$ rim mode, which is the one-dimensional representation odd under the rim rotation by 60° , and the two doublets. Let η denote the five-dimensional complement $2A_1 \oplus B \oplus E_2$ of the retained E_1 doublet. After exact current elimination the reduced density action has the form

$$\Gamma[d] = \frac{1}{2} d^T H d + \sum_i d_i^3 + \frac{1}{4} \sum_i d_i^4, d = x + \eta.$$

Solving $\partial\Gamma/\partial\eta = 0$ order by order gives $\eta = \eta_2 R^2 + \eta_3 R^3 + \eta_4 R^4 + \dots$; the cubic term feeds the E_1 doublet into the alternating mode at $O(R^3)$ with an amplitude proportional to $\cos 3\varphi$, and its elimination contributes $-c R^6 \cos^2 3\varphi$ with $c > 0$. Substitution into Γ generates the first allowed angular invariant at degree six; D_6 invariance permits $\cos 6\varphi$ and the explicit elimination fixes its coefficient and sign. The negative $\cos 6\varphi$ coefficient selects the six symmetry axes rather than the six bisectors.

Source note: the labelling of the one-dimensional alternating representation and the assignment of the two A_1 copies are to be checked against the archived basis and Hessian H in the supplementary material (Section 14.2). The dimension count alone does not re-execute the elimination, and the numerical sixth-order coefficient is treated as independently reproduced only once that basis and Hessian are supplied.

Appendix B: Why CFB-2 gives a scale-free ratio

On CFB-2 the neutrino completion operator is a real positive matrix and $Y_\nu = \mathcal{R}_f^{(v)}$. The sterile block is $M_R = 2 Y_\nu$ and the Dirac block is proportional to the same Y_ν , so all three generation eigenvectors diagonalise the same two-by-two active/sterile block. If y_i are the three eigenvalues of Y_ν , each light state has mass $|\lambda_-| y_i$ and each heavy state has mass $\lambda_+ y_i$, where $\lambda_\pm$ depend only on the common active/sterile normalisation. The light mass ratios and R_ν therefore depend only on $\{y_i\}$ and the common neutral scale cancels exactly. This is why the CFB-2 ratio is a clean capacity diagnostic even though its absolute masses are not physical.

Appendix C: Claim ledger

Claim	Grade	Qualification
Half/quarter generation sequences	Conditional structural derivation	Requires the half-lazy generation operator to act on the neutral Fold coordinate.
Sixth-order $\cos 6\varphi$ anisotropy	Executed source result	Sign robust on two source metrics; relative magnitude metric-dependent; amplitude not fixed by the truncation.
Same-wheel $R_\nu = 0.117564$	Blind diagnostic; falsified	Construction rejected under its assumptions.
NSD principle	Foundational VERSF principle	Identifies unrecorded representatives; relational orientations are unaffected; does not by itself license projection.

Claim	Grade	Qualification
Orbit-invariant $R_\nu \approx 0.0536$	Executed; incompatible with the benchmark at retained order	Uses the two quadratic images of the source-selected minima; source equivariance unresolved.
Hypothesis P (factorisation through the projection, with generation rule)	Modelling assumption, not derived	Halves the source magnitude; image of a bisector; generation rule part of the hypothesis; decisive open object.
Projected $R_\nu = 0.02822$	Conditional numerical candidate	Placed on record as the frozen prediction of the declared projected-source construction; no continuous coefficient fitted; frozen before comparison; approximately two-sigma local residual. Its grade remains conditional on Hypothesis P, the unresolved generation provenance, and the declared source and parent normalisation.
Sterile-radius cancellation	Theorem for the leading Schur complement	Equivalent full parents require transformed kinetics and currents; independence from other parents not established.
$B_L^{irr} = 0$ at retained tree order	Declared-order source result	Does not prove an all-orders zero.
Radiative effects	Illustrative sensitivity estimate	Potentially relevant in size; sign and magnitude of the ratio shift require the correction matrix.
$\sum m_\nu \simeq 0.1127$ eV (projected shape)	Post-freeze consequence of the shape	Above the Λ CDM benchmark, below the w_0 w_a CDM benchmark; model dependent.
Direct zero responses	Established on returned configurations only	Extension to the projected source requires a factorisation theorem or direct evaluation.
Absolute masses, PMNS	Not returned from source	Requires typed same-branch pairing, residues, current map.

Appendix D: Arithmetic identities used in the text

Ratio from mass ratios. With $m_1 = 1$, $R_\nu = (m_2^2 - 1)/(m_3^2 - 1)$. Projected spectrum: $(1.04682488^2 - 1)/(2.09668389^2 - 1) = 0.095842330/3.396083333 = 0.02822143$.

Single orientation: $(1.09041985^2 - 1)/(2.12785637^2 - 1) = 0.1890155/3.5277726 = 0.0535793$. Same-wheel: $(1.62735454^2 - 1)/(3.87560756^2 - 1) = 1.6482829/14.0203340 = 0.1175637$. CFB-2: $(0.1528470774/0.01681873236)^2 = 82.590$ and $(1/0.01681873236)^2 = 3535.19$, giving $81.590/3534.19 = 0.0230859$.

Comparator uncertainty. See Appendix H: $\sigma_R^2 \simeq \sigma_S^2/D^2 + S^2\sigma_D^2/D^4$ with $S = 7.49 \times 10^{-5}$, $D = 2.513 \times 10^{-3}$, $\sigma_S = 0.19 \times 10^{-5}$, $\sigma_D = 0.020 \times 10^{-3}$ gives $\sigma_R = 7.92 \times 10^{-4}$.

Sensitivities. $c_2 = 2m_2^2/A = 2(1.095842330)/0.095842330 = 22.86761$; $c_1 = -2/A + 2/B = -20.86760 + 0.58891 = -20.27869$; $c_3 = -2m_3^2/B = -2(4.396083333)/3.396083333 = -2.58891$; $c_1 + c_2 + c_3 = 0$.

Half-lazy eigenmode. For $L(a, b)$ on three nodes and the centred profile $(5, -1, -4)/12$, $L v \propto (6a, -6a + 3b, -3b)$; proportionality to $(5, -1, -4)$ requires $\lambda = 6a/5$ from the first row and $3b = 24a/5$ from the third, so $b/a = 8/5$, and the second row then gives $-6a/5 = -\lambda$, consistent.

Absolute masses. $s^2 = \Delta m_{31}^2/(r_3^2 - r_1^2)$ with $\Delta m_{31}^2 = 2.513 \times 10^{-3} \text{ eV}^2$: projected shape $r = (1, 1.04682488, 2.09668389)$, $r_3^2 - r_1^2 = 3.396083333$, $s = 0.027202389 \text{ eV}$, sum 0.112713345 eV ; single orientation $s = 0.026690 \text{ eV}$; same-wheel $s = 0.013388 \text{ eV}$; CFB-2 shape $r = (0.01681873236, 0.1528470774, 1)$, $s = 0.050137 \text{ eV}$, $m_1 = 0.000843 \text{ eV}$. No further normalisation is used.

Appendix E: Projection, reflection invariance and the generation average

Let $q_{\pm} = (-1/2, \pm\sqrt{3}/2)$ and $Pq_{\pm} = (-1/2, 0)$. For an analytic scalar response in a neighbourhood of the source, reflection invariance $F(s, t) = F(s, -t)$ allows exactly the form

$$F(s, t) = \sum_{k \geq 0} f_{2k}(s) t^{2k}.$$

Projection removes every term with $k \geq 1$. On an open neighbourhood, factorisation of F through P requires all coefficient functions f_{2k} , $k \geq 1$, to vanish; accidental cancellation of their sum at a single orbit is not a factorisation theorem. The example $F = s^2 + t^2$ gives $F(q_{\pm}) = 1$ and $F(Pq_{\pm}) = 1/4$.

Generation average. Each selected unit-amplitude representative has $u^2 = 1/4$, $v^2 = 3/4$, and the four values of uv sum to zero. With $\lambda_g = 2^{-(g-1)}$,

$$\frac{1}{4} \sum_{\varphi \in \mathcal{A}_F} Q(T_g x(\varphi)) = \left(\frac{1}{4} - \frac{3}{4} \lambda_g^2, 0\right),$$

giving $(-1/2, 0)$, $(1/16, 0)$ and $(13/64, 0)$ for $g = 1, 2, 3$. Writing $r = u^2 + v^2$, $s = u^2 - v^2$, $t = 2uv$,

$$Q(T_g x) = \left(\frac{1 - \lambda_g^2}{2} r + \frac{1 + \lambda_g^2}{2} s, \lambda_g t\right),$$

so the refined first component depends on the scalar invariant r as well as on the E_2 source. The unprojected orbit has $r = 1$, whereas the preimage $x_* = (0, 1/\sqrt{2})$ of Pq has $r = 1/2$, and refining that preimage gives $Q(T_g x_*) = (-\lambda_g^2/2, 0)$, a different sequence again. An explicit source prescription must therefore state whether the original scalar invariant is retained when the E_2 component is projected; the archived rerun recovers the numerical source vector used in Section 11.3 but does not determine whether the original scalar invariant is retained or which refine/project ordering generated that vector. That upstream provenance therefore remains open.

Appendix F: Non-extension of finite-set zeros to the projected source

Expand a scalar component of a local response near the reference background as

$$C(x) = C_0 + Lx + \frac{1}{2} x^T A x + O(\|x\|^3), A = A^T.$$

If $Lx = 0$ on all four selected representatives then $L = 0$, because the representatives span the two-dimensional Fold plane. For the quadratic term, on the unit-amplitude selected configurations,

$$x^T A x = \frac{1}{4} A_{uu} + \frac{3}{4} A_{vv} \pm \frac{\sqrt{3}}{2} A_{uv}.$$

A zero on all four requires $A_{uv} = 0$ and $A_{uu} + 3A_{vv} = 0$, but not $A = 0$. For example $C_2(u, v) = 3u^2 - v^2$ vanishes on every selected representative while $C_2(0, 1/\sqrt{2}) = -1/2$. The example is not a proposed VERSF correction; it proves that the finite-set zero does not license the quadratic extension. For matrix-valued responses the argument applies component by component.

Sufficient condition. If the retained response is known independently to have the form $C_2(x) = \mathcal{L}(Q(x))$ with $\mathcal{L}$ linear on E_2 , and $\mathcal{L}(q_+) = \mathcal{L}(q_-) = 0$, then $\mathcal{L} = 0$ because q_+ and q_- span E_2 , and $\mathcal{L}(Pq) = 0$ follows. This requires the pure- E_2 factorisation; a scalar channel or a nonlinear dependence on q must first be excluded by the actual source tensor.

Appendix G: Schur cancellation versus active-sterile mixing

For canonical blocks M_D, M_R write a light Takagi column as $(\ell, s)^T$. To leading order in the light-to-sterile mass ratio the lower-block equation gives $s = -M_R^{-1} M_D^T \ell + \dots$. With $T = M_R^{-1} M_D^T$ and v the normalised leading active component, $w_{act} \simeq 1/(1 + \|Tv\|^2)$: the leading active-weight deficit is quadratic in a dimensionless mixing amplitude, and a change of source can change T linearly. When the unwhitened factorisation of Section 11.1 is already canonical, $T = (a/b) H_N^{-1} P_N^{-1} X^T$; the combination a^2/b sets the light scale while a/b , the matrices and their kinetic normalisation determine mixing.

The one-generation real parent

$$M(h) = \begin{pmatrix} 0 & ah \\ ah & bh^2 \end{pmatrix}, a, b, h > 0,$$

has Schur mass a^2/b independent of h , while its exact light mass and active weight are

$$m_\ell(h) = \frac{\sqrt{b^2 h^4 + 4a^2 h^2} - bh^2}{2}, w_\ell(h) = \frac{1}{2} \left(1 + \frac{bh^2}{\sqrt{b^2 h^4 + 4a^2 h^2}} \right).$$

For $a = 0.05, b = 1$ the Schur mass is 0.0025 at both $h = 1$ and $h = 0.1$, whereas the exact light masses are 0.00249378 and 0.00207107 and the active weights 0.997519 and 0.853553. These dimensionless parents illustrate the limitation of the cancellation theorem; they are not replacements for the VERSF parent.

Appendix H: First-order variation of the ratio under a mass and kinetic correction

Comparator variance. For $R = S/D$, first-order propagation gives

$$\sigma_R^2 \simeq \frac{\sigma_S^2}{D^2} + \frac{S^2 \sigma_D^2}{D^4} - \frac{2S \text{Cov}(S, D)}{D^3}.$$

A two-observable comparison that profiles the common scale predicts $(S, D) = (Rd, d)$ for a fixed ratio R and, for a joint Gaussian likelihood with constant covariance and an interior profiled scale,

$$\chi_{min}^2(R) = \frac{(S_{obs} - R D_{obs})^2}{\sigma_S^2 + R^2 \sigma_D^2 - 2R \text{Cov}(S, D)},$$

the variance of the single linear residual $S - RD$. The projected ratio gives a signed residual of about -2.008 standard deviations (theory minus observation) in the uncorrelated symmetric approximation. Far from the measured region the actual likelihood must be used rather than this approximation.

Matrix correction. In the canonical basis of the unperturbed parent let $M = M^T$, let $\delta B = \delta B^T$ be the returned mass correction and $\delta K = \delta K^\dagger$ the kinetic correction. To first order $(I + \delta K)^{-1/2} = I - \frac{1}{2} \delta K + O(\delta K^2)$, so the corrected canonical mass matrix is $M + \delta M$ with

$$\delta M = \delta B - \frac{1}{2} \delta K^T M - \frac{1}{2} M \delta K.$$

For nondegenerate positive Takagi values and columns u_i of U ,

$$\delta m_i = \text{Re}(u_i^T \delta M u_i).$$

Writing the infinitesimal unitary change as $U \rightarrow U(I + S)$ with $S^\dagger = -S$, the diagonal first-order equation is $\delta D = U^T \delta M U + S^T D + DS$, whose last two diagonal terms are purely imaginary; taking real parts yields the formula. At a degeneracy the perturbation must be resolved within the degenerate subspace, or an exact updated Takagi factorisation used. Substituting the δm_i into the ratio derivative of Section 12.4 tests any returned loop matrix without fitting its direction to the residual. If $\delta M = \varepsilon M$ for real ε then every $\varepsilon_i = \varepsilon$ and the ratio shift vanishes to first order; the same holds exactly for a finite common positive rescaling. The required calculation is to return δB and δK at a specified matching scale, whiten them consistently, diagonalise the corrected parent, identify the same light states through their currents, and compare the resulting ratio.

Appendix I: Consistent-basis test of representative equivalence

Choose the convention $\psi_+ = G\psi_-$ with G invertible. Equivalent data satisfy

$$B_- = G^T B_+ G, K_- = G^\dagger K_+ G, C_- = C_+ G,$$

where C maps the neutral fields to the physical current. With $Z_\pm = K_\pm^{-1/2}$ and $W = K_+^{1/2} G K_-^{-1/2}$ one has $W^\dagger W = I$, $G Z_- = Z_+ W$, and

$$\mathcal{M}_- = W^T \mathcal{M}_+ W,$$

so the Takagi values coincide. For nondegenerate values a corresponding choice is $U_- = W^\dagger U_+$ up to ordering and permitted phases; at a degeneracy the subspaces are compared. The canonical current amplitudes also coincide: $C_- Z_- U_- = C_+ Z_+ U_+$.

This is the explicit test of the claimed equivalence of q_+ and q_- in Section 11.2. The matrix G must be supplied by the actual reflection action on the typed carrier, not chosen by fitting one spectrum to another. If no such transformation exists, the claimed physical equivalence is unproved or a relational frame has to be included. The mass congruence is checked with the dimensionless residual

$$r_B = \frac{\|B_- - G^T B_+ G\|_F}{\|B_-\|_F + \|G^T B_+ G\|_F},$$

with corresponding residuals for K and C and a declared convention for an identically zero denominator; tolerances follow the solver precision and conditioning.

References

- [1] K. Taylor, Physical Bath Selection and Neutrino Mixing in VERSF, September 2026.
- [2] K. Taylor, The VERSF Master Substrate Action and Constitutive-Descent Theorem, VERSF corpus, 2026 edition.
- [3] K. Taylor, The VERSF History-Measure, Generation-Geometry and Branch-Closure Theorem (HMGBC-1), VERSF corpus, 2026 edition.
- [4] K. Taylor, Primitive-Matter Selection in VERSF: Cumulative Executed Theorems on Fact, Commitment, Record Retention and Renewal, cumulative 2026 edition. The primitive nonlinear third-jet result $g_{\chi} = 1/2$ is given by the Primitive Nonlinear Axle-Source Return Theorem (Theorem 32.62); PMS-EX97 gives the canonical master-Hessian E2 stiffness $\kappa_{E2} = 3$; PMS-EX98–EX99 give the clock-invariant susceptibility $g_{\chi}/\kappa_{E2} = 1/6$; PMS-EX106 gives the type-correct same-history E2 stiffness $\kappa_{E2} = 2.924147572697$ and susceptibility 0.170990002238 .
- [5] K. Taylor, Uniform Binary Admissibility and the Half-Lazy Closure Operator: A Structural Route to the χ -Halving Law in VERSF.
- [6] K. Taylor, SMD-G0A: Gauge-Generation Intertwiner Census.
- [7] K. Taylor, Generation History to Neutrino Response, execution note, 22 September 2026.
- [8] K. Taylor, Native History Sources and the Full Neutral Response in VERSF, execution note, 21 September 2026.
- [9] K. Taylor, Neutral Bilinear and Bath Response, execution note, 22 September 2026.
- [10] K. Taylor, Evaluation of the Bath/Completion Source, Kinetic Residues and Neutral Majorana Functionals, execution note, 22 September 2026.
- [11] K. Taylor, RCCSJ-1: Representation-Resolved Current-Cycle Source Jet and Positive-History Realisation in VERSF, 3 September 2026.
- [12] K. Taylor, CFB-2 archived finite-regulator candidate result, branch VERSF-CFB2-K7xT4-2026-07.
- [13] K. Taylor, CFB-3 archived finite-regulator candidate result, branch VERSF-CFB3-K7xT4-MagneticGeneration-2026-07.
- [14] Particle Data Group, Review of Particle Physics, neutrino masses, mixing and oscillations; used only for post-freeze comparison (see also [22]).
- [15] Supporting calculation notes for SFNMR-1 (September 2026): Sixth-Order Fold Anisotropy; Mixed Fourth-Order Fold-Fermion Attack; MA9-prime and Source-Linked Neutrino Test; Universal Fold-History Completion Test. To be supplied with the reproducibility package of Section 14.2.
- [16] K. Taylor, Gravity from Fold Density Gradients: A Unified VERSF Derivation, VERSF corpus, 2026 edition, section defining the Fold as the minimal unit of distinguishability carrying one irreducible committed binary outcome. Specifically Section 1, pp. 3 to 4, and Section 3, Step 1, p. 5. These passages support the foundational premise of Principle NSD (one irreducible committed distinction per Fold); they do not address, and are not cited in support of, the projection prescription of Hypothesis P.
- [17] K. Taylor, VERSF D-Free Source-Linked Neutral Compression, execution note, 25 September 2026.
- [18] K. Taylor, VERSF Active Majorana Retained-Order Audit, execution note, 25 September 2026.
- [19] K. Taylor, VERSF LLHH One-Loop Size Diagnostic, execution note, 25 September 2026.

- [20] I. Esteban, M. C. Gonzalez-Garcia, M. Maltoni, I. Martinez-Soler, J. P. Pinheiro and T. Schwetz, NuFit-6.0: Updated global analysis of three-flavor neutrino oscillations, JHEP 12 (2024) 216, arXiv:2410.05380v2, Table 1 (normal ordering, IceCube 2024 with Super-Kamiokande atmospheric data). The comparison retains this specified benchmark and does not claim it is the latest available global fit.
- [21] W. Elbers et al. (DESI Collaboration), Constraints on Neutrino Physics from DESI DR2 BAO and DR1 Full Shape, arXiv:2503.14744v3 (7 October 2025), abstract and Sections IV.1 and IV.3. The quoted bounds refer to the stated data combinations, cosmological models and priors, not to a model-independent mass limit.
- [22] Particle Data Group, Neutrino Masses, Mixing and Oscillations, Review of Particle Physics, 2025 update, Sections 14.2, 14.3 and Table 14.7, <https://pdg.lbl.gov/2025/reviews/rpp2025-rev-neutrino-mixing.pdf>. Where the atmospheric splitting is tabulated as Δm_{32}^2 it is converted with $\Delta m_{31}^2 = \Delta m_{32}^2 + \Delta m_{21}^2$.
- [23] K. Taylor, The Twelve-Sector Current-Rigidity, Source-Generated Rank-Six Current-Stiffness and Current-Link Schur Reciprocity Theorem in VERSF, September 2026.
- [24] K. Taylor, PCRS-1: The Primitive Current-Record Scale and Absolute Gauge-Normalisation Boundary Theorem in VERSF, September 2026.
- [25] UFG Blind Stage-A Source-Only Execution and Generation History Neutrino Reproduction packages, September 2026; independently rerun for the SFNMR-1 final normalisation audit.