The Neutrino Mass and Charge-Conjugation Completion Theorem in VERSF

▲ Programme Milestone — Neutral-Fermion and Lepton-Completion Series Gate NMCC-1 / Neutral Weyl Carrier, Typed Charge-Conjugation Completion, Gauge-Admissible Dirac–Majorana Classification, Canonical Nambu Normalisation, Generalised Takagi Mass Spectrum, Exact Schur-Complement Seesaw, Lepton-Number and CP Firewalls, Planck-to-Neutrino Descent, Mixing Non-Insertion Audit, and Numerical Non-Insertion Closure

Keith Taylor VERSF Theoretical Physics Programme Neutral-Fermion and Lepton-Completion Series Tier A conditional gate-theorem paper, NMCC-1 July 2026

General Reader Summary

Every particle of matter we know gets its mass in a particular way, and for almost all of them the rulebook is strict. An electron carries electric charge, and that charge acts like a chaperone: as long as electromagnetism holds, the electron is forbidden from taking the exotic mass route described below, so its mass must be of the ordinary kind — whatever deeper machinery ultimately generates it.

The neutrino is the exception, because the neutrino carries no electric charge. Remove the chaperone and a second possibility quietly appears. A neutrino could gain mass the ordinary way, by pairing up with a partner particle — this is called a *Dirac* mass, and it is how the electron does it. But a neutral particle could also, in a precise mathematical sense, pair up with its own mirror image — its antiparticle twin. That second option is called a *Majorana* mass, and it is only available to particles with no charge. Whether the neutrino uses the first route, the second, or a blend of both is one of the great unanswered questions in physics.

This paper does something deliberately unglamorous and, we would argue, necessary: it writes the complete rulebook for that question before anyone in the VERSF programme is allowed to claim an answer.

The first rule is a warning against a tempting shortcut. It is easy to slide from "the neutrino is neutral" to "so it is probably a Majorana particle," or worse, to "so the mirror-pairing explains its mass." The paper proves this reasoning is empty. Neutrality only unlocks a door; it does not push anything through it, and it certainly does not generate a mass. A mass always needs three separately earned ingredients: the right cast of particles must actually exist, the pairing must be allowed by every conservation law in force, and something must set the size. Each ingredient has to be derived; none comes free with neutrality.

The second contribution is the correct machinery. Because of the mirror-pairing option, the neutrino mass problem is mathematically a different species of problem from the electron's — reusing the electron's toolkit with the labels swapped gives wrong answers. The paper builds the right toolkit: how to organise every allowed mass term into a single tidy object, how to read the true physical masses out of that object, and how its internal pattern reveals — unambiguously — whether the outcome is Dirac, Majorana, a near-miss hybrid called *pseudo-Dirac*, or the celebrated *seesaw*. The seesaw is the elegant idea that if the neutrino has an extremely heavy hidden partner, the observed neutrino masses are pushed down as the partner's mass goes up — one end of a seesaw drops because the other end rises. That would explain, in one stroke, why neutrino masses are millions of times smaller than the electron's.

The third contribution is a pair of proofs about what standard oscillation measurements do not determine. Neutrino oscillation experiments — the ones that watch neutrinos change identity mid-flight — are magnificent, but the paper proves they are blind in two specific ways. They measure only *differences*: technically, differences between the squares of the masses. Adjust all the masses so that every squared mass grows by the same amount, and every standard oscillation experiment gives identical readings. And ordinary flavour-change measurements are insensitive to the phase fingerprints that would reveal Majorana character — those could only surface in a different, lepton-number-violating kind of process. So the two headline questions — how heavy is the lightest neutrino, and is it its own mirror twin — cannot be settled by standard oscillations. They need different evidence or a genuine theoretical derivation.

There is also a proof of a subtler frustration: even a perfect measurement of the light neutrino masses could never be traced back to a unique deep-level origin. The paper exhibits a whole continuous family of different hidden-sector stories — eighteen bookkeeping directions wide before redundant relabellings are even removed — that all produce identical observable outcomes, and a genuinely continuous family survives after the bookkeeping is stripped away. Anyone who fits the data and announces "therefore the hidden sector looks like this" has picked one story from an infinite shelf.

Two possible answers to "what sets the size?" are then put on the table, carefully labelled as candidates. The first is the seesaw route, anchored at the enormous Planck scale, where the tininess of neutrino masses reflects the vastness of the hidden partner's mass. The second is new to this gate and comes from elsewhere in the VERSF programme: an energy scale associated with the framework's commitment field works out to about three thousandths of an electronvolt — a plausible size for the *lightest* neutrino, or for a smaller secondary contribution, though the observed oscillation splittings demand that at least one neutrino be roughly eighteen times heavier. That partial coincidence is striking, and the paper treats it with matching suspicion. The commitment field is a different kind of field from a neutrino, and a number being the right size is not the same as a mechanism being derived. So the coincidence is installed as a named, testable candidate — with the missing link explicitly listed as a debt — sitting alongside the seesaw route, not replacing it. Which route dominates becomes a question the programme can eventually answer and be judged on.

Finally, the honesty ledger. The paper closes the mathematics — the rulebook, the classification, the machinery, the impossibility proofs — and then lists, by name, twenty-six debts that must be paid before any actual neutrino mass can be claimed: does the hidden partner exist, what sets each ingredient's size, how does the commitment scale couple to neutrinos, and so on. Nothing on that list may be quietly filled in from the measured data and then presented as a prediction.

In one sentence: this paper builds the complete, rigorous rulebook for the neutrino mass problem — what the structure must be, what experiments can and cannot decide, which two scale anchors — feeding three distinct contributions — are on the table, and exactly what VERSF still owes — so that when a number finally emerges, it is a genuine prediction and not a disguised fit.

Scope and Closure Box

Claim	NMCC-1 status
Active neutrino carrier must be distinguished from any sterile carrier	Closed as requirement
Charge-conjugate carrier is a conjugate representation, not a second observed particle	Closed
A neutral quadratic mass bilinear is complex symmetric	Closed
General Dirac expansion requires independent particle and antiparticle CAR families	Closed
The mode-spinor relation $v = C\,\bar{u}^T$ or the Dirac equation forces identification of particle and antiparticle sectors	Disproved
One-particle conjugation, Fock-space charge conjugation and Majorana self-conjugacy are distinct	Closed
Nambu left-handed notation imposes a Majorana identification	Disproved
Positive neutral kinetic metrics admit canonical normalisation	Closed
The canonically normalised neutral mass operator remains symmetric	Closed
Physical tree-level neutral masses are Takagi singular values	Closed
The unordered Takagi mass multiset is unique	Closed
Generalised Takagi singular-value equation with noncanonical kinetic metric	Closed
Neutrality alone implies a Majorana mass	Disproved
Charge conjugation alone generates a mass	Disproved
A Majorana field condition implies charge-conjugation invariance of all interactions	Disproved
Pure Dirac completion produces pairwise-degenerate Takagi singular values	Closed
Exact continuous lepton number exists iff a nontrivial generator satisfies both the kinetic and mass-invariance equations	Closed

Claim	NMCC-1 status
Standard lepton number is preserved iff M_L = M_R = 0 and [Q_L, K_ν] = 0	Closed
Nonzero Majorana blocks necessarily violate the standard lepton-number assignment	Closed
Neutral completions are taxonomised — nested, per invariant subspace	Closed
A bare active-neutrino Majorana block is electroweak gauge invariant	Disproved
A Dirac mass requires an independently admitted right-handed neutral carrier	Closed as requirement
A sterile Majorana block is gauge admissible merely because its carrier is named	Disproved
Sterile right-handed carrier is, within the minimal representation ansatz, the minimal anomaly-neutral electroweak-singlet chiral carrier	Inherited conditionally
A bare sterile Majorana block is admissible under an exact unbroken gauged B−L	Disproved
Exact unbroken gauged B−L forces M_L = M_R = 0 (Dirac-class or massless)	Closed conditionally
Gauge-admissible pre-breaking completions are classified	Closed structurally
Exact algebraic Schur-complement identity	Closed
Schur split is a congruence, not a unitary congruence	Closed
Schur complement alone equals the exact physical light spectrum	Disproved
Hierarchical seesaw interpretation under controlled light–heavy mixing	Closed conditionally
Type-I rank bound rank(M_ν,eff) ≤ rank(M_D)	Closed
Fewer sterile than active carriers force tree-level massless light modes in type I	Closed
Low-energy neutrino data uniquely determine M_D and M_R	Disproved
Oscillations determine the absolute neutrino mass scale	Disproved
Standard lepton-number-conserving oscillations determine Majorana phases	Disproved
Oscillations distinguish Dirac from Majorana neutrinos by themselves	Disproved
Mixing matrix is defined without reference to the charged-lepton frame	Disproved
PMNS-type mixing arises from the relative charged-lepton and neutrino frames	Closed

Claim	NMCC-1 status
A Hermitian frame kernel determines masses, absolute scale, ordering or Majorana phases	Disproved
Frame–Takagi compatibility criterion for inheriting mixing angles from a Hermitian kernel	Closed conditionally
Weak-commitment kernel inherited as candidate active mixing frame	Inherited conditionally — mixing-only
Charged-lepton localisation exponent $\kappa = 8/3$ may be exported to neutrinos	Disproved absent a separate neutral-sector derivation
Planck-seesaw normal form	Closed conditionally on a Planck-related heavy completion
Planck mass alone fixes the neutrino scale	Disproved
κ-closure scale $m_\kappa c^2 = \sqrt{(4/3)}\, \hbar c/\xi \simeq 2.8$ meV as candidate direct-completion anchor	Inherited conditionally
The κ scalar mass may be relabelled as a neutrino mass	Disproved
Hybrid κ–Planck two-anchor, three-contribution completion normal form	Closed conditionally
κ-field Lagrangian-uniqueness and nonlinear-exclusion claims	Not inherited — open
Existence and number of right-handed neutrino carriers	Open
Whether the physical completion is Dirac, Majorana, pseudo-Dirac or mixed	Open
Derivation of M_L, M_D, M_R from VERSF substrate dynamics	Open
Derivation of Φ_L, Φ_R, Y_ν, $\hat{L}_\nu$, $\hat{R}_\nu$	Open
Light-neutrino mass ordering	Open
Absolute lightest neutrino mass	Open
Active–sterile mixing and heavy spectrum	Open
Dirac and Majorana CP phases	Open
Radiative and on-shell matching	Open
Empirical confirmation	Not established

Notation

Equations are set in Unicode. Superscripts: ^{T} (transpose), † (Hermitian adjoint), $^{-1}$ (inverse), * (complex conjugate), ^{c} (charge conjugate or canonical, fixed by context). Composite powers are written with a caret, as in K_ν^(−T/2), with the definition K_ν^(−T/2) ≔ (K_ν^(−1/2))^T. The Planck symbol m_P denotes the unreduced Planck mass (ℏc/G)^(1/2) throughout; the reduced mass m̄_P = m_P/√(8π) is never used without its bar. Carriers are script objects: $\mathscr{H}$ (carrier space), $\mathcal{N}$_L (neutral Nambu carrier), $\mathscr{B}$_ν (completion bilinear). Three distinct operator types recur and are deliberately given distinct symbols: K_ν is the positive Hermitian *kinetic metric*; $\mathcal{K}_\nu$ is the Hermitian *weak-commitment frame kernel* inherited from the flavour-frame programme; $M_{\nu,\mathrm{eff}} = M_{\nu,\mathrm{eff}}^T$ is the complex-symmetric *light mass operator*, with Hermitian magnitude square $H_{\nu,m} \coloneqq M_{\nu,\mathrm{eff}}^\dagger M_{\nu,\mathrm{eff}}$. No two of these may be conflated (Section 21). $\oplus$ is the direct sum, $\otimes$ the tensor product, $\|\cdot\|$ an operator norm declared where used, $\sigma(\cdot)$ the spectrum, $\mathrm{diag}(\cdot)$ a diagonal matrix, and $\mathrm{rank}(\cdot)$, $\det(\cdot)$ the usual matrix invariants.

Abstract

A neutral fermion admits mass completions that have no analogue in a charged sector, because a neutral Weyl carrier may couple either to an independently existing opposite-chirality carrier or to a charge-conjugate carrier. Neutrality, however, does not decide which completion is realised and does not itself generate a mass. This paper establishes the Neutrino Mass and Charge-Conjugation Completion Theorem in VERSF by separating carrier existence, charge-conjugation typing, gauge admissibility, canonical normalisation, physical mass extraction, lepton-number classification and absolute-scale anchoring — and by installing an explicit firewall at every joint where these are conflated in practice.

Let $\mathscr{H}_{\nu L}$ be the active left-handed neutrino carrier and, if independently admitted, let $\mathscr{H}_{NR^c}$ be the left-handed charge conjugate of a sterile right-handed carrier. On the neutral left-handed Nambu space $\mathcal{N}_L = \mathscr{H}_{\nu L} \oplus \mathscr{H}_{NR^c}$, let $K_\nu > 0$ be the Hermitian kinetic metric and let the dimension-one completion bilinear be

$$\mathcal{B}_\nu = \begin{pmatrix} M_L & M_D \\ M_D^T & M_R \end{pmatrix}, \qquad \mathcal{B}_\nu^T = \mathcal{B}_\nu .$$

Canonical normalisation gives the symmetric operator $M_\nu^c = K_\nu^{-T/2}\, \mathscr{B}_\nu\, K_\nu^{-1/2}$. Its physical tree-level masses are the Takagi singular values,

$$U_\nu^T M_\nu^c U_\nu = \mathrm{diag}(m_i), \quad m_i \geq 0,$$

or equivalently the positive square roots of the generalised eigenvalues of the pair $(\mathscr{B}_\nu^\dagger K_\nu^{-T} \mathscr{B}_\nu, K_\nu)$.

The paper proves that a pure off-diagonal completion $M_L = M_R = 0$ yields pairwise-degenerate Takagi masses and an exact Dirac pairing; that continuous lepton number exists exactly when a nontrivial Hermitian generator Q satisfies $[Q, K_\nu] = 0$ and $Q^T \mathscr{B}_\nu + \mathscr{B}_\nu Q = 0$; that nonzero Majorana blocks violate the standard lepton-number assignment; and that the Dirac, Majorana, pseudo-Dirac and mixed regimes are distinguished by block structure and symmetry, never by neutrality alone.

Gauge admissibility is treated before electroweak breaking. A Dirac block requires an independently justified gauge-singlet right-handed carrier and a Higgs-type bridge; a bare active Majorana block is forbidden, while an effective left-handed completion requires a gauge-

admissible operator such as the dimension-five neutral completion or an equivalent VERSF closure interface; a sterile Majorana block is admissible only after the sterile carrier and all of its charges — Standard Model and VERSF alike — are independently derived.

For invertible M_R, the quadratic form admits the exact Schur complement

$$M_{\nu,\mathrm{eff}} = M_L - M_D M_R^{-1} M_D^T.$$

The paper proves the associated rank, nullity and determinant constraints, exhibits the Schur split as a non-unitary congruence, and installs a firewall between the exact algebraic identity and the approximate identification of its Takagi spectrum with the physical light spectrum, which additionally requires a controlled light–heavy hierarchy and momentum-dependent matching.

The paper also proves two independent scale obstructions. First, the rescaling $\mathcal{B}_\nu \mapsto a\,\mathcal{B}_\nu$ leaves all dimensionless mass ratios and Takagi projectors unchanged. Second, neutrino oscillations are invariant under $m_i^2 \mapsto m_i^2 + \mu_0^2$, so they cannot determine the lightest mass or the common absolute scale; Majorana phases cancel from all standard lepton-number-conserving flavour-oscillation probabilities.

When a Planck-related heavy completion is admitted, the neutral blocks have the conditional Planck normal form

$$\mathcal{B}_\nu^{(P)} = \begin{pmatrix} (v_{cl}^2/m_P)\ \Phi_L \hat{L}_\nu & (v_{cl}/\sqrt{2})\ Y_\nu \\ (v_{cl}/\sqrt{2})\ Y_\nu^T & m_P\ \Phi_R \hat{R}_\nu \end{pmatrix},$$

which gives

$$M_{\nu,\mathrm{eff}}^{(P)} = (v_{cl}^2/m_P) \cdot [\ \Phi_L \hat{L}_\nu - (1/2\Phi_R)\, Y_\nu \hat{R}_\nu^{-1} Y_\nu^T\ ].$$

When the commitment-field scalar scale $m_\kappa = \sqrt{(4/3)}\ \hbar/(c\xi)$ is additionally admitted, the paper installs the hybrid two-anchor, three-contribution normal form

$$M_{\nu,\mathrm{eff}} = m_\kappa \Phi_{\kappa\nu} \hat{L}_\nu^{(k)} + (v_{cl}^2/m_P)\, \Phi_L \hat{L}_\nu - (v_{cl}^2/2m_P \Phi_R)\, Y_\nu \hat{R}_\nu^{-1} Y_\nu^T,$$

in which closure-scale direct, Planck-suppressed direct and Planck-seesaw completions coexist as physically distinct, separately falsifiable contributions, and any reduction to fewer terms is a named premise rather than a default. The κ scale is inherited strictly as a candidate anchor with an owed coupling bridge; a scalar mass parameter is not a fermion mass.

This is an explicit scale-typed candidate architecture, not a numerical derivation. NMCC-1 also supplies the field-level layer beneath the carrier typing: the corrected general Dirac expansion with independent particle and antiparticle CAR families, the trichotomy separating one-particle antilinear conjugation, Fock-space unitary charge conjugation and Majorana self-conjugacy, and the proof that identification of the particle and antiparticle sectors is a completion condition implied neither by the field equation nor by the mode-spinor relation $v = C\,\bar{u}^T$. NMCC-1 closes the neutral-fermion carrier, charge-conjugation, Takagi, symmetry-classification and seesaw logic, while leaving the sterile census, completion operators, dimensionless descent, mixing texture, absolute spectrum and observable matching as named bridge debts.

Contents

Part I — Neutral Carrier, Charge Conjugation and Canonical Spectrum

1. Aim, Inheritance and Novelty

1.1 Inherited results

NMCC-1 inherits several upstream results but must not over-credit them.

From the Standard Model census and occupancy programme (SC-1, SC-2) it inherits the existence of three active left-handed neutrino addresses inside the lepton-doublet carrier. It does **not** inherit an independently existing right-handed neutrino carrier unless such a carrier is separately derived by the world-to-framework occupancy map or a subsequent neutral-sector census.

From the spectral-frame and charged-lepton absolute-scale programmes (SF-1, CLAS-1) it inherits:

- positive kinetic-metric discipline;
- canonical normalisation before mass extraction;
- singular values, not raw matrix entries, as the only frame-invariant physical magnitudes;
- the common-scale no-go for dimensionless flavour data;
- the distinction between a dimensionless magnitude pattern and an independent dimensionful anchor;
- the requirement that flavour addresses precede the observed mass ordering;
- the Planck-basis and closure-basis normal forms for dimensionful scales.

From the closure-interface and Higgs-radial programme (HR-1) it inherits the structural necessity of a gauge-admissible electroweak completion interface. It does not inherit the numerical vacuum scale as a first-principles VERSF result.

From the electroweak representation and connection-closure gate it inherits the doublet typing $L = (\nu_L, e_L) \sim (\mathbf{1}, \mathbf{2}, -½)$, the intrinsic left-handedness and weak-doublet address of the active neutrino, the statement that, within the inherited minimal representation ansatz, $\nu_R \sim (\mathbf{1}, \mathbf{1}, 0)$ is the *minimal anomaly-neutral electroweak-singlet chiral carrier* permitting the conventional Dirac bridge (Section 2.2), and the anomaly-admissibility of a gauged B−L grade once ν_R is included (Section 12.4).

From the substrate-CAR and relativistic-field gate it inherits the field-theoretic layer beneath this paper: substrate CAR algebra → relativistic fermion field → positive- and negative-frequency mode sectors → the distinction between a carrier and its charge-conjugate representation. That gate explicitly reserves antiparticle structure and charge conjugation for a later completion and marks its V_α modes as placeholders pending the relation to charge-conjugate modes. NMCC-1 is that reserved completion, executed for the neutral sector (Sections 3.5–3.8) — including a correction to the general Dirac expansion, which requires independent particle and antiparticle CAR families.

From the Schrödinger-to-Dirac gate it inherits the substrate origin of the relativistic spinor carrier: fermionic mass appears in a first-order relativistic operator, requires Clifford generators, takes the orientation-coupling form $m c^2 \beta$, and joins opposite chiralities — so that a massless fermion retains decoupled chiral propagation. This is inherited as a relativistic carrier and bilinear-form theorem, giving M_D substrate provenance (Section 12.2); it is not inherited as a neutrino mass derivation.

From the κ-field/commitment-field equivalence gate it inherits, conditionally, the scalar commitment mode and its mass scale $m_\kappa = \sqrt{(4/3)}\, \hbar/(c\xi)$ (Section 17.4). It does **not** inherit that gate's Lagrangian-uniqueness or nonlinear-exclusion claims, which remain open (Section 17.4).

From the weak-commitment and flavour-frame gates it inherits the Hermitian frame operator $T_\nu = D_0 I + \varepsilon \mathcal{K}_\nu$, the scale–shape separation it encodes, the relative-frame definition of lepton mixing, and a conditional candidate kernel for the active mixing frame (Section 21). It does **not** inherit a mass operator, a mass spectrum or an absolute scale from any of them: a Hermitian frame kernel is a different type of object from a complex-symmetric Majorana mass operator, and the boundary between them is policed throughout this paper.

1.2 Why a separate neutrino gate is required

A charged Dirac mass is a map between two inequivalent charged chiral representations. The physical mass problem is therefore naturally a singular-value problem for a rectangular or square chiral map, and the polar decomposition of that map exhausts the structure.

A neutral mass may additionally pair a Weyl field with a charge-conjugate carrier. The resulting bilinear is a symmetric form on a left-handed Nambu carrier. Its canonical diagonalisation is a unitary congruence — the Takagi problem — not an ordinary similarity transformation, and not the charged-fermion polar decomposition repeated with neutral labels. The change from a chiral map to a symmetric form is the mathematical signature of the new physical channel opened by neutrality.

The following questions must therefore be separated and answered in order, because each one silently presupposes the ones before it:

1. Does an independent sterile right-handed carrier exist?
2. Which charge-conjugate carriers are gauge admissible?
3. Which completion blocks are generated?
4. Does any exact continuous lepton-number symmetry survive?
5. What are the Takagi singular masses?
6. What determines the absolute scale?
7. What relative frame produces observable mixing?
8. What bridge connects the running operator to oscillation, beta-decay, lepton-number-violation and cosmological observables?

1.3 Central aim

To prove the necessary and sufficient structural conditions under which an independently derived neutral Weyl carrier admits a gauge-admissible charge-conjugation completion, a canonically normalised complex-symmetric mass operator, a unique Takagi mass spectrum and a rigorous Dirac–Majorana–pseudo-Dirac–seesaw classification, while preventing neutrality, oscillation data, sterile-carrier assumptions or Planck-scale insertion from being mistaken for a mass derivation.

1.4 Core novelty

NMCC-1 establishes:

1. a typed active–sterile–charge-conjugate carrier architecture;
2. a non-doubling theorem for the left-handed Nambu representation;
3. the symmetry of the neutral completion bilinear from fermionic statistics;
4. canonical normalisation for a nontrivial neutral kinetic metric;
5. a generalised Takagi mass theorem with a uniqueness clause;
6. a neutral common-scale no-go theorem;
7. an oscillation absolute-scale no-go theorem;
8. a pure-Dirac pair-degeneracy theorem with exact zero-mode counting;
9. an exact continuous lepton-number criterion;
10. a nested, invariant-subspace completion taxonomy;
11. a pre-breaking gauge-admissibility theorem;
12. an exact algebraic Schur-complement theorem, identified as a non-unitary congruence;
13. a controlled seesaw interpretation theorem with explicit error order;
14. rank, nullity and determinant constraints;
15. a Planck–closure neutral-completion normal form;
16. an ultraviolet factorisation non-uniqueness theorem with a quantified degeneracy group;
17. a charged-lepton-relative mixing theorem;
18. a Majorana-phase oscillation-invisibility theorem;
19. basis covariance, perturbative stability and numerical non-insertion firewalls.

1.5 Standard mathematics and VERSF-specific content

Standard mathematics includes conjugate representations, positive Hermitian forms, complex symmetric matrices, the Autonne–Takagi factorisation, singular values, Schur complements, rank inequalities, unitary congruence, perturbation theory for singular values and effective quadratic actions. No priority is claimed for any of it.

The VERSF-specific burden is to derive the physical neutral carriers, their completion maps, the origin and scale of each mass block, the Planck or closure descent factors, the relative charged-lepton/neutrino frame and the observable matching. Every place where that burden has not yet been discharged is named as a debt in Appendix G.

1.6 Inheritance matrix

Each upstream gate is graded by exactly what it supplies to NMCC-1 — and by what it does not.

Prior gate	Value to NMCC-1
Electroweak Representation and Connection Closure	Load-bearing: active doublet carrier, minimal anomaly-neutral sterile admissibility within the stated representation ansatz, pre-breaking gauge admissibility, B−L firewall
Electroweak Flavour-Frame Operator	Load-bearing: relative-frame mixing $U_{CC} = U_{eL}^\dagger P_A U_\nu$, weak-commitment structure, Majorana caveat
Weak-Commitment Neutrino Operator	Load-bearing: scale–shape separation; Hermitian frame kernel versus complex-symmetric mass operator typing
Projected Weak-Doublet Closure Hamiltonian / Projection Theorem	Strong support: the frame kernel descends from H_{cl} via $H_W = P_W^\dagger H_{cl} P_W$ rather than being invented
Weak-Commitment Closure Kernel	Strong conditional support: candidate active-frame texture and its invariants
Projected Leakage Support Trace	Strong conditional support: $\beta = \sqrt{3}/20$ with the denominator derived as a projected support trace; explicitly mixing-only
First Projection Audit	Provenance and audit architecture; not a mass theorem
C_3 Substrate Reduction / CKM Curvature	Methodological pattern only: separation of exact algebra, conditional substrate premises and fitted branch choices
Neutrinos as Entropy Carriers	Outside the inheritance chain: transport and decoherence phenomenology, not mass completion or charge conjugation
From Substrate CAR Algebra to Relativistic Fermion Fields	Load-bearing: field-level CAR algebra, frequency sectors, carrier-versus-conjugate distinction; its reserved antiparticle programme is completed here, with the Dirac expansion corrected to independent b, d families
From Schrödinger to Dirac	Load-bearing: first-order relativistic spinor carrier and orientation-coupling mass bilinear; substrate provenance of M_D; no species content
On the Equivalence of the κ-Field and the Commitment Field	Conditional: scalar commitment mode and candidate scale $m_\kappa \simeq 2.8$ meV; Lagrangian-uniqueness and nonlinear-exclusion claims explicitly not inherited

No entry in this matrix closes the sterile census, the Dirac-versus-Majorana selection, any block M_L, M_D, M_R, any descent factor Φ_L, Φ_R, or any absolute mass. Those remain the open debts of Appendix G.

---

2. Typed Active, Sterile and Charge-Conjugate Carriers

2.1 Active carrier

Let 𝒳_νL be the active left-handed neutrino flavour carrier inherited from the lepton-doublet occupancy map. At the present gate its rank is three, with independently fixed world-to-framework addresses assigned before any neutrino mass ordering is imposed.

2.2 Sterile right-handed carrier

Let 𝒳_NR be a candidate right-handed neutral carrier of rank r. Its existence, rank and multiplicity are not consequences of the three active neutrino addresses. They require an independent VERSF census or closure theorem.

Inherited admissibility (conditional). The electroweak representation-closure gate establishes an admissibility statement whose scope must be stated exactly: arbitrarily many gauge singlets can be added without disturbing Standard Model gauge anomalies, and other anomaly-free chiral extensions exist, so "uniqueness" holds only inside a declared representation ansatz. NMCC-1 therefore inherits a statement that is stronger than bare invention and weaker than existence:

Within the inherited minimal representation ansatz, N_R ~ (**1**, **1**, 0) is the minimal electroweak-singlet chiral carrier that permits the conventional neutrino Dirac bridge without contributing to Standard Model gauge anomalies. The ansatz *permits* this carrier; it does not *force* its physical existence, its generation multiplicity or its mass completion.

Admissibility is a licence to ask, not an answer. The same gate distinguishes ordinary electroweak neutrality from complete neutrality under every possible VERSF or B−L charge; that distinction becomes load-bearing in Section 12.4.

Premise N1 — Sterile-carrier independence. No right-handed neutrino carrier may be added merely because a Dirac or seesaw model requires one.

Falsifier: the number of sterile carriers is chosen to reproduce the observed number of light masses or mixing angles.

2.3 Charge-conjugate carrier

The charge conjugate of 𝒳_NR is a left-handed carrier

𝒳_NR^c ≅ conj(𝒳_NR),

transforming in the conjugate gauge representation. Charge conjugation is anti-linear on amplitudes and exchanges a representation with its conjugate.

For a true gauge singlet, the representation and its conjugate are equivalent, but the chirality label and the charge-conjugation typing remain distinct bookkeeping data. Equivalence of representations is not identity of carriers.

Premise N2 — Charge-conjugate typing. Every conjugate carrier used in a completion is typed by Theorem 1: representation, chirality and conjugation labels are tracked as independent data and are never collapsed by an appeal to neutrality.

2.4 Theorem 1 — Charge-Conjugate Carrier Typing

Let $\mathscr{H}_R$ be a right-handed Weyl carrier in gauge representation R. Its charge conjugate is a left-handed carrier in the conjugate representation $\bar{R}$. A bilinear between a left-handed carrier in representation S and $\mathscr{H}_R^c$ is gauge admissible only if the full completion operator, including any interface field in representation R_{int}, contains a singlet in

$$S \otimes \bar{R} \otimes R_{int}.$$

Proof. Gauge invariance of a local operator is precisely the statement that the operator transforms in the trivial representation. Charge conjugation maps R to $\bar{R}$ and flips chirality; the tensor product of the participating representations must therefore contain a singlet, and no property of any single factor — in particular electric neutrality of one component — can substitute for this condition on the product. ■

Consequence. Electric neutrality is not enough. Gauge representation, chirality and interface content all remain load-bearing.

---

3. Neutral Nambu Completion and the Non-Doubling Firewall

3.1 Left-handed Nambu carrier

Define

$$\mathcal{N}_L = \mathscr{H}_{\nu L} \oplus \mathscr{H}_{NR}^c.$$

Every coordinate in $\mathcal{N}_L$ is left-handed. This permits a uniform two-component treatment of all neutral mass terms: Dirac, Majorana and mixed completions become blocks of one symmetric form on one carrier.

3.2 Lemma — Chirality-flip bookkeeping

A right-handed Weyl carrier of rank r contributes 2r complex field components. Its left-handed charge conjugate contributes the same 2r components, related by an invertible anti-linear map. The map is a change of description, not a source of new components.

3.3 Theorem 2 — Nambu Non-Doubling Theorem

Writing N_R as its left-handed charge conjugate N_R^c does not introduce a second independent physical field. It is an invertible change of chiral description between the right-handed carrier and its charge-conjugate left-handed carrier, and the physical degree-of-freedom count of $\mathcal{N}\_L$ is exactly the sum of the counts of $\mathscr{H}\_\nu L$ and $\mathscr{H}\_NR$.

Proof. By the Lemma, the conjugation map is a bijection on field components. A doubling would require the direct sum of a carrier with an *independently propagating* copy; the direct sum in $\mathcal{N}\_L$ is between the active carrier and the conjugate description of the sterile carrier, which share no components with each other. ■

Falsifier: counting N_R and N_R^c as two independent sterile species.

3.4 Charge-conjugation completion is not a mass

The operation $\psi \mapsto \psi^c$ changes representation and chirality. It carries no unit of mass and does not select a bilinear coefficient. Charge conjugation is therefore a typing operation — it declares which bilinears may be written — not a scale anchor and not a dynamical mass-generation mechanism. Every mass claim requires, in addition, a gauge-admissible completion operator and an independently derived dimensionful scale.

3.5 Field-level inheritance: completing the reserved antiparticle programme

The substrate-CAR gate constructs relativistic fermion fields from the substrate CAR algebra but explicitly reserves antiparticle structure and charge conjugation for a later completion, marking its negative-frequency V_α modes as placeholders. The present section discharges that reservation for the neutral sector: it types the conjugate modes, corrects the general field expansion, and separates three structures that are routinely conflated.

3.6 The corrected general Dirac expansion

The general Dirac field requires *independent* particle and antiparticle CAR families:

$\psi(x) = \Sigma\_s \int d\Pi\_p$ [$b\_s(p)$ $u\_s(p)$ $e^{\wedge}(-ip\cdot x)$ + $d\_s^\dagger(p)$ $v\_s(p)$ $e^{\wedge}(ip\cdot x)$],

with $\{b, b^\dagger\}$ and $\{d, d^\dagger\}$ independent CAR sectors. An expansion using a single family $a, a^\dagger$ for both terms is not the general Dirac construction: it is either an incomplete placeholder or a silent move toward a Majorana identification before that identification has been derived (falsifier F45).

Sign and self-conjugacy correction. The momentum spinors obey

$(\not{p} - m)\, u(p) = 0$, $(\not{p} + m)\, v(p) = 0$, but both complete field modes — $u_s(p)e^{(-ip\cdot x)}$ and $v_s(p)e^{(ip\cdot x)}$ — satisfy the *same* field equation

$$(i\gamma^\mu \partial_\mu - m)\, \psi = 0.$$

The field equation therefore requires both positive- and negative-frequency solution sectors but does not identify their CAR operators. A spinor basis may be chosen such that

$$v_s(p) = C\, \bar{u}_s^T(p);$$

this is a relation between charge-conjugate *mode functions* and may be used in both Dirac and Majorana quantisations. Majorana self-conjugacy requires the additional *operator-level* identification of the independent particle and antiparticle sectors, so that $\psi^c = e^{(i\eta_M)}\, \psi$. It is this identification — not the mode-spinor relation alone, and not the field equation — that constitutes the Majorana completion (falsifier F46). The field equation is neutral between the Dirac and Majorana completions, which is exactly why the choice between them must be derived.

3.7 Theorem 2a — Field-Level Charge-Conjugation Trichotomy

Three structures must be distinguished; none implies the next.

1. **One-particle conjugation.** An antilinear map $j_C : \mathcal{H}_R \to \mathrm{conj}(\mathcal{H}_R)$ sends a right-handed carrier into its left-handed conjugate representation. This acts on classical wavefunctions and representation labels.
2. **Fock-space charge conjugation.** A *unitary* operator $\mathcal{C}$ on Fock space exchanges the particle and antiparticle sectors:

$$\mathcal{C}\, b_s(p)\, \mathcal{C}^{-1} = \eta_C\, d_s(p), \quad \mathcal{C}\, d_s(p)\, \mathcal{C}^{-1} = \eta_C^*\, b_s(p).$$

The map on classical wavefunctions is antilinear while its quantum implementation is unitary; conflating the two levels is a category error.

3. **Majorana self-conjugacy.** A Majorana field satisfies

$$\psi^c = e^{(i\eta_M)}\, \psi,$$

which consistently *identifies* the previously independent particle and antiparticle sectors, up to the Majorana phase. This is an additional physical completion condition — never a consequence of electric neutrality, never a consequence of the field equation, and never a consequence of choosing a charge-conjugate mode-spinor basis (Section 3.6).

Proof. (1) is representation theory on the one-particle space and exists for every Weyl carrier. (2) requires a Fock space with two CAR sectors to exchange. CAR preservation alone does not fix the implementation — time reversal preserves the CAR relations *antiunitarily* — so the discriminator must be stated: charge conjugation is linear on amplitudes at the second-quantised level, commuting with i and with the time-evolution phase rather than conjugating them, and this linearity is what forces $\mathcal{C}$ to be unitary. It is precisely the opposite behaviour to the antilinear one-particle map of clause (1), which is the content of the antilinear-classical/unitary-quantum distinction. (2) presupposes (1) but adds the second-quantised structure. (3) is a constraint equation that halves the independent CAR content and is consistent only when the b and d sectors have identical quantum numbers under every exact charge; imposing it is a physical choice with the gauge and B−L obligations of Sections 12 and 12.4. Each level adds independent structure, so no level implies its successor. ■

3.8 Theorem 2b — CAR-Sector Completion Classification

At the field level, the neutral completion classes of Theorem 10 correspond to:

- **Dirac case:** b and d are independent CAR sectors, related by $\mathcal{C}$ but never identified; the protecting continuous charge of Theorem 9 rotates them oppositely.
- **Majorana case:** the particle and antiparticle sectors are consistently identified, $\psi^c = e^{i\eta_M} \psi$, up to a Majorana phase; the independent CAR content is halved.
- **Nambu notation:** writing both sectors as left-handed coordinates on $\mathcal{N}_L$ is a change of description with no identification imposed — it is not a physical doubling (Theorem 2) and not a Majorana identification (this theorem). The notation is completion-neutral; the completion is selected by the block structure, not by the bookkeeping.

Proof. The Dirac clause is Section 3.6 with the continuous charge of Theorem 9 acting as $b \mapsto e^{i\alpha}b$, $d \mapsto e^{-i\alpha}d$. The Majorana clause is level 3 of Theorem 2a, and the halving of CAR content is the field-level image of the pairwise Takagi structure collapsing to single self-conjugate modes. The Nambu clause is Theorem 2 lifted to Fock space: the conjugate description of the d sector introduces no identification between sectors. ■

---

4. Symmetric Neutral Completion Bilinears

4.1 General quadratic mass term

On $\mathcal{N}_L$, the most general local Lorentz-scalar quadratic neutral mass term is

$$\mathcal{L}_{\nu,m} = -\tfrac{1}{2} \mathcal{N}_L^T C \mathcal{B}_\nu \mathcal{N}_L + \text{h.c.},$$

where C denotes the spinor charge-contraction matrix and $\mathcal{B}_\nu$ has mass dimension one.

4.2 Theorem 3 — Symmetric Completion Theorem

Fermionic statistics require

$$\mathcal{B}_\nu^T = \mathcal{B}_\nu.$$

Proof. Write the bilinear in two-component form, $\mathcal{B}_{\nu,ij}\, \varepsilon_{\alpha\beta}\, n_i^\alpha n_j^\beta$, where $\varepsilon_{\alpha\beta}$ is the antisymmetric Lorentz-invariant spinor tensor and the n_i are Grassmann-valued left-handed components. Exchanging the flavour labels $i \leftrightarrow j$ exchanges the spinor factors, producing one minus sign from the antisymmetry of $\varepsilon_{\alpha\beta}$ and a second minus sign from Grassmann anticommutation. The spinor contraction is therefore *symmetric* under flavour exchange, so only the symmetric part of $\mathcal{B}_\nu$ contributes: the antisymmetric flavour part vanishes identically against the contraction. Without loss of generality $\mathcal{B}_\nu^T = \mathcal{B}_\nu$. ■ (Appendix B gives the index computation in full.)

Remark. Symmetry of $\mathcal{B}_\nu$ is a theorem of statistics, not a modelling choice. Any construction in which the neutral completion matrix fails to be symmetric has made a bookkeeping error upstream, and its spectrum is meaningless (falsifier F5).

4.3 Block decomposition

Relative to $\mathcal{N}_L = \mathcal{H}_{\nu L} \oplus \mathcal{H}_{NR^c}$,

$$\mathcal{B}_\nu = \begin{pmatrix} M_L & M_D \\ M_D^T & M_R \end{pmatrix},$$

with

$M_L^T = M_L$, $M_R^T = M_R$.

The off-diagonal block M_D need not be symmetric; the symmetry of $\mathcal{B}_\nu$ forces only its transpose to appear in the conjugate position.

4.4 Type firewall

- M_D is a chiral bridge between distinct active and sterile carriers.
- M_L is a symmetric completion on the active carrier.
- M_R is a symmetric completion on the sterile charge-conjugate carrier.

These blocks have different gauge requirements, different admissible origins and different natural scales. They may not be conflated, even when they contribute to the same low-energy effective operator.

---

5. Kinetic Metrics and Canonical Nambu Normalisation

5.1 General neutral quadratic action

At a declared matching scale $\mu\star$, let

$\mathscr{L}_\nu^{(2)} = i\, \mathcal{N}_L^\dagger K_\nu \bar{\sigma}^\mu D_\mu \mathcal{N}_L - \tfrac{1}{2} \mathcal{N}_L^T C \mathscr{B}_\nu \mathcal{N}_L + \text{h.c.},$

with

$K_\nu = K_\nu^\dagger > 0.$

5.2 Premises

Premise N3(a) — Positive kinetic form. K_ν is positive definite on the retained finite-dimensional carrier.

Premise N3(b) — Common scale and scheme. K_ν and $\mathscr{B}_\nu$ are evaluated in the same effective theory, scheme and matching convention.

Premise N3(c) — Typed kinetic compatibility. Before electroweak breaking, the kinetic metric preserves the gauge- and charge-typed active/sterile decomposition: $[P_A, K_\nu] = 0$, and $[Q_L, K_\nu] = 0$ whenever standard lepton number is claimed. Without N3(c), $K_\nu^{-1/2}$ mixes active and sterile coordinates, the block labels M_L, M_D, M_R lose their typed meaning after canonical normalisation, and any derivation must instead carry the transformed current and charge operators through the normalisation explicitly.

Together, N3(a)–(c) constitute master premise N3 of Sections 13.1 and 27. The label N2 is reserved throughout for charge-conjugate typing (Section 2.3); local declarations and the master table use one numbering.

5.3 Theorem 4 — Canonical Nambu Normalisation

Let $K_\nu^{1/2}$ be the unique positive Hermitian square root of K_ν, and define

$\mathcal{N}_L^c = K_\nu^{1/2} \mathcal{N}_L.$

Then the kinetic term is canonical, and the neutral mass operator is

$$\mathbf{M_\nu^c = K_\nu^{-T/2} \mathscr{B}_\nu K_\nu^{-1/2},}$$

where $K_\nu^{-T/2} \coloneqq (K_\nu^{-1/2})^T$. Moreover,

$(M\_\nu^c)^T = M\_\nu^c$.

Proof. Substituting $\mathcal{N}\_L = K\_\nu^{(-1/2)}\, \mathcal{N}\_L^c$ into the kinetic term gives $i\, \mathcal{N}\_L^{c\dagger}\, K\_\nu^{(-1/2)}\, K\_\nu\, K\_\nu^{(-1/2)}\, \bar{\sigma}^\mu D\_\mu\, \mathcal{N}\_L^c = i\, \mathcal{N}\_L^{c\dagger}\, \bar{\sigma}^\mu D\_\mu\, \mathcal{N}\_L^c$, since $K\_\nu^{(-1/2)}$ is Hermitian. The mass bilinear becomes $-\tfrac{1}{2}\, \mathcal{N}\_L^{cT}\, C\, [K\_\nu^{(-T/2)}\, \mathscr{B}\_\nu\, K\_\nu^{(-1/2)}]\, \mathcal{N}\_L^c$ + h.c. Transposition gives

$$(M\_\nu^c)^T = K\_\nu^{(-T/2)}\, \mathscr{B}\_\nu^T\, K\_\nu^{(-1/2)} = M\_\nu^c,$$

using Theorem 3. ■

Remark. Because K_ν is Hermitian, $K\_\nu^{(-T/2)} = (K\_\nu^{(-1/2)})^*$ — the canonical sandwich is a conjugate congruence, exactly the transformation law of a symmetric bilinear form, and the symmetry of the completion survives normalisation automatically.

5.4 Canonical-normalisation firewall

The physical neutral mass matrix is M_ν^c, not the raw bilinear $\mathscr{B}\_\nu$, unless $K\_\nu = I$. Raw spectral data of $\mathscr{B}\_\nu$ are not invariant under nonunitary field normalisation and carry no physical meaning on their own (falsifier F6).

---

6. Physical Neutral Masses as Generalised Takagi Singular Values

6.1 Lemma — Autonne–Takagi factorisation

Every finite-dimensional complex symmetric matrix $M = M^T$ admits a factorisation

$$\mathbf{U^T\, M\, U = D = diag(m_1, \ldots, m_n),\ m\_i \geq 0,}$$

with U unitary; equivalently $M = U^*\, D\, U^\dagger$. The diagonal entries are exactly the singular values of M. Existence follows from the singular-value decomposition together with symmetry; the diagonal factor is unique as an unordered multiset because it coincides with the singular-value multiset, which is determined by the spectrum of $M^\dagger M$.

6.2 Theorem 5 — Neutral Physical-Mass Theorem

The physical tree-level neutral masses are the Takagi singular values m_i of M_ν^c, equivalently the ordinary singular values of M_ν^c:

$$m\_i^2 \in \sigma(\, (M\_\nu^c)^\dagger\, M\_\nu^c\,).$$

They are nonnegative, invariant under unitary changes of canonical neutral flavour frame, and unique as an unordered multiset.

Uniqueness clause. The multiset $\{m_i\}$ is uniquely determined by the pair $(\mathscr{B}_\nu, K_\nu)$. The Takagi matrix U_ν is unique only up to (i) real orthogonal rotations within degenerate positive singular subspaces, (ii) signs on nonzero nondegenerate modes — for $U \mapsto UP$ with $P = \mathrm{diag}(e^{i\alpha_i})$ and fixed $D > 0$, $P^T D P = D$ forces $e^{2i\alpha_i} = 1$ — and (iii) arbitrary phases on exact zero modes. No physical claim may depend on the residual freedom.

6.3 Generalised eigenvalue form

The masses solve

$$\mathscr{B}_\nu^\dagger K_\nu^{-T} \mathscr{B}_\nu v_i = m_i^2 K_\nu v_i,$$

and the eigenvectors may be chosen K_ν-orthonormal. This form makes explicit that the physical spectrum is a joint invariant of the pair $(\mathscr{B}_\nu, K_\nu)$, never of $\mathscr{B}_\nu$ alone. (Proof in Appendix C.)

6.4 Majorana field representatives

For each nonzero Takagi mode, a four-component self-conjugate representative may be formed from the corresponding left-handed mass eigenfield and its charge conjugate. This field representation does not imply that the full theory is invariant under charge conjugation: self-conjugacy of a free mass eigenfield and C-invariance of interactions are logically independent statements (Section 22.4).

6.5 Degeneracy and phase caveats

If the Takagi singular values are nondegenerate, the mass eigendirections are fixed up to signs or convention-dependent phases compatible with the Majorana condition. Degenerate subspaces admit further orthogonal rotations and require additional commuting observables for branch identification. Branch labels assigned inside an unresolved degenerate subspace are not stable data (Section 25.4).

---

7. Common-Scale and Oscillation-Scale No-Go Theorems

7.1 Theorem 6 — Neutral Common-Scale Orbit

For any $a > 0$, the rescaling

$\mathscr{B}\_\nu \mapsto a\, \mathscr{B}\_\nu$

implies

$M\_\nu^c \mapsto a\, M\_\nu^c$, $m\_i \mapsto a\, m\_i$,

while preserving Takagi projectors, all mass ratios and every homogeneous degree-zero spectral invariant.

Proof. The canonical sandwich of Theorem 4 is linear in $\mathscr{B}\_\nu$, and singular values are absolutely homogeneous of degree one. Projectors onto singular subspaces are degree-zero. ■

Consequence. Purely dimensionless neutral-flavour data cannot determine an absolute mass scale. An anchor must be dimensionful and independently derived.

7.2 Oscillation Hamiltonian

For relativistic active propagation in vacuum,

$H\_osc = (1/2E)\, U\_act\, \mathrm{diag}(m\_i^2)\, U\_act^\dagger$,

up to a multiple of the identity, and with the appropriate active projection applied if sterile modes exist.

7.3 Theorem 7 — Oscillation Absolute-Scale No-Go

For any $\mu_0^2 \geq 0$, define $m\_i'^2 = m\_i^2 + \mu_0^2$. Then

$H\_osc' = H\_osc + (\mu_0^2/2E)\, I$

on the propagating subspace. The additional term produces only a common phase in the evolution operator and leaves every oscillation probability unchanged — in vacuum, and in matter once the same identity shift is carried through the full matter Hamiltonian, since matter potentials add to H_osc and commute with multiples of the identity.

Oscillations therefore cannot determine the lightest mass or the common absolute mass offset.

Proof. $P(\alpha\to\beta; L) = |\langle\beta|\, \exp(-i\, H\_osc\, L)\, |\alpha\rangle|^2$. Replacing H_osc by H_osc + cI multiplies the amplitude by the pure phase $\exp(-icL)$, of unit modulus. ■

7.4 Corollary — Existence of a symmetric shifted completion

Given a Takagi matrix U_ν,

$M\_\nu' = U\_\nu^*\, \mathrm{diag}(\, \sqrt{(m\_i^2 + \mu_0^2)}\,)\, U\_\nu^\dagger$

is complex symmetric and has the same Takagi eigendirections with shifted masses. The oscillation degeneracy therefore lies *within* the admissible symmetric mass-operator class: no structural principle of the neutral sector excludes the shifted spectra.

7.5 Consequences

Oscillation data do not by themselves determine:

- the lightest neutrino mass;
- the sum of masses;
- the Dirac-versus-Majorana character;
- the Majorana phases;
- the heavy sterile scale;
- the ultraviolet factorisation of the effective mass operator.

Every one of these therefore requires an independent derivation or an independent observable, and each is tracked as a named debt in Appendix G.

Part II — Dirac, Majorana and Gauge-Admissible Completion

8. Block Structure of the Neutral Mass Operator

8.1 General block form

The canonical operator may be written, after carrying the kinetic normalisation into each block, as

$$M_{\nu^c} = \begin{pmatrix} M_L & M_D \\ M_D^T & M_R \end{pmatrix}.$$

The classification below is invariant under unitary congruences that preserve the active–sterile typing.

8.2 Completion channels

1. **Active self-completion:** M_L.
2. **Active–sterile chiral completion:** M_D.

3. **Sterile self-completion:** M_R.

The mere mathematical possibility of a block does not establish its gauge or VERSF admissibility. Admissibility is settled in Section 12; here the blocks are only typed.

---

9. Pure Dirac Pairing and Pairwise Takagi Degeneracy

9.1 Pure Dirac operator

If $M_L = 0$ and $M_R = 0$, then

$$M_D^{(N)} = \begin{pmatrix} 0 & M_D \\ M_D^T & 0 \end{pmatrix},$$

with M_D of size $n_L \times n_R$.

9.2 Theorem 8 — Dirac Pair-Degeneracy Theorem

Let the nonzero singular values of M_D be $s_1, \ldots, s_q$, where $q = \mathrm{rank}(M_D)$. Then the nonzero Takagi singular values of $M_D^{(N)}$ are

$$\mathbf{s_1, s_1, s_2, s_2, \ldots, s_q, s_q},$$

each appearing exactly twice, and the number of zero Takagi values is exactly

$$n_L + n_R - 2q.$$

Proof. Take the singular-value decomposition $M_D = U_L \Sigma U_R^\dagger$, with Σ the $n_L \times n_R$ diagonal matrix of singular values. The block unitary

$$W = \begin{pmatrix} U_L & 0 \\ 0 & U_R^* \end{pmatrix}$$

acts by congruence:

$$W^T M_D^{(N)} W = \begin{pmatrix} 0 & \Sigma \\ \Sigma^T & 0 \end{pmatrix}.$$

Reordering coordinates pairs the a-th active and a-th sterile directions and reduces the operator to a direct sum of 2×2 symmetric blocks $\begin{pmatrix} 0 & s_a \\ s_a & 0 \end{pmatrix}$ for $a = 1, \ldots, q$, plus a zero block of dimension $n_L + n_R - 2q$. Each 2×2 block has $(\text{block})^\dagger(\text{block}) = s_a^2 I_2$, so both of its Takagi values equal s_a; explicitly, the unitary $(1/\sqrt{2})\cdot\begin{pmatrix} 1 & i; & 1 & -i \end{pmatrix}$ Takagi-diagonalises it to diag(s_a, s_a). Congruence by the unitary W and the permutation preserves Takagi values, which proves the multiset claim and the zero-mode count. ■ (Appendix D restates the computation.)

9.3 Dirac reconstruction

Each equal pair may be combined into one four-component Dirac field. Pairwise Takagi degeneracy is therefore the neutral Nambu signature of an exact Dirac completion of the quadratic mass sector — with full Dirac character requiring, in addition, that the protecting charge extend to the interactions.

9.4 Firewall

A pair of nearly equal masses is not automatically Dirac. Exact Dirac character additionally requires the protecting continuous charge of Theorem 9 and the corresponding interaction structure. Degeneracy is necessary, not sufficient (falsifier F14).

---

10. Lepton-Number Symmetry Criterion

10.1 General continuous charge

Let $Q = Q^\dagger$ generate a continuous phase transformation

$\mathcal{N}_L \mapsto \exp(i\alpha Q)\, \mathcal{N}_L$, $\alpha \in \mathbb{R}$.

10.2 Theorem 9 — Neutral Continuous-Symmetry Criterion

The quadratic action is invariant under this one-parameter group for all α if and only if

$[Q, K_\nu] = 0$ and $Q^T \mathcal{B}_\nu + \mathcal{B}_\nu Q = 0$.

Proof. *(Kinetic term.)* Invariance requires $\exp(-i\alpha Q)\, K_\nu \exp(i\alpha Q) = K_\nu$ for all α, which holds iff $[Q, K_\nu] = 0$. *(Mass term.)* The bilinear transforms as $\mathcal{B}_\nu \mapsto \exp(i\alpha Q^T)\, \mathcal{B}_\nu \exp(i\alpha Q)$. Differentiating at $\alpha = 0$ gives the necessary condition $Q^T \mathcal{B}_\nu + \mathcal{B}_\nu Q = 0$. Conversely, if the derivative condition holds, then $d/d\alpha\, [\exp(i\alpha Q^T)\, \mathcal{B}_\nu \exp(i\alpha Q)] = i \exp(i\alpha Q^T)(Q^T \mathcal{B}_\nu + \mathcal{B}_\nu Q) \exp(i\alpha Q) = 0$ for every α, so the finite transformation is a symmetry: first-order invariance exponentiates exactly for a one-parameter group. ■

10.3 Standard lepton number

On the left-handed Nambu carrier, the standard Dirac lepton-number generator is

$$Q_L = \begin{pmatrix} I_{\nu L} & 0 \\ 0 & -I_{N R^c} \end{pmatrix}.$$

For the block mass matrix,

$$Q_L^T \mathcal{B}_\nu + \mathcal{B}_\nu Q_L = \begin{pmatrix} 2M_L & 0 \\ 0 & -2M_R \end{pmatrix}.$$

For the standard generator Q_L, the *mass bilinear* is therefore invariant precisely when

$M_L = M_R = 0$.

The complete quadratic action preserves standard continuous lepton number if and only if, in addition, the kinetic condition of Theorem 9 holds:

$[Q_L, K_\nu] = 0$.

Thus, provided the kinetic metric respects the active–sterile lepton-number grading (Premise N3(c)), standard lepton number is preserved exactly when the Majorana blocks vanish. A kinetic metric that mixes the +1 and −1 lepton-number sectors breaks the symmetry even at $M_L = M_R = 0$; the mass condition alone is not the symmetry condition.

10.4 Corollary — Majorana-block violation

A nonzero M_L or M_R violates the standard lepton-number assignment by two units. A different exact continuous charge may survive only if it is independently derived and satisfies Theorem 9; it may not be postulated to rescue a broken assignment.

---

11. Completion Taxonomy and Hierarchical Regimes

11.1 Majorana completion

A neutral mass eigenmode is Majorana-type when it is represented by one self-conjugate four-component field rather than by two independent Weyl carriers forming a Dirac pair.

This field property does not imply that weak interactions conserve charge conjugation or parity.

11.2 Pseudo-Dirac regime

If

$\|M_L\|, \|M_R\| \ll \|M_D\|$,

the exact Dirac pairs are split by small Majorana perturbations. The continuous lepton-number symmetry becomes approximate, and the Takagi spectrum consists of nearly degenerate pairs. For one generation with real positive parameters $\varepsilon_L, \varepsilon_R \ll m_D$, the pair is

$$m_\pm = m_D \pm (\varepsilon_L + \varepsilon_R)/2 + O(\varepsilon^2/m_D),$$

so the fractional splitting is controlled directly by the lepton-number-violating blocks. With complex blocks, in the convention where m_D is made real and positive by rephasing, the singular values are

$$m_\pm = m_D \pm |\varepsilon_L^* + \varepsilon_R|/2 + O(\varepsilon^2/m_D):$$

the splitting is controlled by the *phase-dependent* combination $|\varepsilon_L^* + \varepsilon_R|$, which can vanish at leading order — $\varepsilon_R = -\varepsilon_L^*$ — even with both lepton-number-violating blocks nonzero. The phases do not merely enter at the stated order; they can annihilate the leading-order splitting entirely. This is the converse trap to falsifier F14: near-degeneracy is not evidence of Dirac character, and an accidentally unsplit pair is not evidence that the Majorana blocks vanish.

11.3 Mixed regime

When the diagonal and off-diagonal blocks are comparable, no Dirac-pair interpretation is privileged. The physical modes are general Majorana mixtures of active and sterile coordinates.

11.4 Theorem 10 — Neutral Completion Taxonomy

The completion character of a canonically normalised symmetric neutral mass operator is assigned *per invariant subspace* — the subspaces singled out by the Takagi structure and any exact continuous charges — and the taxonomy is nested rather than flat:

1. **Exact Dirac class.** On the subspace, $M_L = M_R = 0$ and a protecting continuous charge satisfying Theorem 9 pairs every massive mode.
2. **Lepton-number-violating Majorana completion.** At least one diagonal block is nonzero on the subspace and no exact continuous charge protects the entire massive spectrum there. Within this single completion class, three *regimes* — limits, not disjoint classes — are distinguished by declared hierarchies:
    - **pseudo-Dirac limit:** the diagonal blocks are perturbatively small relative to M_D;
    - **seesaw limit:** a heavy block is parametrically separated and the light–heavy mixing is controlled (Premise N8);
    - **generic mixed regime:** neither hierarchy applies.
3. **Partial completions.** In multi-generation systems, exact Dirac pairs, Majorana modes and massless modes may coexist on different invariant subspaces. The operator as a whole then belongs to no single class, and assigning it a global label is an error.

The taxonomy is exhaustive; the regimes inside class 2 are limits of one completion and would not be mutually exclusive if treated as global classes; and the assignment on each subspace is determined by block structure, symmetry and hierarchy — never by electric neutrality.

Proof. The Takagi decomposition splits the carrier into singular subspaces, and any exact charge of Theorem 9 commutes with the corresponding projectors, by the following computation. The kinetic clause $[Q, K_\nu] = 0$ licenses passing to the canonical basis, where the mass clause reads $Q^T M + M Q = 0$ with $M = M_{\nu^c}$ and $Q = Q^\dagger$. Its Hermitian conjugate gives $Q M^\dagger = -M^\dagger Q^*$, and its complex conjugate gives $Q^* M = -M Q$; therefore

$$Q M^\dagger M = -M^\dagger Q^* M = M^\dagger M Q,$$

so $[Q, H] = 0$ with $H = M^\dagger M$, and Q commutes with every spectral projector of H — hence with the singular-subspace projectors. Subspace-wise assignment is therefore well defined. On each subspace, either the diagonal blocks vanish and a protecting charge exists (class 1) or not (class 2); this dichotomy is exhaustive. The class-2 regimes are defined by declared norm hierarchies on the same block data and therefore label limits of one object rather than partitioning it. Coexistence (class 3) is exhibited by block-diagonal examples — one exact Dirac pair direct-summed with one Majorana singlet — which no global label describes. Neutrality appears in no criterion. ■

---

12. Gauge Admissibility Before Electroweak Breaking

12.1 Active representation

Before electroweak breaking, ν_L is the neutral component of the lepton doublet

$$L \sim (\mathbf{2}, -\tfrac{1}{2})$$

under $SU(2)_L \times U(1)_Y$.

12.2 Dirac completion

A renormalisable Dirac bridge requires an independently existing

$$N_R \sim (\mathbf{1}, 0)$$

and an interface field with the Higgs-conjugate quantum numbers:

$$-\bar{L}\, Y_\nu\, \tilde{\Phi}\, N_R + \text{h.c.}$$

After the radial mode acquires a vacuum value,

$M_D = (v_{cl}/\sqrt{2})\, Y_\nu$

in the chosen convention.

Substrate provenance. The Schrödinger-to-Dirac gate shows that fermionic mass necessarily appears in a first-order relativistic operator with Clifford structure, in the orientation-coupling form $m c^2 \beta$, joining otherwise decoupled chiral channels. M_D is therefore not an arbitrary off-diagonal matrix: it is the species-resolved, gauge-admissible descendant of that general orientation-coupling bilinear. What the provenance supplies is the *form*; the existence of N_R, its rank, Y_ν and every numerical entry remain underived.

12.3 Active Majorana completion

A bare $\nu_L \nu_L$ term is not electroweak gauge invariant: with $L \sim (\mathbf{2}, -½)$, the product $L \otimes L$ transforms in $(\mathbf{3} \oplus \mathbf{1})$ at hypercharge −1, which contains no singlet. A left-handed Majorana block therefore requires a gauge-admissible completion. The minimal such operator is the dimension-five neutral completion built from two closure-interface doublets $H \sim (\mathbf{2}, +½)$, contracted as gauge singlets:

$\mathcal{O}_5^{ij} = (c_{ij}/\Lambda)\, (\varepsilon_{ab} L_i^a H^b)(\varepsilon_{cd} L_j^c H^d) + \text{h.c.},$

whose total hypercharge is $-½ - ½ + ½ + ½ = 0$, or an equivalent ultraviolet/VERSF completion carrying the same gauge transformation. After the radial mode acquires its vacuum value this induces a symmetric active block of order

$M_L \sim c\, v_{cl}^2 / \Lambda.$

This dimension-five Majorana operator is distinct from the renormalisable $\bar{L}\, \tilde{H}\, N_R$ Dirac bridge of Section 12.2 — the two carry different lepton-number and B−L charges, different scales and different admissibility obligations, and only their coexistence with a heavy M_R produces the $M_D M_R^{-1} M_D^T$ completion after electroweak breaking.

12.4 Sterile Majorana completion and the B−L firewall

A bare $N_R^T C N_R$ term is gauge admissible under the Standard Model gauge group only if N_R is a true singlet. Any additional VERSF charge — topological, orientational or substrate-representational — must also permit the bilinear. Naming a carrier "sterile" does not discharge this check.

Sterile Majorana B−L Firewall. The electroweak representation-closure gate shows that once ν_R is included, a *gauged* B−L grade becomes anomaly-admissible. The bilinear $N_R^T C N_R$ is a Standard Model singlet but carries $|\Delta B-L| = 2$. A bare M_R is therefore admissible only if at least one of the following holds:

1. B−L is not an exact gauged symmetry of the completed theory;
2. B−L is gauged but spontaneously broken, with M_R generated at or below the breaking scale;
3. the mass arises through a field or completion object carrying the compensating B−L charge.

Asserting a bare M_R while an exact unbroken gauged B−L is in force is falsifier F42. The same accounting applies to the active block: the dimension-five completion $\mathcal{O}_5$ also carries |ΔB−L| = 2.

Corollary — B−L rigidity. If an exact unbroken gauged B−L survives in the completed theory, then M_L = M_R = 0 identically, and by Theorem 10 the neutral completion is Dirac-class (if M_D exists) or massless. The Dirac-versus-Majorana question is thereby partly a question about the fate of B−L — which must be independently settled (Premise N12; debt D23), not assumed in either direction.

12.5 Theorem 11 — Gauge-Admissible Neutral Completion

Before electroweak breaking:

- M_D requires an independently admitted opposite-chirality neutral carrier and a gauge-compensating interface;
- M_L requires a higher-dimensional or ultraviolet gauge completion;
- M_R requires an independently admitted carrier neutral under every exact gauge and VERSF charge relevant to the bilinear, including any exact gauged B−L grade.

No block is licensed merely because it is algebraically useful.

Proof. Each clause is Theorem 1 applied to the representation content of the respective bilinear: $(\mathbf{2}, -\tfrac{1}{2}) \otimes (\mathbf{1}, 0)$ needs the $\tilde{\Phi}$ interface for a singlet; $(\mathbf{2}, -\tfrac{1}{2}) \otimes (\mathbf{2}, -\tfrac{1}{2})$ contains an SU(2) singlet at hypercharge −1 but no full electroweak singlet $(\mathbf{1}, 0)$, so a pair of H ~ $(\mathbf{2}, +\tfrac{1}{2})$ interface fields or an equivalent ultraviolet completion is mandatory; $(\mathbf{1}, 0) \otimes (\mathbf{1}, 0)$ is a Standard Model singlet, leaving only the VERSF-charge check. ■

---

13. The Neutrino Charge-Conjugation Completion Theorem

13.1 Premises

N1. The active carrier and every sterile carrier are independently derived. **N2.** The charge-conjugate carriers are correctly typed. **N3.** The kinetic metric is positive and evaluated at a common scale. **N4.** The completion bilinear is gauge admissible and complex symmetric. **N5.**

The origin and dimensionful scale of every nonzero block are independently specified. **N6.** No neutrino mass, splitting, ordering, mixing angle or CP phase is used to choose the carrier, block pattern, rank or scale. **N7.** A valid matching bridge is supplied before observable claims.

13.2 Theorem 12 — Neutrino Mass and Charge-Conjugation Completion

Under N1–N6, the canonically normalised neutral mass operator is

$$M\_\nu^c = K\_\nu^{(-T/2)} \begin{pmatrix} M\_L & M\_D \\ M\_D^T & M\_R \end{pmatrix} K\_\nu^{(-1/2)}\ , \qquad (M\_\nu^c)^T = M\_\nu^c.$$

It admits a Takagi factorisation

$U\_\nu^T M\_\nu^c U\_\nu = \mathrm{diag}(m\_i),\ m\_i \geq 0,$

which uniquely determines the unordered nonnegative tree-level mass spectrum. The Dirac, Majorana, pseudo-Dirac or mixed character — assigned per invariant subspace by the taxonomy of Theorem 10 — is determined by the block structure and the continuous-symmetry criterion of Theorem 9.

Charge conjugation is necessary to type the self-completion channels but is insufficient to determine their existence, their scale or the numerical spectrum.

Proof. Existence and symmetry of M_ν^c: Theorems 3 and 4 under N2–N4. Spectrum and uniqueness: Theorem 5 and its uniqueness clause. Classification: Theorems 8–10 under N1. Insufficiency of conjugation: Section 3.4 — the conjugation map fixes typing but contains no scale and selects no coefficient; a completion with all blocks zero satisfies every typing requirement and has an all-zero mass spectrum. ■

13.3 Outcome taxonomy

1. **Carrier completion only:** active and conjugate carriers are typed, but no mass block is derived.
2. **Structural mass completion:** the nonzero block pattern and Takagi spectrum are derived up to an overall scale.
3. **Absolute running-mass completion:** all dimensionful block scales are independently derived at a declared matching scale.
4. **Observable completion:** running, thresholds, interactions and self-energy effects are matched to measurable quantities.

NMCC-1 closes the theorem linking stages 1–3 but does not claim that the VERSF-specific inputs required for stage 3 have already been discharged. Stage 4 additionally requires N7.

Part III — Seesaw Structure and Planck-to-Neutrino Descent

14. Exact Quadratic Schur Complement

14.1 Heavy-block assumption

Assume M_R is invertible. Work at the algebraic level of the quadratic form defined by

$$\mathcal{B}_\nu = \begin{pmatrix} M_L & M_D \\ M_D^T & M_R \end{pmatrix}.$$

14.2 Theorem 13 — Exact Schur-Complement Identity

Completing the square in the sterile coordinate gives the exact identity

$$\mathbf{M_{\nu,eff} = M_L - M_D\, M_R^{-1}\, M_D^T.}$$

Equivalently, $\mathcal{B}_\nu$ factorises as the congruence

$$\mathcal{B}_\nu = W^T \begin{pmatrix} M_{\nu,eff} & 0 \\ 0 & M_R \end{pmatrix} W, \qquad W = \begin{pmatrix} I & 0 \\ M_R^{-1}\, M_D^T & I \end{pmatrix},$$

Proof. Define $N' = N + M_R^{-1} M_D^T \nu$. The quadratic form $\nu^T M_L \nu + 2\nu^T M_D N + N^T M_R N$ separates exactly into $N'^T M_R N' + \nu^T (M_L - M_D M_R^{-1} M_D^T)\, \nu$; expanding the block product $W^T(M_{\nu,eff} \oplus M_R)W$ reproduces $\mathcal{B}_\nu$ entry by entry. ■

14.3 Exactness firewall — congruence is not unitary congruence

The identity is exact for the algebraic mass quadratic form and for the zero-momentum tree-level elimination of the heavy coordinate. Two independent gaps separate it from the physical light spectrum.

First, the factorising matrix W is invertible but not unitary. Takagi singular values are invariants of *unitary* congruence only; a general congruence changes them. The Takagi spectrum of $\mathcal{B}_\nu$ is therefore not the union of the Takagi spectra of $M_{\nu,eff}$ and M_R. The identification holds only approximately, in the controlled regime of Theorem 14 — and the two factorisations must be kept distinct: the triangular congruence factor W of the present theorem is an exact but non-unitary completion-of-the-square, while Theorem 14 constructs a *unitary* light–heavy block-diagonalising transformation perturbatively, to arbitrary desired order. The unitary rotation and W are structurally different matrices and must not be identified term by term; their leading actions reproduce the same Schur-complement light block, which is why — and only why — the spectra approximately match, with corrections controlled by Θ.

Second, integrating out a propagating field in the full theory produces momentum-dependent higher-dimensional operators. The Schur complement is a zero-momentum statement.

For both reasons the Schur complement is not automatically the exact physical light-pole mass operator (falsifier F15).

---

15. Hierarchical Seesaw Interpretation and Active–Sterile Mixing

15.1 Mixing parameter

Define the dimensionless light–heavy mixing matrix

$\Theta = M_D M_R^{-1}$.

15.2 Premise N8 — Controlled hierarchy

A seesaw interpretation requires

$\|\Theta\| \ll 1$

in an explicitly declared operator norm, together with a heavy-block spectral gap separating the retained light and heavy modes.

15.3 Theorem 14 — Hierarchical Seesaw Interpretation

Under N8, a unitary congruence block-diagonalises the canonical neutral operator order by order in Θ. The light block is

$M_{light} = M_L - M_D M_R^{-1} M_D^T + \Delta M$,

the heavy block is

$M_{heavy} = M_R + O(\|\Theta\| \|M_D\|)$,

and the active–sterile mixing angles are of order $\|\Theta\|$. The correction ΔM must be stated relative to the light scale, not the Dirac scale:

- **Pure type I ($M_L = 0$):** $\Delta M = O(\|\Theta\|^2 \|M_{eff}\|) = O(\|\Theta\|^3 \|M_D\|)$, so the *relative* light-mass error is $O(\|\Theta\|^2)$. The one-generation expansion makes the order explicit:

  $m_{light} = -m_D^2/M_R + m_D^4/M_R^3 - \cdots, \Theta = m_D/M_R,$

  where the first correction is $O(\Theta^3 m_D) = O(\Theta^2 m_{light})$, in agreement with Appendix F.2.

- **General case ($M_L \neq 0$):** mixed terms in M_L, M_D and M_R^{-1} occur, with schematic bound $\Delta M = O(\|\Theta\|^2 \|M_L\| + \|\Theta\|^3 \|M_D\|)$, and the correction order must be tracked per term rather than quoted as a single power.

In this regime — and only in this regime — the Takagi values of the Schur complement approximate the light physical masses. The quoted *relative* error applies only on light singular subspaces whose leading Schur-complement singular value is nonzero and not parametrically suppressed by an independent cancellation, symmetry or rank condition; on suppressed or exactly null subspaces the relative error is undefined or enhanced, while the absolute operator bound remains valid.

Proof sketch. Write the block-diagonalising unitary as exp of an off-diagonal anti-Hermitian generator of order Θ and solve the off-diagonal vanishing condition order by order; the leading generator reproduces the shift N' of Theorem 13. The residual light-block correction begins at second order in the *rotation*, which acts on blocks already carrying their own powers of Θ — whence $O(\Theta^2 M_L)$ from the direct block and $O(\Theta^3 M_D)$ from the Dirac block, reducing to the stated pure type-I order when $M_L = 0$. Weyl's singular-value perturbation bound (Section 25.1) then converts the operator estimate into a mass estimate. ■

The precise correction order in observables additionally depends on the canonical kinetic structure, momentum dependence and matching convention; none of these is discharged here.

15.4 Type labels

- **Type-I-like:** $M_L = 0$, heavy $M_R \neq 0$.
- **Type-II-like:** direct active block $M_L \neq 0$ dominates.
- **Hybrid:** both terms contribute and may interfere, including destructively.

These are mechanism classes, not numerical predictions.

---

16. Rank, Nullity and Determinant Theorems

16.1 Theorem 15 — Type-I Rank Bound

For $M_L = 0$ and invertible M_R,

$M_{\nu,\text{eff}} = -M_D M_R^{-1} M_D^T$

obeys

$\text{rank}(M_{\nu,\text{eff}}) \le \text{rank}(M_D) \le \min(n_L, n_R)$.

Proof. $\text{rank}(ABC) \le \min\{\text{rank } A, \text{rank } B, \text{rank } C\}$, applied with $A = M_D$, $B = M_R^{-1}$, $C = M_D^T$, and $\text{rank}(M_D) = \text{rank}(M_D^T) \le \min(n_L, n_R)$. ■

16.2 Corollary — Forced massless modes

If $n_R < n_L$, then at least

$n_L - n_R$

light modes are exactly massless at tree level in the pure type-I completion.

Proof. By rank–nullity on the n_L-dimensional light space, $\text{nullity}(M_{\nu,\text{eff}}) = n_L - \text{rank}(M_{\nu,\text{eff}}) \ge n_L - n_R$. ■

This is an exact structural consequence. It cannot be repaired without adding another completion block, another carrier or radiative structure — and any such addition must itself pass the admissibility and non-insertion audits.

16.3 Determinant identity

If $n_L = n_R = n$ and M_D is invertible,

$\det(M_{\nu,\text{eff}}) = (-1)^n \cdot \det(M_D)^2 / \det(M_R)$.

Proof. Multiplicativity of determinants on $-M_D M_R^{-1} M_D^T$, together with $\det(M_D^T) = \det(M_D)$. ■

16.4 Nullity firewall

A massless mode may follow from rank deficiency, from an exact symmetry or from a tuned cancellation. These are three different mechanisms with three different stability properties, and they must not be conflated: only the first two are structurally protected.

---

17. Primitive and Planck Scale Normal Forms

17.1 Neutral primitive-scale normal form

If the available VERSF dimensionful primitives are $\hbar$, c, ξ, τ_s, Λ and the dimensionless invariants $\mathscr{J}$, every scalar neutral mass scale has the form

$$\mathfrak{m}_\nu = (\hbar/c\xi) \cdot F_\nu(c\tau_s/\xi, \Lambda\xi^2, \mathscr{J}).$$

This is dimensional analysis, not a derivation of F_ν, and it is recorded only to fix what a genuine derivation must produce.

17.2 Multiple-scale requirement

A general neutral completion may contain more than one independent dimensionful scale:

$$M_D \sim v_{cl},\ M_L \sim v_{cl}^2/\Lambda_L,\ M_R \sim \Lambda_R.$$

The common single-scale factorisation valid in a charged sector need not apply to the full neutral block, and importing it silently is a scale-typing error.

17.3 Planck basis

Convention. Throughout this paper, m_P denotes the *unreduced* Planck mass, $m_P = (\hbar c/G)^{1/2}$. The reduced mass $\bar{m}_P = m_P/\sqrt{8\pi}$ is smaller by a factor ≈ 5.01, so every downstream numerical use of the normal forms below must respect this declaration; silently switching conventions rescales the seesaw output by roughly a factor of five in either direction (falsifier F39).

If the conventional Planck mass is admitted,

$$M_R = m_P\, \Phi_R\, \hat{R}_\nu,$$

where Φ_R is a dimensionless scalar and $\hat{R}_\nu$ a dimensionless symmetric operator. The Planck mass determines neither object.

Similarly, a direct active completion may be written

$$M_L = (v_{cl}^2/m_P)\, \Phi_L\, \hat{L}_\nu.$$

Both are basis declarations. They assign every unknown a name and a type; they derive nothing.

17.4 The κ-closure basis and the direct-completion normal form

The κ-field/commitment-field equivalence gate supplies a scalar commitment mode with mass

$$\mathbf{m\_\kappa = \sqrt{(4/3)} \cdot \hbar/(c\xi),\ m\_\kappa\, c^2 = \sqrt{(4/3)} \cdot \hbar c/\xi.}$$

At the programme's current calibration $\xi \simeq 80\text{–}82\ \mu m$,

$m\_\kappa\, c^2 \simeq 2.78\text{–}2.85\ \text{meV}$,

which lies in a phenomenologically plausible range for the *lightest* eigenvalue or for a subleading contribution to the light-neutrino mass operator — although the atmospheric splitting requires at least one eigenvalue of roughly 0.05 eV, roughly eighteen times larger, so the κ scale cannot supply the full observed spectrum without an independently derived amplification or a coexisting route. The numerical proximity is genuinely interesting; it is not a derivation of a neutrino mass. The κ field is a *scalar* field, and its mass parameter cannot be relabelled as a fermion mass without deriving a neutral-sector coupling or completion map. The strongest safe object at this gate is therefore the κ-anchored direct-completion normal form

$$\mathbf{M\_L^{(k)} = m\_\kappa\, \Phi\_\kappa\nu\, \hat{L}\_\nu^{(k)},}$$

where $\Phi\_\kappa\nu$ is a dimensionless commitment-to-neutrino completion factor and $\hat{L}\_\nu^{(k)}$ a dimensionless symmetric active-sector texture — a *distinct operator* from the Planck-route texture $\hat{L}\_\nu$ of Section 18.1, with its own provenance (Section 17.5). Neither may be selected from neutrino masses or oscillation data (falsifier F49). Deriving $\Phi\_\kappa\nu$ and the scalar-to-neutral-fermion completion map is the named debt D24.

Non-inheritance note. The κ-field gate additionally claims a Lagrangian-uniqueness and nonlinear-exclusion result. That claim is not inherited here as closed: a local nonlinear equation such as $(\square + m^2)s + \lambda s^3 = \rho$ remains causal and local — it violates superposition without coupling spacelike-separated histories — and Lagrangians depending nonlinearly on $X = (\partial s)^2$ can still yield second-order field equations. NMCC-1 inherits the κ-field equation, the scale m_κ and the identification of a commitment-related scalar mode; it inherits nothing about uniqueness (falsifier F48).

17.5 Mass-symbol firewall

The programme now carries several superficially similar mass symbols that must remain typed and separate:

- $m\_spin^{(0)}$ — the uniform substrate spinorial mass appearing before species decomposition;
- m_κ — the commitment-field scalar mass;

- $\hat{L}\_\nu^{(k)}$ and $\hat{L}\_\nu$ — the κ-route and Planck-route direct-completion textures: distinct operators despite parallel roles, never to share a symbol;
- κ (commitment-field symbol) versus κ = 8/3 (the charged-lepton localisation exponent quarantined in Section 26.4): an accidental glyph collision with no physical content;
- M_L, M_D, M_R — the species-resolved neutral completion blocks;
- m_P — the Planck anchor;
- m_i — the physical neutrino Takagi masses.

No equality between any two may be assumed merely because their dimensions match (falsifier F47). Every asserted equality is a theorem obligation with its own bridge.

---

18. The Planck–Closure Neutral Completion Normal Form

18.1 Scale-typed block operator

The strongest Planck-specialised neutral normal form available at this gate is

$$\mathcal{B}\_\nu^{(P)} = \begin{pmatrix} (v\_cl^2/m\_P)\ \Phi\_L\ \hat{L}\_\nu & (v\_cl/\sqrt{2})\ Y\_\nu \\ (v\_cl/\sqrt{2})\ Y\_\nu^T & m\_P\ \Phi\_R\ \hat{R}\_\nu \end{pmatrix}.$$

Here

- Y_ν is a dimensionless active–sterile completion operator;
- $\hat{L}\_\nu$ and $\hat{R}\_\nu$ are dimensionless symmetric magnitude/texture operators;
- Φ_L and Φ_R are dimensionless descent or amplification invariants.

Every object remains proof-gated.

18.2 Theorem 16 — Planck-Seesaw Normal Form

If $\hat{R}\_\nu$ is invertible, the algebraic light operator is

$$\mathbf{M\_\nu,eff^{(P)} = (v\_cl^2/m\_P) \cdot [\ \Phi\_L\ \hat{L}\_\nu - (1/2\Phi\_R)\ Y\_\nu\ \hat{R}\_\nu^{-1}\ Y\_\nu^T\ ].}$$

Proof. Substitute the blocks of $\mathcal{B}\_\nu^{(P)}$ into Theorem 13: $M\_D\ M\_R^{-1}\ M\_D^T = (v\_cl^2/2)\ Y\_\nu\ (m\_P\ \Phi\_R\ \hat{R}\_\nu)^{-1}\ Y\_\nu^T = (v\_cl^2/2m\_P\ \Phi\_R)\ Y\_\nu\ \hat{R}\_\nu^{-1}\ Y\_\nu^T$. ■

The overall Planck-to-electroweak suppression v_cl^2/m_P is fixed once the two dimensionful bases are admitted. The physical light scale remains controlled by the independently derived dimensionless operator in brackets — which is exactly where the open physics lives.

18.2a Hybrid κ–Planck completion normal form

Admitting the κ-closure direct scale of Section 17.4 alongside the full Planck-anchored blocks of Section 18.1 gives the two-anchor, three-contribution hybrid

$$\mathbf{M_{\nu,eff} = m_\kappa\, \Phi_{\kappa\nu}\, \hat{L}_\nu^{(k)} + (v_{cl}^2/m_P)\, \Phi_L\, \hat{L}_\nu - (v_{cl}^2/2m_P\, \Phi_R)\, Y_\nu\, \hat{R}_\nu^{-1}\, Y_\nu^T.}$$

Nomenclature. The hybrid is a *two-anchor, three-contribution* architecture: two independent dimensionful anchors — m_κ and m_P — support three separately typed contributions, since the Planck anchor feeds both a suppressed direct contribution and a seesaw contribution. All downstream apparatus entries (scope box, certificate, debts, conclusion) use this name; "two routes" without qualification is retired.

The three contributions are typed separately: a closure-scale direct completion, a Planck-suppressed direct completion and an electroweak–Planck seesaw completion, with independently owed bridges (D24 for the κ coupling; D10 for the Planck descent). The two direct textures $\hat{L}_\nu^{(k)}$ and $\hat{L}_\nu$ are distinct operators with distinct provenance and must never share a symbol or be silently identified (Section 17.5). Any reduction of the hybrid — $\Phi_L = 0$ (absence of the Planck-direct contribution is independently established) or $\Phi_{\kappa\nu} = 0$ (pure Planck architecture) — is a *named premise choice*, not a default; setting a term to zero does not itself show that another mechanism supersedes it, and dropping a term silently is exactly the ansatz-versus-theorem error of Section 18.3. Relative dominance among the three contributions is a testable structural question: any contribution may dominate, they may be comparable, and they may interfere, including destructively. Dominance claims must be derived from the dimensionless factors, never selected from the observed spectrum.

18.3 Scale firewall

Setting $\Phi_R = 1$, $Y_\nu = I$, $\hat{R}_\nu = I$ or $\Phi_L = 0$ merely for simplicity is an ansatz, not a theorem. Selecting any of these values *from the observed neutrino scale* is numerical insertion (falsifier F28).

18.4 Planck-identification firewall

A VERSF metric-certification scale and the conventional Planck mass inferred from Newton's constant may be identified only after a matching theorem. Until then, the Planck normal form is an audit basis, not a completed substrate derivation (falsifier F38).

---

19. Ultraviolet Factorisation Non-Uniqueness

19.1 Theorem 17 — Seesaw Factorisation Degeneracy

Let

$S_\nu = M_D M_R^{-1} M_D^T$.

For any invertible matrix $X \in GL(n_R, \mathbb{C})$,

$M_D \mapsto M_D X, M_R \mapsto X^T M_R X$

leaves S_ν unchanged.

Proof. $M_D X (X^T M_R X)^{-1} X^T M_D^T = M_D X X^{-1} M_R^{-1} (X^T)^{-1} X^T M_D^T = M_D M_R^{-1} M_D^T$. ■

Degeneracy count — pre-quotient. The algebraic action of $GL(n_R, \mathbb{C})$ has real dimension $2n_R^2$ — 18 for $n_R = 3$. This is a *pre-quotient* degeneracy: the orbit directions include sterile-coordinate changes with no physical content, transformations that must be accompanied by a kinetic-metric retransformation to preserve canonical normalisation, and possible stabilisers that reduce individual orbit dimensions. The number of *physically inequivalent* ultraviolet parameters remaining after fixing kinetic normalisation, sterile-basis conventions, the heavy spectrum and stabilisers must be counted separately; Casas–Ibarra-type parametrisations are the standard display of exactly this residual freedom. The non-uniqueness conclusion survives the quotient intact: even after every convention is fixed, a continuous family of ultraviolet completions reproduces the same light operator.

19.2 Corollary — Overall scale degeneracy

The one-parameter subgroup $X = aI$ gives

$M_D \mapsto a M_D, M_R \mapsto a^2 M_R$,

which also leaves the type-I effective operator unchanged. The heavy scale itself is therefore invisible to the light operator.

19.3 Consequence

Even exact knowledge of the low-energy complex symmetric operator does not uniquely determine the heavy scale, the Dirac completion or the sterile texture. Additional VERSF structure or independent observables are necessary, and their absence must be declared, not papered over.

19.4 Inverse-seesaw firewall

A reconstructed heavy matrix obtained by algebraically solving the seesaw relation from observed low-energy data is an inverse fit, not a forward derivation. By Theorem 17, it is one arbitrary point on a continuous orbit of pre-quotient dimension $2n_R^2$ (falsifier F20).

Part IV — Mixing, CP, Stability and Audit Apparatus

20. Charged-Lepton-Relative Mixing and Active Projection

20.1 Charged-current frame

Let U_{eL} be the unitary transformation that places the canonically normalised charged-lepton left-handed carrier in its mass frame. Let U_ν be the Takagi matrix of the full neutral operator.

20.2 Active projection

If sterile modes exist, let

$$P_A : \mathcal{N}_L \to \mathscr{H}_{\nu L}$$

be the active projection. The charged-current mixing matrix is

$$\mathbf{U_{CC} = U_{eL}^\dagger P_A U_\nu.}$$

For exactly three active modes and no sterile admixture, this reduces to the usual unitary three-by-three form.

20.3 Theorem 18 — Relative-Frame Mixing Theorem

Observable lepton mixing is determined by the relative orientation of the charged-lepton left-handed mass frame and the active part of the neutrino Takagi frame. Neither sector alone determines the mixing matrix.

Proof. U_{CC} is a product of frame data from both sectors. Rotating either frame independently by a unitary V changes U_{CC} by V or $V^\dagger$; only the relative orientation is invariant under simultaneous re-derivations of the two sectors, so any single-sector claim about U_{CC} is frame-dependent and hence unphysical. ■

20.4 Consequences

- A neutrino matrix written in a chosen flavour basis does not predict mixing unless the charged-lepton frame is independently fixed.
- Sorting neutrino masses does not derive the electron, muon and tau flavour labels.
- With heavy or sterile mixing, the active submatrix $P_A U_\nu$ need not be exactly unitary, and nonunitarity is itself an observable subject to matching.

---

21. Weak-Commitment Frame Inheritance and Frame–Takagi Compatibility

21.1 Inherited frame objects

The weak-commitment programme derives a Hermitian neutrino frame operator

$$T_\nu = D_0 I + \varepsilon \mathcal{K}_\nu, \mathcal{K}_\nu = \mathcal{K}_\nu^\dagger,$$

in which the small weak-commitment factor ε controls the overall scale while the residual kernel $\mathcal{K}_\nu$ controls the eigenframe. The eigenvectors of T_ν are independent of ε. This is the structural reason very small neutrino masses and large lepton mixing can coexist: scale and shape live in different objects, and suppressing one does not suppress the other. NMCC-1 inherits this scale–shape separation as prior support for Sections 7 and 20.

21.2 The Hermitian/symmetric typing firewall

Three operator types must now be held apart, and the notation of this paper separates them deliberately:

- K_ν — the positive Hermitian *kinetic metric* (Section 5);
- $\mathcal{K}_\nu$ — the Hermitian *weak-commitment frame kernel* (inherited);
- $M_{\nu,eff} = M_{\nu,eff}^T$ — the complex-symmetric *light mass operator*, with Hermitian magnitude square $H_{\nu,m} \coloneqq M_{\nu,eff}^\dagger M_{\nu,eff}$.

A Hermitian kernel is diagonalised by a similarity transformation and carries no Majorana phases. A complex-symmetric mass operator is diagonalised by a Takagi congruence and carries phases that the Hermitian left-square object $H_{\nu,m}$ *cannot* reconstruct. Treating a PMNS-fitting Hermitian kernel as a Majorana mass matrix is therefore a type error, not an approximation (falsifier F41). The earlier flavour papers already flag this asymmetry; NMCC-1 promotes it to a firewall.

21.3 Theorem 19 — Weak-Commitment Frame–Takagi Compatibility

Let $\mathcal{K}_\nu = \mathcal{K}_\nu^\dagger$ be the independently derived weak-commitment kernel and let $M_{\nu,\mathrm{eff}} = M_{\nu,\mathrm{eff}}^T$ be the canonically normalised effective light-neutrino mass operator, with $H_{\nu,m} \coloneqq M_{\nu,\mathrm{eff}}^\dagger M_{\nu,\mathrm{eff}}$. Assume

$$[\mathcal{K}_\nu, H_{\nu,m}] = 0,$$

that the spectrum of $\mathcal{K}_\nu$ is nondegenerate, **and that the spectrum of $H_{\nu,m}$ is nondegenerate** (equivalently, the light Takagi masses are distinct). Then the eigenvectors of $\mathcal{K}_\nu$ determine the Takagi eigendirections of $M_{\nu,\mathrm{eff}}$ and their ordering. Diagonalising $H_{\nu,m}$ leaves arbitrary eigenvector phases, but imposing the Takagi equation

$$M_{\nu,\mathrm{eff}}\, u_i = m_i\, u_i^*$$

fixes those phases up to signs for nonzero modes; zero modes retain arbitrary phase freedom, and degenerate positive singular subspaces retain the real-orthogonal freedom of Theorem 5. Consequently the charged-current mixing matrix $U_{CC} = U_{eL}^\dagger P_A U_\nu$ — which for the three-dimensional effective light operator considered here restricts to

$$U_{CC} = U_{eL}^\dagger U_\nu$$

on the active sector (Section 20.2) — may inherit its mixing *angles* from the weak-commitment kernel, while the kernel determines none of the following: the neutrino masses, the common absolute scale, the Majorana phases, or the Dirac-versus-Majorana completion.

Degeneracy counterexample. The commutator condition with nondegenerate $\mathcal{K}_\nu$ is *not* sufficient on its own. Take

$$\mathcal{K}_\nu = \begin{pmatrix} 1 & 0 \\ 0 & 2 \end{pmatrix}, \qquad M = \begin{pmatrix} 0 & m \\ m & 0 \end{pmatrix}, \qquad H_{\nu,m} = |m|^2\, I.$$

Then $[\mathcal{K}_\nu, H_{\nu,m}] = 0$ trivially and $\mathcal{K}_\nu$ is nondegenerate, yet the $\mathcal{K}_\nu$ eigenbasis leaves M off-diagonal: inside the degenerate singular subspace of $H_{\nu,m}$, being an eigenvector of $H_{\nu,m}$ does not make a vector a Takagi vector of M. The nondegeneracy hypothesis on $H_{\nu,m}$ is therefore load-bearing, not decorative.

Proof. Commuting Hermitian operators are simultaneously diagonalisable. From the Takagi factorisation $M_{\nu,\mathrm{eff}} = U_\nu^* D_\nu U_\nu^\dagger$ one has $H_{\nu,m} = U_\nu D_\nu^2 U_\nu^\dagger$, so the Takagi columns are eigenvectors of $H_{\nu,m}$. If the spectrum of $H_{\nu,m}$ is nondegenerate, its eigenvectors are unique up to phase and ordering; by commutation the nondegenerate $\mathcal{K}_\nu$ eigenvectors coincide with them, hence fix the Takagi eigendirections up to the phase ambiguity of the $H_{\nu,m}$ eigenvectors — which the Takagi condition $M_{\nu,\mathrm{eff}}\, u_i = m_i\, u_i^*$ subsequently restricts to signs on nonzero nondegenerate modes, per the uniqueness clause of Theorem 5. The masses enter only through D_ν, which the kernel does not constrain, and the phases cancel from $H_{\nu,m}$ entirely. ■

Equivalent stronger conditions. Either of the following bypasses the spectral hypothesis on $H\_{\nu,m}$ directly:

- **Congruence-diagonality:** $V\_\nu^T M\_{\nu,eff} V\_\nu$ is diagonal, where V_ν is the $\mathcal{K}\_\nu$ eigenvector matrix — this simply *is* the statement that V_ν is a Takagi frame;
- **Intertwining:** $\mathcal{K}\_\nu^T M\_{\nu,eff} = M\_{\nu,eff} \mathcal{K}\_\nu$. In the eigenbasis of $\mathcal{K}\_\nu$ (real eigenvalues k_i), this reads $(k\_i - k\_j) M\_{ij} = 0$, forcing the symmetric mass matrix to be diagonal by congruence whenever $\mathcal{K}\_\nu$ is nondegenerate.

Sufficient bridge. A stronger, checkable sufficient condition is

$$H_{\nu,m} = m_{\nu,0}^2 \cdot f(\mathcal{K}_\nu),$$

with f positive and *strictly* monotone on the spectrum of $\mathcal{K}\_\nu$. Strict monotonicity does double duty: it makes the commutation automatic, and — because a strictly monotone function of a nondegenerate spectrum is nondegenerate — it discharges the spectral hypothesis on $H\_{\nu,m}$ as well, additionally transferring the kernel's ordering to the mass-squared ordering. It still determines neither the scale $m\_{\nu,0}$ nor the phases. Deriving the intertwining condition or the monotone map from VERSF structure is the named frame-bridge debt D22. Assuming either is bridge laundering (falsifier F44).

21.4 Corollary — Reconstruction degeneracy of a Hermitian frame

Let V_ν diagonalise the inherited Hermitian kernel with nondegenerate spectrum. The complete family of complex-symmetric mass operators sharing that mass-squared eigenframe is

$$M_{\nu,eff} = U_\nu D_\nu U_\nu^\dagger,\ U_\nu = V_\nu P,\ P = \mathrm{diag}(e^{i\alpha_1}, e^{i\alpha_2}, e^{i\alpha_3}),*$$

with $D\_\nu = \mathrm{diag}(m\_1, m\_2, m\_3)$ an *arbitrary* nonnegative diagonal — nondegenerate choices give exactly this family, while degenerate choices enlarge it further through the orthogonal freedom of Theorem 5. The Hermitian kernel therefore determines an eigenframe but leaves entirely free:

- the individual masses m_1, m_2, m_3;
- the absolute scale;
- the ordering (absent the monotone bridge);
- the independent Majorana phases encoded in the column-phase freedom α_i.

Proof. By Theorem 5 the Takagi frame of any such operator consists of eigenvectors of $H\_{\nu,m}$; nondegeneracy of the kernel spectrum fixes those up to phases and ordering, and every choice of (D_ν, P) yields a valid symmetric operator with the required frame. ■

Clarification. The phases in P parametrise *different* symmetric mass operators sharing the same $H\_{\nu,m}$ eigenframe; they are not residual rephasings of the Takagi vectors of one fixed nonzero mass operator, which by Theorem 5 are restricted to signs on nonzero nondegenerate modes. For three massive Majorana neutrinos, after the permitted charged-lepton rephasings only two relative Majorana phases remain physical. The field self-conjugacy phase η_M of Section 3.7 is a separate, convention-fixable field phase and must not be conflated with these two physical low-energy phases.

This is the exact, theorem-level boundary between the prior PMNS-frame papers and a mass derivation: a frame is a direction field, not a spectrum.

21.5 Conditional inheritance of the candidate kernel

The projection and leakage-support gates supply the candidate active-frame kernel not as an arbitrary matrix but as the endpoint of a derivation chain from the parent closure object:

$$H_{cl} \to H_W = P_W^\dagger H_{cl} P_W \to \mathcal{K}_\nu \to [\text{ bridge, debt D22 }] \to H_{\nu,m} \to M_{\nu,\mathrm{eff}}.$$

Every arrow except the bracketed one is supplied, conditionally, by prior gates. The bracketed arrow is precisely what NMCC-1 does not yet possess, and it is named rather than assumed.

The inherited candidate kernel has the Hermitian texture

$$\mathcal{K}_\nu = \begin{pmatrix} r_e & A & A(1+\beta e^{i\psi}) \\ A & r_s+\delta & B \\ A(1+\beta e^{-i\psi}) & B & r_s-\delta \end{pmatrix},$$

with conditional structural targets inherited under their own named premises:

$\rho_\odot = (r_s + B - r_e)/A = 3/2$, $A/B = 6/5$, $\delta = 0 + O(\beta^2)$, $\psi = 3\pi/4$, $\beta = \sqrt{3}/20$,

where the denominator 20 is derived as a projected support trace in the leakage-support gate — which itself states explicitly that it derives no mass mechanism and no absolute scale.

Inheritance grade. NMCC-1 inherits this kernel strictly as *an explicit conditional candidate for the active neutrino mixing frame* — never as the mass operator, never as the mass spectrum, never as a scale. Inserting its invariants into the mass blocks M_L, M_D, M_R would be mixing-to-mass insertion (falsifier F43).

---

22. Majorana Phases, CP and Oscillation Invisibility

22.1 Generalised CP transformation

In a canonical neutral basis, a generalised CP transformation has the form

$\mathcal{N}_L(x) \mapsto X \cdot i\sigma^2 \mathcal{N}_L^*(x_P)$,

with unitary X and x_P the parity-reflected point of the underlying geometry.

22.2 CP-invariance condition

The mass term is CP invariant iff

$$\mathbf{X^T M_{\nu^c} X = (M_{\nu^c})^*.}$$

Interaction terms must satisfy their own corresponding conditions; mass-term CP invariance alone certifies nothing about the interacting theory.

22.3 Theorem 20 — Majorana-Phase Invisibility in Standard Lepton-Number-Conserving Oscillations

Column phases associated purely with Majorana rephasings cancel from standard lepton-number-conserving flavour-oscillation probabilities.

Proof. Oscillation probabilities depend on the mixing matrix only through the quartic invariants

$U_{\alpha i} U_{\beta i}^* U_{\alpha j}^* U_{\beta j}$.

A Majorana rephasing acts as $U \mapsto U \cdot \mathrm{diag}(e^{i\varphi_1}, \ldots, e^{i\varphi_n})$. In each quartic invariant the phase of column i enters once via $U_{\alpha i}$ and once, conjugated, via $U_{\beta i}^*$, and likewise for column j; every φ cancels identically. ■

Standard flavour oscillations therefore cannot determine the Majorana phases and cannot distinguish a Majorana completion from a Dirac completion through propagation probabilities alone. The scope restriction matters: the phases need not cancel from hypothetical lepton-number-violating neutrino–antineutrino transition amplitudes, which lie outside the standard oscillation kernel and carry their own matching obligations.

22.4 Charge-conjugation firewall

A self-conjugate Majorana field representation does not imply that the weak interaction is invariant under charge conjugation. Field self-conjugacy, lepton-number violation, C symmetry and CP symmetry are four distinct statements with four distinct verification obligations.

---

23. Observable Bridge Firewalls

23.1 Beta-decay endpoint

In the standard light-neutrino interpretation,

$m_\beta^2 = \Sigma_i |U_{ei}|^2 m_i^2$.

This is an observable bridge — a weighted moment of the spectrum — not a definition of the mass operator, and not any single mass eigenvalue (falsifier F32).

23.2 Neutrinoless double beta decay

Under the standard light-Majorana exchange mechanism,

$m_{\beta\beta} = | \Sigma_i U_{ei}^2 m_i |$.

A null or positive result cannot be interpreted without nuclear-matrix-element, operator and mechanism matching. Other lepton-number-violating mechanisms may contribute, interfere or dominate (falsifier F33).

23.3 Cosmological mass sum

A cosmological bound or inference on $\Sigma_i m_i$ is model dependent and does not by itself determine the microscopic completion blocks (falsifier F34).

23.4 Heavy-sector observables

Active–sterile mixing, heavy decays and nonunitarity constraints probe combinations of M_D and M_R but, by Theorem 17, do not uniquely invert the ultraviolet factorisation.

23.5 Matching firewall

No experimental number may be compared with the tree-level VERSF operator until the scale, scheme, running, thresholds, interactions and relevant observable kernel are declared. Comparison without declared matching is scheme laundering (falsifier F35).

---

24. Basis Covariance

24.1 General field transformations

Under an invertible neutral coordinate transformation

$\mathcal{N}\_L \mapsto S^{-1} \mathcal{N}\_L$,

the kinetic and completion forms transform as

$K\_\nu \mapsto S^\dagger K\_\nu S$, $\mathscr{B}\_\nu \mapsto S^T \mathscr{B}\_\nu S$.

24.2 Theorem 21 — Neutral Basis Covariance

The generalised Takagi singular values obtained from the pair ($\mathscr{B}\_\nu$, K_ν) are invariant under the simultaneous transformations above.

Proof. Direct substitution into the canonical sandwich: M_ν^c transforms by unitary congruence, $M\_\nu^c \mapsto V^T M\_\nu^c V$ with $V = K\_\nu^{(1/2)} S \tilde{K}\_\nu^{(-1/2)}$ unitary ($\tilde{K}\_\nu$ the transformed metric), and Takagi singular values are unitary-congruence invariants. ■

After canonical normalisation, unitary congruence

$M\_\nu^c \mapsto U^T M\_\nu^c U$

preserves the Takagi singular values.

24.3 Typed-basis firewall

A transformation mixing active and sterile carriers may be used only after electroweak breaking and only if the corresponding gauge and interaction operators are transformed simultaneously. A mass-only basis change cannot erase the physical active/sterile distinction in the charged current (falsifier F25).

---

25. Perturbative Stability and Seesaw Conditioning

25.1 Takagi singular-value stability

For a symmetric perturbation $E = E^T$, Weyl's inequality for singular values gives

$|\tilde{m}\_i - m\_i| \le \|E\|_2$

under consistent ordering. Takagi masses are therefore Lipschitz-stable in the operator norm of the perturbation — the spectrum is robust even where the eigenframe is not.

25.2 Seesaw variation

For $M_{eff} = M_L - M_D M_R^{-1} M_D^T$, the exact first-order variation is

$$\delta M_{eff} = \delta M_L - \delta M_D M_R^{-1} M_D^T - M_D M_R^{-1} \delta M_D^T + M_D M_R^{-1} \delta M_R M_R^{-1} M_D^T,$$

using $\delta(M_R^{-1}) = -M_R^{-1} \delta M_R M_R^{-1}$. Therefore

$$\|\delta M_{eff}\| \lesssim \|\delta M_L\| + 2\|M_D\| \|M_R^{-1}\| \|\delta M_D\| + \|M_D\|^2 \|M_R^{-1}\|^2 \|\delta M_R\|.$$

25.3 Conditioning firewall

The heavy-block sensitivity is quadratic in $\|M_R^{-1}\|$: a nearly singular M_R makes the effective light operator violently sensitive to small heavy-sector perturbations. Any numerical prediction must report the condition number or the smallest singular value of the heavy block (falsifier F17).

25.4 Degeneracy stability

Pseudo-Dirac splittings and nearly degenerate light masses are especially sensitive to perturbations. Branch labels inside an unresolved degenerate subspace are not stable without another commuting observable (falsifier F37).

---

26. Numerical and Structural Non-Insertion

26.1 Forbidden target set

Define

$$\mathcal{T}_\nu = \{ m_i, \Delta m^2_{ij}, \text{mass ordering, mixing angles, CP phases, absolute mass bounds} \}.$$

A forward VERSF derivation requires

$$\mathcal{T}_\nu \cap \text{Inputs}(M_L, M_D, M_R, K_\nu, \Phi_L, \Phi_R, \Phi_{\kappa\nu}, Y_\nu, \hat{L}_\nu, \hat{L}_\nu^{(k)}, \hat{R}_\nu) = \emptyset,$$

except for quantities explicitly labelled as calibration and quarantined as such in the closure certificate.

26.2 Theorem 22 — Oscillation-Calibration Limitation

If the model uses the measured Δm^2 values to set its common neutral scale, the absolute mass scale is not predicted. If it uses the mixing matrix to choose Y_ν or the Takagi frame, the mixing is not predicted. Calibration consumes exactly the predictive content of whatever it touches.

26.3 Discrete insertion

Choosing:

- three sterile carriers because three light masses are observed;
- a zero eigenvalue because current data permit one;
- normal or inverted ordering because it fits a preferred texture;
- a rank, topology or closure count because it reproduces an angle;

is target insertion unless independently derived. Discrete choices are insertions exactly as continuous ones are; the audit does not distinguish cardinality.

26.4 Charged-lepton ladder firewall

The charged-lepton localisation exponent $\kappa = 8/3$ (a localisation exponent, unrelated to the commitment-field symbol κ of Section 17.4), the tau suppression factor and the charged magnitude operator may not be imported into the neutral sector. A neutral-sector localisation or admissibility law must be independently derived and may have different direction, exponent, rank or saturation behaviour (falsifier F26).

26.5 Blind protocol

1. Derive the active and sterile carriers.
2. Derive the gauge-admissible block pattern.
3. Derive K_ν and canonically normalise.
4. Derive the dimensionless completion operators.
5. Derive every dimensionful block scale.
6. Publish the running Takagi spectrum and relative frame.
7. Derive radiative and observable matching.
8. Only then compare with oscillation, beta-decay, lepton-number-violation and cosmological data.

Each step is frozen before the next begins; comparison precedes nothing.

---

27. Premise-Dependency Map

27.1 Named premises

Premise	Content
N1	Active and sterile carriers independently derived
N2	Charge-conjugate carriers correctly typed
N3	Positive kinetic metric at a common scale and scheme, with typed kinetic compatibility — declared as N3(a), N3(b) and N3(c) in Section 5.2
N4	Gauge-admissible complex-symmetric completion bilinear
N5	Independent origin and scale for every nonzero mass block
N6	No target-mass, splitting, ordering, angle or phase insertion
N7	Observable matching before empirical claims
N8	Controlled light–heavy hierarchy for seesaw interpretation
N9	Charged-lepton mass frame independently derived
N10	Active projection and interaction operators transformed consistently
N11	Weak-commitment kernel independently derived and typed as a Hermitian frame object, not a mass operator
N12	B−L gauge status — exact, spontaneously broken, or absent — independently declared
N13	Particle and antiparticle CAR sectors independent until a Majorana identification is derived
P1	Conventional or VERSF Planck scale admitted and correctly typed
P2	Dimensionless Planck descent factors independently derived
P3	κ scale admitted and typed as a scalar commitment scale, with the fermionic coupling bridge owed

27.2 Dependency table

● denotes an algebraic dependency. ○ denotes a derivational or predictive-grade dependency.

Result	N1	N2	N3	N4	N5	N6	N7	N8	N9	N10	P1	P2
Symmetric completion theorem		●		●								
Canonical Nambu normalisation			●	●								
Generalised Takagi masses			●	●		○						
Common-scale no-go			●	●								
Oscillation-scale no-go			●	●						●		
Dirac pair-degeneracy	●	●	●	●		○						
Lepton-number criterion		●	●	●								
Completion classification	●	●	●	●		○						
Gauge-admissibility theorem	●	●		●								

Result	N1	N2	N3	N4	N5	N6	N7	N8	N9	N10	P1	P2
Charge-conjugation completion theorem	●	●	●	●	●	○						
Exact Schur complement				●								
Hierarchical seesaw interpretation			●	●	●	○		●				
Rank/nullity theorem	●			●								
Planck normal form				●	●	○					●	●
Relative-frame mixing	●		●	●		○			●	●		
Observable masses and rates	●	●	●	●	●	○	●	as required	●	●	as used	as used
Empirical confirmation	●	●	●	●	●	○	●	as required	●	●	as used	as used

Extension premises. The Frame–Takagi compatibility theorem (Theorem 19) depends algebraically on N3, N4 and N11 and predicatively on N6 and N9. The Sterile Majorana B−L Firewall and its rigidity corollary (Section 12.4) depend on N12 through Theorem 11. Theorems 2a and 2b depend on N2 and N13. The κ-closure normal form (Section 17.4) depends on P3 and, predicatively, on N6; the hybrid κ–Planck form (Section 18.2a) depends additionally on P1 and P2 through its Planck-anchored terms. None of N11–N13 or P3 is used by any other result in this paper.

---

28. Falsification Conditions

#	Failure condition	Consequence
F1	A sterile carrier is inserted because a seesaw model needs one	Carrier circularity
F2	N_R and N_R^c are counted as independent species	Nambu doubling error
F3	Charge conjugation is claimed to generate a mass scale	Type/scale error
F4	Neutrality is claimed to prove Majorana character	Completion overclaim
F5	The neutral completion matrix is not symmetric	Fermionic-statistics violation
F6	The raw bilinear is diagonalised before kinetic normalisation	Field-normalisation error
F7	Ordinary eigenvalues of a complex symmetric matrix are called physical masses	Wrong spectral invariant
F8	Takagi singular values are not used	Neutral mass extraction invalid
F9	A bare active Majorana block is inserted before electroweak breaking	Gauge-invariance violation
F10	A Dirac block is written without an independently admitted right-handed carrier	Hidden carrier insertion

#	Failure condition	Consequence
F11	A sterile Majorana block ignores an exact VERSF charge	Gauge/charge mismatch
F12	A Majorana field condition is described as exact C symmetry of the theory	Symmetry conflation
F13	Nonzero Majorana blocks are said to preserve standard lepton number	Violates Theorem 9
F14	Near-degenerate modes are called Dirac without a protecting continuous charge	Pseudo-Dirac/Dirac conflation
F15	The Schur complement is called the exact physical light-pole spectrum	Effective-action overclaim
F16	Seesaw hierarchy is assumed without controlling $\|M_D M_R^{-1}\|$	Uncontrolled expansion
F17	A nearly singular heavy block is used without conditioning analysis	Numerical instability
F18	Fewer sterile carriers yield full-rank type-I light masses	Rank contradiction
F19	Low-energy masses are claimed to determine M_D and M_R uniquely	Violates factorisation degeneracy
F20	An inverse seesaw reconstruction is presented as a forward derivation	Direction-of-inference error
F21	Oscillation splittings are claimed to fix the lightest mass	Violates Theorem 7
F22	Oscillations are claimed to determine Majorana phases	Observable mismatch
F23	Oscillations alone are claimed to establish Majorana character	Observable overclaim
F24	The neutrino mixing matrix is defined without the charged-lepton frame	Relative-frame error
F25	Active/sterile mixing is removed by a mass-only basis choice	Interaction-basis error
F26	The charged-lepton $\kappa = 8/3$ ladder is exported to neutrinos	Sector-domain violation
F27	Planck mass is inserted as M_R without a descent theorem	Hidden scale assumption
F28	$\Phi_R = 1$ or $Y_\nu = I$ is chosen from numerical convenience	Target insertion
F29	The observed ordering is used to label VERSF branches	Retrospective labelling
F30	Mixing angles or CP phases are used to choose the completion texture	Texture fitting
F31	A massless neutrino is chosen because data allow it	Discrete target insertion
F32	A beta-decay quantity is identified directly with one mass eigenvalue	Observable-kernel error
F33	$m_{\beta\beta}$ is interpreted without a mechanism firewall	Lepton-number-violation overclaim
F34	A cosmological mass sum is treated as model independent	Cosmological matching error
F35	Running masses are compared with observables without matching	Scheme/scale mismatch

#	Failure condition	Consequence
F36	Radiative corrections are invoked qualitatively to excuse a discrepancy	Uncalculated-residual laundering
F37	A degenerate Takagi subspace is given unique branch labels without another observable	Degeneracy ambiguity
F38	A VERSF certification scale is equated with the Newton Planck scale without matching	Planck-identification laundering
F39	Reduced and unreduced Planck conventions are mixed	Planck-convention error
F40	A numerical agreement obtained after using Δm^2, angles or masses as inputs is called a prediction	Prediction-grade inflation
F41	A Hermitian frame kernel is used as, or silently converted into, the complex-symmetric mass operator	Hermitian/symmetric type conflation
F42	A bare M_R is asserted while an exact unbroken gauged B−L is in force	B−L charge-violation error
F43	Frame-kernel invariants are inserted into the mass blocks M_L, M_D, M_R	Mixing-to-mass insertion
F44	Frame–Takagi commutation or the monotone bridge is assumed rather than derived	Bridge laundering
F45	A single-CAR-family field expansion is presented as the general Dirac construction	Premature Majorana identification
F46	The Dirac equation or the mode-spinor relation $v = C \bar{u}^T$ is claimed to force identification of the independent b and d CAR sectors	Completion/equation-of-motion conflation
F47	Two mass symbols (e.g. m_κ and a neutrino mass) are equated on dimensional grounds	Mass-symbol conflation
F48	The κ-field Lagrangian-uniqueness or nonlinear-exclusion claim is cited as a closed theorem	Uninherited-claim laundering
F49	$\Phi_{\kappa\nu}$ or $\hat{L}_\nu^{(k)}$ is selected from the observed neutrino scale or oscillation data	Target insertion

---

29. NMCC-1 Closure Certificate

NMCC-1 requirement	Result
Active carrier typed	Closed structurally
Sterile carrier separated from active census	Closed as requirement
Charge-conjugate carrier typed	Closed
Nambu non-doubling theorem	Closed
Symmetric completion bilinear proved	Closed
Corrected general Dirac expansion installed	Closed

NMCC-1 requirement	Result
Field-level charge-conjugation trichotomy (Theorem 2a)	Closed
CAR-sector completion classification (Theorem 2b)	Closed
Positive kinetic metric installed	Closed as premise
Canonical Nambu normalisation proved	Closed
Symmetry preserved under canonical normalisation	Closed
Takagi mass theorem proved, with uniqueness clause	Closed
Generalised eigenvalue form proved	Closed
Neutral common-scale no-go proved	Closed
Oscillation absolute-scale no-go proved	Closed
Dirac pair-degeneracy theorem proved, with zero-mode count	Closed
Continuous lepton-number criterion proved	Closed
Majorana-block violation of standard lepton number proved, with typed kinetic condition	Closed
Neutral completion taxonomy (nested, per invariant subspace)	Closed
Gauge-admissibility theorem	Closed structurally
Sterile admissibility as minimal anomaly-neutral singlet within the stated ansatz	Inherited conditionally
Sterile Majorana B−L firewall and rigidity corollary	Closed
Charge-conjugation completion theorem	Closed conditionally
Exact Schur-complement identity, with congruence factorisation	Closed
Schur-complement/physical-spectrum firewall	Closed
Hierarchical seesaw interpretation	Closed conditionally on N8
Type-I rank and nullity constraints	Closed
Algebraic GL(n_R, ℂ) seesaw invariance and pre-quotient dimension	Closed
Existence of continuous ultraviolet non-uniqueness	Closed
Relative-frame mixing theorem	Closed
Hermitian-frame/symmetric-mass typing firewall	Closed
Frame–Takagi compatibility criterion	Closed conditionally
Reconstruction-degeneracy corollary	Closed
Weak-commitment kernel as candidate mixing frame	Inherited conditionally — mixing-only
Majorana-phase invisibility in standard lepton-number-conserving oscillations	Closed
Basis covariance	Closed
Perturbative stability apparatus	Closed
Planck–closure neutral normal form	Closed as conditional normal form

NMCC-1 requirement	Result
κ-closure direct-completion normal form	Closed as conditional normal form
Hybrid κ–Planck two-anchor, three-contribution normal form	Closed conditionally
Mass-symbol firewall	Closed
Active right-handed sterile carrier existence	Open
Sterile carrier rank	Open
Physical block pattern (M_L, M_D, M_R)	Open
Independent derivation of K_ν	Open
Independent derivation of Y_ν	Open
Independent derivation of $\hat{L}_\nu$, $\hat{R}_\nu$	Open
Independent derivation of Φ_L, Φ_R	Open
Planck-certification/Newton-scale matching	Open
Light–heavy hierarchy and heavy spectrum	Open
Absolute light-neutrino scale	Open
Light-neutrino ordering	Open
Charged-lepton-relative mixing	Open numerically
B−L gauge status (exact, broken or absent)	Open
Physical quotient dimension of the factorisation degeneracy for a specified VERSF sterile completion	Open
Frame-to-Takagi bridge (commutation or monotone map)	Open — load-bearing
κ-to-neutrino coupling bridge $\Phi_{\kappa\nu}$	Open — load-bearing
Relative dominance among the κ-direct, Planck-direct and Planck-seesaw contributions	Open
VERSF Fock-space conjugation operator constructed from substrate	Open
Dirac and Majorana CP phases	Open
Observable and radiative matching	Open
Empirical confirmation	Open

Closure grade

NMCC-1 closes the algebraic, field-typing, gauge-admissibility and canonical-spectrum framework required for a VERSF neutral-fermion mass derivation. It does not derive the physical sterile census, completion mechanism, absolute scale, mass ordering or observable neutrino spectrum.

Specifically, the closed framework comprises: the neutral-carrier typing and charge-conjugation non-doubling theorems, the field-level conjugation trichotomy with CAR-sector completion classification, the symmetric-completion and canonical Nambu-normalisation theorems, the generalised Takagi mass theorem with uniqueness clause, the common-scale and oscillation-scale no-go theorems, the Dirac-pairing theorem with exact zero-mode counting, the continuous lepton-number criterion, the neutral-completion taxonomy (nested, per invariant subspace), the gauge-admissibility theorem with the sterile Majorana B−L firewall and rigidity corollary, the Schur-complement theorem with congruence factorisation and corrected seesaw error orders, the rank/nullity theorem, the ultraviolet factorisation-degeneracy theorem with pre-quotient degeneracy count, the relative-frame mixing theorem, the repaired weak-commitment Frame–Takagi compatibility criterion with its reconstruction-degeneracy corollary, and the κ-closure and hybrid κ–Planck completion normal forms. The paper delivers the exact structural bridge from neutral occupancy to a physical Takagi mass problem while leaving all numerical neutrino masses, ordering, mixing phases, heavy scales and empirical claims open.**

---

30. Conclusion

A neutral particle is not automatically a Majorana particle, and charge conjugation is not a mass-generation mechanism. Neutrality enlarges the admissible completion space; it does not choose a completion within it.

NMCC-1 constructs the correct typed carrier for that problem. The active left-handed neutrino carrier and any independently existing sterile right-handed carrier are represented on one left-handed Nambu space by using the sterile charge-conjugate carrier. This is not a doubling of physical species. Fermionic statistics then force the neutral completion bilinear to be complex symmetric.

With a positive kinetic metric, canonical normalisation gives

$$M_\nu^c = K_\nu^{(-T/2)}\ \mathscr{B}_\nu\ K_\nu^{(-1/2)},$$

and the physical tree-level masses are its Takagi singular values,

$$U_\nu^T M_\nu^c U_\nu = \mathrm{diag}(m_i).$$

The block operator

$$\mathcal{B}_\nu = \begin{pmatrix} M_L & M_D \\ M_D^T & M_R \end{pmatrix}$$

then carries the complete Dirac–Majorana classification. A pure off-diagonal block gives pairwise Takagi degeneracy and exact Dirac pairing when protected by continuous lepton number. Diagonal blocks give Majorana-type completion and violate the standard lepton-number assignment. Small diagonal blocks give pseudo-Dirac pairs. A hierarchical heavy block gives the seesaw regime.

For invertible M_R,

$$M_{\nu,\mathrm{eff}} = M_L - M_D M_R^{-1} M_D^T$$

is the exact algebraic Schur complement — a congruence identity, not a unitary-congruence identity. Its interpretation as the physical light operator requires controlled active–sterile mixing and effective-theory matching. The type-I rank theorem shows immediately that fewer sterile carriers than active carriers force massless tree-level modes.

Two no-go results sharply limit what data can establish. Rescaling the complete mass operator leaves all neutral mass ratios unchanged. Adding a common constant to all squared masses leaves all oscillation probabilities unchanged. Oscillations therefore cannot determine the lightest mass, the absolute scale or the Majorana phases, and they cannot by themselves decide whether the physical completion is Dirac or Majorana.

The Planck basis supplies a natural scale architecture but not a numerical answer. Under a Planck-related heavy completion,

$$\mathcal{B}_\nu^{(P)} = \begin{pmatrix} (v_{cl}^2/m_P)\, \Phi_L \hat{L}_\nu & (v_{cl}/\sqrt{2})\, Y_\nu \\ (v_{cl}/\sqrt{2})\, Y_\nu^T & m_P\, \Phi_R \hat{R}_\nu \end{pmatrix},$$

so that

$$M_{\nu,\mathrm{eff}}^{(P)} = (v_{cl}^2/m_P) \cdot [\, \Phi_L \hat{L}_\nu - (1/2\Phi_R)\, Y_\nu \hat{R}_\nu^{-1} Y_\nu^T \,].$$

This identifies the exact objects VERSF must derive: the sterile census, the neutral completion texture, the Planck descent, the electroweak bridge, the relative charged-lepton frame and the observable matching. None may be selected from the intended neutrino masses or mixing data.

The weak-commitment programme now supplies a conditional candidate for one of these objects — the active mixing frame — with substrate provenance through the chain $H_{cl} \to H_W \to \mathcal{K}_\nu$. The Frame–Takagi compatibility theorem locates the exact point at which that inheritance becomes physical: the bridge from the Hermitian frame kernel to the symmetric mass operator (debt D22), together with the B−L status determination (debt D23) that partly controls the Dirac-versus-Majorana selection. A frame is a direction field, not a spectrum; NMCC-1 accepts the direction field and continues to withhold the spectrum.

With the field-level layer now installed, the full inheritance chain reads

first-order closure flow → Clifford spinor carrier → substrate CAR algebra → relativistic particle/antiparticle fields → typed charge conjugation → active/sterile neutral Nambu carrier → Dirac–Majorana completion matrix → canonical Takagi spectrum,

and the corresponding scale architecture is

κ-closure direct ⊕ Planck-suppressed direct ⊕ electroweak–Planck seesaw,

a two-anchor, three-contribution architecture whose relative weights, phases and interference are open, derivable and falsifiable quantities.

The programme advance is therefore

active neutral census → typed conjugate completion → gauge-admissible symmetric bilinear → canonical Takagi spectrum → Dirac/Majorana/seesaw classification → κ-closure and Planck–closure scale targets.

The natural next gate after NMCC-1 is not another classification paper. It is a derivation paper aimed at the first genuinely open physical objects, for example:

The Planck–Closure Seesaw and Neutrino Absolute-Scale Theorem in VERSF

or, if the sterile census is not yet established:

The Right-Handed Neutral Carrier and Sterile-Occupancy Theorem in VERSF.

Final Theorems

Neutral completion theorem. Given independently derived active and sterile carriers, correctly typed charge conjugation, a positive kinetic metric and a gauge-admissible symmetric bilinear,

$$M_{\nu^c} = K_\nu^{(-T/2)} \begin{pmatrix} M_L & M_D \\ M_D^T & M_R \end{pmatrix} K_\nu^{(-1/2)}$$

has a unique unordered nonnegative Takagi mass spectrum, and its Dirac, Majorana, pseudo-Dirac or mixed character is fixed by the block structure and the continuous-symmetry criterion.

Seesaw corollary. If M_R is invertible, the exact algebraic light completion is

$$M_{\nu,\mathrm{eff}} = M_L - M_D M_R^{-1} M_D^T,$$

with a physical light-spectrum interpretation only under controlled light–heavy mixing and matching.

Planck-basis corollary. If the heavy and direct-active completions are independently shown to be Planck related, then

$$M_{\nu,\mathrm{eff}}^{(P)} = (v_{cl}^2/m_P) \cdot [\, \Phi_L \hat{L}_\nu - (1/2\Phi_R)\, Y_\nu \hat{R}_\nu^{-1} Y_\nu^T \,],$$

and the entire neutrino absolute-scale problem is transferred to the independent derivation of the dimensionless operator in brackets and its radiative bridge to observables.

Appendix A — Minimal Algebraic Skeleton

Active carrier: $\mathscr{H}_{\nu L}$. Sterile carrier: $\mathscr{H}_{NR}$, if independently derived. Charge-conjugate carrier: $\mathscr{H}_{NR^c}$. Nambu carrier: $\mathcal{N}_L = \mathscr{H}_{\nu L} \oplus \mathscr{H}_{NR^c}$. Kinetic metric: $K_\nu > 0$.

Completion bilinear:

$$\mathcal{B}_\nu = \begin{pmatrix} M_L & M_D \\ M_D^T & M_R \end{pmatrix}, \qquad \mathcal{B}_\nu^T = \mathcal{B}_\nu.$$

Canonical mass operator:

$$M_\nu^c = K_\nu^{-T/2}\, \mathscr{B}_\nu\, K_\nu^{-1/2}.$$

Takagi masses:

$$U_\nu^T M_\nu^c U_\nu = \mathrm{diag}(m_i),\ m_i \geq 0.$$

Generalised mass equation:

$$\mathscr{B}_\nu^\dagger K_\nu^{-T} \mathscr{B}_\nu v_i = m_i^2 K_\nu v_i.$$

Lepton-number criterion:

$$[Q, K_\nu] = 0,\ Q^T \mathscr{B}_\nu + \mathscr{B}_\nu Q = 0.$$

Schur complement:

$$M_{\nu,\mathrm{eff}} = M_L - M_D M_R^{-1} M_D^T.$$

κ-closure scale and direct completion:

$$m_\kappa = \sqrt{4/3} \cdot \hbar/(c\xi),\ M_L^{(k)} = m_\kappa\, \Phi_{\kappa\nu}\, \hat{L}_\nu^{(k)}.$$

Hybrid two-anchor, three-contribution light operator:

$$M_{\nu,\mathrm{eff}} = m_\kappa\, \Phi_{\kappa\nu}\, \hat{L}_\nu^{(k)} + (v_{cl}^2/m_P)\, \Phi_L\, \hat{L}_\nu - (v_{cl}^2/2m_P\, \Phi_R)\, Y_\nu\, \hat{R}_\nu^{-1}\, Y_\nu^T.$$

Planck normal form:

$$\mathcal{B}_\nu^{(P)} = \begin{pmatrix} (v_{cl}^2/m_P)\, \Phi_L\, \hat{L}_\nu & (v_{cl}/\sqrt{2})\, Y_\nu \\ (v_{cl}/\sqrt{2})\, Y_\nu^T & m_P\, \Phi_R\, \hat{R}_\nu \end{pmatrix}.$$

Appendix B — Charge-Conjugation and Symmetry Proof

For left-handed Weyl fields n_i^α, the Lorentz scalar is

$$n_i n_j = \varepsilon_{\alpha\beta} n_i^\alpha n_j^\beta.$$

Exchanging $i \leftrightarrow j$ gives one minus sign from the antisymmetric spinor tensor and one from Grassmann exchange, so

$$n_i n_j = n_j n_i.$$

Therefore the contracted mass term

$$\tfrac{1}{2} \mathscr{B}_{\nu,ij} n_i n_j$$

depends only on the symmetric part $\mathscr{B}_{\nu,(ij)}$, and the antisymmetric flavour part contracts to zero identically:

$$\mathscr{B}_{\nu,[ij]} n_i n_j = \tfrac{1}{2} \mathscr{B}_{\nu,[ij]} (n_i n_j - n_j n_i) = 0.$$

This proves $\mathscr{B}_\nu^T = \mathscr{B}_\nu$ without loss of generality. ■

Appendix C — Canonical Takagi and Generalised Eigenvalue Proofs

Setup. Let $A = K_\nu^{(-1/2)}$, the positive Hermitian inverse square root, so that

$$M_\nu^c = A^T \mathscr{B}_\nu A.$$

Generalised eigenvalue form. The right singular operator of M_ν^c is

$$(M_\nu^c)^\dagger M_\nu^c = A^\dagger \mathscr{B}_\nu^\dagger (A^T)^\dagger A^T \mathscr{B}_\nu A = A \mathscr{B}_\nu^\dagger K_\nu^{(-T)} \mathscr{B}_\nu A,$$

using Hermiticity of A in the outer factors and, in the middle, $(A^T)^\dagger A^T = A^* A^T = (A A)^* = (K_\nu^{(-1)})^* = K_\nu^{(-T)}$, since A is Hermitian and $K_\nu^{(-1)}$ is Hermitian. If w_i is an eigenvector of $(M_\nu^c)^\dagger M_\nu^c$ with eigenvalue m_i^2, define $v_i = A w_i$. Multiplying the eigenvalue equation on the left by $A^{-1} = K_\nu^{(1/2)}$ and substituting gives

$$\mathscr{B}_\nu^\dagger K_\nu^{(-T)} \mathscr{B}_\nu v_i = m_i^2 K_\nu v_i.$$

Conversely, every solution of the generalised problem pulls back to a singular vector of M_{ν^c}, so the two spectra coincide with multiplicity. Since $(M_{\nu^c})^\dagger M_{\nu^c}$ is Hermitian positive semidefinite in the canonical frame, the w_i may be chosen orthonormal; the $v_i = A w_i$ are then K_ν-orthonormal: $v_i^\dagger K_\nu v_j = w_i^\dagger A K_\nu A w_j = w_i^\dagger w_j = \delta_{ij}$.

Takagi factorisation. Existence follows from the singular-value decomposition together with symmetry of M_{ν^c} (Autonne–Takagi): if $M = M^T$ has SVD $M = V\Sigma W^\dagger$, symmetry forces a unitary relation between V and W* on each singular subspace, from which a single unitary U with $U^T M U = \Sigma$ is constructed; the diagonal is the singular-value multiset and is therefore unique. ■

Appendix D — Dirac-Pairing Proof

Let $M_D = U_L \Sigma U_R^\dagger$ be a singular-value decomposition with $q = \mathrm{rank}(M_D)$ nonzero singular values $s_1, \dots, s_q$.

On the active–sterile Nambu carrier, apply the block unitary congruence built from U_L and the conjugate of U_R:

$$W = \begin{pmatrix} U_L & 0 \\ 0 & U_R^* \end{pmatrix}, \qquad W^T \begin{pmatrix} 0 & M_D \\ M_D^T & 0 \end{pmatrix} W = \begin{pmatrix} 0 & \Sigma \\ \Sigma^T & 0 \end{pmatrix}.$$

A permutation of coordinates reduces the result to

$$\bigoplus_a \begin{pmatrix} 0 & s_a \\ s_a & 0 \end{pmatrix} \oplus 0_{(n_L+n_R-2q)},$$

a direct sum over $a = 1, \dots, q$ of 2×2 symmetric blocks plus an exact zero block of the stated dimension. For each nonzero block, the unitary Takagi matrix

$$(1/\sqrt{2}) \cdot \begin{pmatrix} 1 & i \\ 1 & -i \end{pmatrix}$$

produces $\mathrm{diag}(s_a, s_a)$. Unitary congruence and permutation preserve Takagi values, so the nonzero Takagi multiset of the full pure-Dirac operator is $\{s_1, s_1, \dots, s_q, s_q\}$ and the zero-mode count is $n_L + n_R - 2q$. Each equal pair combines into one four-component Dirac field when the continuous charge of Theorem 9 is present. ■

Appendix E — Schur-Complement and Rank Proofs

Schur identity. For invertible M_R,

$$\nu^T M_L \nu + 2\nu^T M_D N + N^T M_R N = (N + M_R^{-1} M_D^T \nu)^T M_R (N + M_R^{-1} M_D^T \nu) + \nu^T (M_L - M_D M_R^{-1} M_D^T)\nu,$$

verified by expansion; the cross terms cancel exactly. In block form this is the congruence

$$\mathscr{B}_\nu = W^T (M_{\nu,\mathrm{eff}} \oplus M_R) W, \quad W = \begin{pmatrix} I & 0 ; M_R^{-1} M_D^T & I \end{pmatrix},$$

with $\det(W) = 1$ but W not unitary unless $M_D = 0$.

Rank inequality.

$$\mathrm{rank}(ABC) \le \min\{ \mathrm{rank}\, A, \mathrm{rank}\, B, \mathrm{rank}\, C \},$$

applied to $A = M_D$, $B = M_R^{-1}$, $C = M_D^T$, gives $\mathrm{rank}(M_{\nu,\mathrm{eff}}) \le \mathrm{rank}(M_D)$ in the type-I case. Rank–nullity on the light space then forces at least $n_L - n_R$ exact zero modes when $n_R < n_L$.

Determinant identity. For square invertible M_D and M_R, multiplicativity of determinants gives

$$\det(-M_D M_R^{-1} M_D^T) = (-1)^n \cdot \det(M_D)^2 / \det(M_R). \ \blacksquare$$

Appendix F — Worked Structural Examples

F.1 One-generation Dirac completion

$$\mathcal{B} = \begin{pmatrix} 0 & m_D \\ m_D & 0 \end{pmatrix}$$

has Takagi masses ($|m_D|$, $|m_D|$) and exact standard lepton number: $Q_L = \mathrm{diag}(1, -1)$ satisfies both conditions of Theorem 9.

F.2 One-generation type-I seesaw

$$\mathcal{B} = \begin{pmatrix} 0 & m_D \\ m_D & M_R \end{pmatrix}, \qquad |m_D| \ll |M_R|,$$

has a light algebraic completion $-m_D^2/M_R$, a heavy mode of order $|M_R|$ and active–sterile mixing of order $|m_D/M_R|$. The light Takagi mass is $|m_D^2/M_R| \cdot (1 + O(m_D^2/M_R^2))$, exhibiting the Theorem 14 error order explicitly.

F.3 Pseudo-Dirac pair

$$\mathcal{B} = \begin{pmatrix} \varepsilon_L & m_D \\ m_D & \varepsilon_R \end{pmatrix}, \qquad |\varepsilon_L|,\ |\varepsilon_R| \ll |m_D|,$$

produces a nearly degenerate pair split by the small lepton-number-violating blocks; for real positive parameters the Takagi masses are $m_D \pm (\varepsilon_L + \varepsilon_R)/2 + O(\varepsilon^2/m_D)$, matching Section 11.2. With complex blocks the splitting is $|\varepsilon_L^* + \varepsilon_R|$, which can vanish at leading order even with both blocks nonzero — the converse trap noted there.

F.4 Rank-deficient type I

With three active carriers and two sterile carriers, M_D is 3×2. The type-I light operator has rank at most two, and at least one exactly massless tree-level mode is forced — an illustration of the Corollary to Theorem 15, not a prediction that nature realises it.

F.5 Same light operator, different heavy completions

Any invertible X transforms (M_D, M_R) as in Theorem 17 without changing the effective light operator. For $n_R = 3$ this is a pre-quotient 18-real-parameter algebraic family per low-energy operator; the physically inequivalent residue after convention-fixing remains a continuous family. The low-energy mass matrix therefore cannot identify a unique ultraviolet completion.

Appendix G — Open Bridge-Debt Ledger

Debt	Required content	Status
D1 — Active-carrier debt	Pin the three active neutrino branches to the occupancy map and interaction frame	Inherited conditionally
D2 — Sterile-census debt	Derive whether right-handed neutral carriers exist and determine their number	Open — load-bearing
D3 — Charge-conjugation typing debt	Derive the exact VERSF conjugation map and all additional substrate charges	Open beyond standard representation typing
D4 — Kinetic-metric debt	Derive K_ν at the matching scale	Open
D5 — Dirac-block debt	Derive Y_ν and the active–sterile interface	Open
D6 — Active-Majorana debt	Derive or exclude M_L through a gauge-admissible VERSF completion	Open

Debt	Required content	Status
D7 — Sterile-Majorana debt	Derive or exclude M_R and its scale	Open
D8 — Lepton-number debt	Determine whether an exact or approximate continuous charge survives	Open physically; criterion closed
D9 — Planck-identification debt	Match the VERSF certification scale to the Newton Planck scale	Open
D10 — Planck-descent debt	Derive Φ_L and Φ_R	Open — load-bearing
D11 — Neutral-texture debt	Derive $\hat{L}_\nu$, $\hat{R}_\nu$ and their ranks	Open
D12 — Seesaw-hierarchy debt	Derive $\lVert M_D M_R^{-1} \rVert \ll 1$, if claimed	Open
D13 — Active–sterile mixing debt	Derive the full Takagi frame and active projection	Open
D14 — Charged-lepton-relative frame debt	Combine the neutral frame with the independently derived charged-lepton frame	Open
D15 — Ordering debt	Derive the light mass ordering without target insertion	Open
D16 — Absolute-scale debt	Derive the lightest mass and full running spectrum	Open
D17 — CP debt	Derive Dirac and Majorana phases from VERSF structure	Open
D18 — RG and threshold debt	Run all blocks and match heavy thresholds	Open
D19 — Observable-kernel debt	Derive beta, lepton-number-violation and cosmological matching kernels	Open
D20 — Blind-comparison debt	Freeze the derivation before comparison	Open procedural requirement
D21 — Empirical debt	Compare the matched theory with data	Open
D22 — Frame-bridge debt	Derive $[\mathcal{K}_\nu, H_{\nu,m}] = 0$ or the monotone map $H_{\nu,m} = m_{\nu,0}^2 f(\mathcal{K}_\nu)$ from VERSF structure	Open — load-bearing
D23 — B−L status debt	Determine whether B−L is exact gauged, spontaneously broken, or absent in the completed theory	Open
D24 — κ-coupling bridge debt	Derive $\Phi_{\kappa\nu}$ and the scalar-to-neutral-fermion completion map from commitment dynamics	Open — load-bearing
D25 — Contribution-dominance debt	Derive the relative weights, phases and interference among the κ-direct, Planck-direct and	Open

Debt	Required content	Status
	Planck-seesaw contributions from the dimensionless factors	
D26 — Fock-conjugation debt	Construct the VERSF Fock-space charge-conjugation operator from the substrate CAR algebra	Open

---

Appendix H — References

Internal VERSF programme papers

1. K. Taylor, **The Generation and Species Census Theorem in VERSF**, SC-1. [Active lepton-generation carrier and species census.]
2. K. Taylor, **The World-to-Framework Occupancy Map for Standard Model Matter in VERSF**, SC-2. [Active neutrino occupancy addresses and interaction typing.]
3. K. Taylor, **The Minimal Hermitian Lift and Frame-Rotation Firewall in VERSF**, SF-1. [Singular-value, kinetic-metric and frame-covariance discipline.]
4. K. Taylor, **The Closure-Interface Potential Necessity and Higgs-Radial Theorem in VERSF**, HR-1. [Gauge-admissible radial completion interface.]
5. K. Taylor, **The Charged-Lepton Absolute Scale Theorem in VERSF**, CLAS-1. [Common-scale no-go, canonical chiral mass structure, Planck-basis scale architecture and sector-domain firewall.]
6. K. Taylor, **The Electroweak Representation and Connection Closure in VERSF**. [Doublet typing of the active carrier, unique anomaly-neutral ν_R admissibility, B−L anomaly-admissibility.]
7. K. Taylor, **The Electroweak Flavour-Frame Operator in VERSF**. [Relative-frame mixing, weak-commitment structure, Majorana caveat.]
8. K. Taylor, **The Weak-Commitment Neutrino Operator in VERSF**. [Scale–shape separation; Hermitian frame kernel versus symmetric mass operator.]
9. K. Taylor, **The Projected Weak-Doublet Closure Hamiltonian and Projection Theorem in VERSF**. [Descent of the frame kernel from H_cl.]
10. K. Taylor, **The Weak-Commitment Closure Kernel in VERSF**. [Candidate active-frame texture and invariants.]
11. K. Taylor, **The Projected Leakage Support Trace in VERSF**. [Support-trace derivation of the leakage denominator; mixing-only scope.]
12. K. Taylor, **The First Projection Audit in VERSF**. [Provenance and audit architecture.]
13. K. Taylor, **From Substrate CAR Algebra to Relativistic Fermion Fields in VERSF**. [Field-level CAR algebra, frequency sectors, reserved antiparticle programme completed here.]
14. K. Taylor, **From Schrödinger to Dirac in VERSF**. [First-order relativistic spinor carrier and orientation-coupling mass bilinear.]
15. K. Taylor, **On the Equivalence of the κ-Field and the Commitment Field in VERSF**. [Scalar commitment mode and candidate scale m_κ; uniqueness claims not inherited.]

16. Relevant VERSF Planck-certification, Route-M and Fold–Planck–Closure papers. [Planck/closure basis typing and independence firewalls.]

Standard references

17. M. E. Peskin and D. V. Schroeder, *An Introduction to Quantum Field Theory*, Westview Press, 1995.
18. S. Weinberg, *The Quantum Theory of Fields, Volume II*, Cambridge University Press, 1996.
19. S. Weinberg, *Baryon- and Lepton-Nonconserving Processes*, Phys. Rev. Lett. 43, 1566 (1979). [Dimension-five neutrino mass operator.]
20. R. N. Mohapatra and P. B. Pal, *Massive Neutrinos in Physics and Astrophysics*, World Scientific.
21. S. M. Bilenky, *Introduction to the Physics of Massive and Mixed Neutrinos*, Springer.
22. W. Grimus and L. Lavoura, *The Seesaw Mechanism at Arbitrary Order*, JHEP 0011, 042 (2000). [Systematic arbitrary-order seesaw block diagonalisation.]
23. J. A. Casas and A. Ibarra, *Oscillating Neutrinos and* $\mu \to e, \gamma$, Nucl. Phys. B 618, 171 (2001). [Parametrisation of residual ultraviolet freedom.]
24. H. E. Haber, *A Tale of Three Diagonalizations*, arXiv:2009.03990. [Takagi versus eigenvalue and singular-value diagonalisation; phase conventions.]
25. T. Takagi, original factorisation theorem; L. Autonne, prior treatment; modern matrix-analysis treatments of unitary congruence for complex symmetric matrices.
26. R. Bhatia, *Matrix Analysis*, Springer, 1997.
27. T. Kato, *Perturbation Theory for Linear Operators*, Springer, 1976.
28. Particle Data Group, *Neutrino Masses, Mixing, and Oscillations* (review). [Majorana-phase counting; oscillation conventions.]