

The Microscopic Closure-Hessian and Full Yukawa Evaluation Theorem in VERSF

▲ Programme Milestone — Microscopic Flavour Evaluation and Fermion-Magnitude Completion Series

Gate MCHY-1 / Hessian-Generator Type Separation, Exact Relaxed Mixed-Variation Yukawa Vertex, Canonical Polar Embeddings, First Blind Baseline Branch, and Numerical Non-Closure Certificate

Keith Taylor VERSF Theoretical Physics Programme Tier A conditional microscopic-evaluation theorem — July 2026

Principal Claim

The inverse test of this paper delivers three results that could each have gone the other way, and none did. The half-lazy quark fold turned out to be an exact algebraic identity of QCMP-1's frozen constants — a previously unstated 2:1 ratio that nothing in the test's construction forces. The natural six-dimensional binary carrier is falsified by the data themselves, by a factor of thirty-seven. And the strictly positive nine-dimensional realisation sits within a factor 1.39 of the conditioning bound the physical hierarchy makes unavoidable. Results at that level of precision that survive when they could have failed are real evidence that the conditional stack is pointing at something.

But consistency is corroboration, not derivation. The claim becomes a derivation only when the substrate returns the ingredients itself — a checklist of exactly three named steps (§38.6): B_{col} from the pre-confinement closure and record response, V_q from the blind projected Hessian, and the CY3' certificate through the dressed functional. And this paper proves (§38.7) that none of the three can be discharged from the corpus as it stands: their targets are uniquely fixed, their values provably non-identifiable until the constitutive maps, response metrics, microscopic readout and chiral fermion bilinear are added to the action. The keyhole is completely specified; the metal for the key is not yet in the theory.

General Reader Summary

The question

Every particle of matter has a mass, and physics cannot explain a single one of them. The Standard Model — the most successful theory ever tested — describes the masses and blending patterns of matter with tables of numbers called Yukawa matrices, one per family:

- **Y_u** — the up-type quarks (up, charm, top); the up quark is one ingredient of protons and neutrons.
- **Y_d** — the down-type quarks (down, strange, bottom); the down quark is the other.
- **Y_e** — the charged leptons: the electron and its heavier siblings, the muon and tau.
- **Y_ν** — the neutrinos, the ghostly partners of the charged leptons. (The original Standard Model leaves neutrinos massless; this fourth table joins once the theory is extended for their observed masses.)

Each table is a 3×3 grid: the diagonal sets how heavy each of the three generations is, the off-diagonal sets how generations blend in interactions. Every physically independent number in them is measured and copied in by hand. Why the masses span a factor of a million, why the blending looks the way it does, why matter and antimatter differ subtly — the theory is silent.

The VERSF proposal

VERSF proposes these numbers are readings of one underlying quantity: the *stiffness* of the closure substrate. Think of a bell — its notes are not chosen, they are consequences of its shape and stiffness. If VERSF is right, the four tables can be calculated rather than measured. An earlier paper proved that, if so, the tables must have one specific mathematical shape. This paper begins the calculation and stress-tests that shape from both ends.

What the paper does

It lays the exact groundwork. It separates two objects that had shared one name (a stiffness is always positive; a steering object averages to zero; nothing nonzero is both), computes the "spring" part of the stiffness exactly — discovering that flavour structure *cannot* come from it — derives the exact recipe for the final answer including how the substrate re-settles when the Higgs is disturbed, and closes every loophole through which the measured answers could be smuggled in: attachments are fixed by choice-free constructions, and every input must be frozen before any comparison with experiment.

It runs the first blind calculation. On the simplest explicit substrate configuration, the full pipeline runs end to end with nothing borrowed from experiment. Four tables come out full-rank but wrong in a revealing way: down quarks identical to electrons, up quarks to neutrinos, all with one common spectrum and no matter–antimatter asymmetry. Because nothing was fitted, this failure is information — it says precisely which missing ingredient must do which job.

It runs a dress rehearsal with real anchors. Borrowing exactly two measured constants (the strength of electromagnetism and the Higgs scale) — still no particle masses — the programme's inherited formulas produce a complete trial mass table for all nine charged particles, from an electron at 0.52 MeV to a top-quark trial value near 292 GeV, plus quark- and neutrino-mixing

tables. These are pre-conversion numbers resting on assumptions the source papers list as unproven: a rehearsal with borrowed props, but one where the whole cast shows up.

It tests the architecture in reverse — and then re-expresses it. Could *any* single universal object, read through fixed particle addresses, produce those trial tables? Three sharp answers. One universal rule for all nine charged values exists and works exactly, with a genuinely new fact hidden inside it: one ratio comes out as exactly $\frac{1}{2}$ — an algebraic identity that nothing in the construction forces, not a fit. The simplest six-slot home for that rule is proved *impossible*: the top quark and the tau cannot coexist in it, by a factor of thirty-seven. And a nine-slot home — separating the lepton channel from the two quark channels, as the theory's own particle census always suggested — always fits, for any conceivable data; its value lies in the failures it survives, not the fit it guarantees. The object was then re-expressed entirely in the theory's own named constants, each coefficient identified with a specific frozen quantity from earlier papers. A natural but not-yet-proved assumption about how colour is realised in the substrate separates its lepton and quark parts — and under that assumption, six is the minimum possible number of readout channels.

What is and is not claimed

The reverse question — could this architecture work at all — is now closed, conditionally, with a yes. The forward question — does the theory's own microscopic calculation produce it — remains open, and the paper proves *why*, exactly: with the theory's equations as currently written, the three missing pieces are mathematically underdetermined — many different microscopic stories would produce the same ingredients — so no amount of further calculation can pin them down until specific new terms are added to the theory's action. One piece of the mixing story did come out of pure symmetry. For the rest: the keyhole is now completely specified, but the metal for the key is not yet in the theory.

"the answer must have this shape" → "the shape has an explicit working model, built forwards from named parts — what remains is to make the substrate return those parts."

Headline Results and Conditionality

Claim	MCHY-1 status
The Yukawa mediator is the positive physical closure Hessian	Closed as typing
The positive Hessian is identical to a traceless transport generator	Disproved except for the zero operator
Exact radial potential Hessian $H_V = 2\lambda_{\text{sub}} \rho_{\text{c}^2} P_{\text{rad}}$	Closed
Universal residual-Hessian formula $H_X = 2J_X^\dagger W_X J_X$	Closed

Claim	MCHY-1 status
Explicit bare positive Hessian H_{bare}	Closed as formula; Jacobians not yet numerically supplied
Isotype-resolved normal form $H_{\text{cl}} = rP_{\text{rad}} + \beta_A P_A + \beta_L P_L + \beta_{\perp} P_{\perp}$	Inherited; exact conditional on multiplicity-freeness (else matrix coefficients on multiplicity spaces)
Real-second-variation convention declared; factor-of-two ambiguity removed	Closed
Full trace, angular and auxiliary fluctuation carrier	Constructed as required parent space
Exact auxiliary relaxation by Schur complement	Closed
Exact relaxed mixed-variation definition of Y_f	Closed
CY3' becomes a testable equality of two independently calculable operators	Closed
Rank-one total scalar projector can produce rank-three Yukawa matrix	Disproved
Rank-three total projector can produce four distinct full-rank spectra under isometric embeddings	Disproved (common-spectrum no-go)
Common-spectrum no-go establishes rank $P_Y > 3$ (not a rank-six minimum)	Closed
Six-dimensional generation-by-completion pre-readout address carrier $P_{\Omega}^{(3 \times 2)}$	Corpus-supported as completion subcarrier; conditionally excluded as the complete charged carrier (§38); succeeded by $P_{\Omega}^{(3 \times 3)}$, rank $P_Y^{\text{phys}} \leq 9$
Exact address selectors S_{fX} with $S_{fX}^\dagger S_{fX} = I_3$ (N_R inherited separately)	Inherited as exact
One universal address-realisation bridge J_{Ω}	Named reduction premise (UAR), not yet derived
Minimal flat covariant transport carries the flavour answer alone	Disproved (diamond degeneracy)
Symmetric-diamond 32-dimensional Hessian defined as first executable stage	Closed as protocol
First blind baseline branch $H_{\text{cl}}^{(0)}$, $P_Y^{(0)}$, $\mathcal{E}_{fL/R}^{(0)}$ frozen and evaluated	Closed

Claim	MCHY-1 status
Four baseline matrices $Y_{f^{(0)}}$ returned without flavour input	Closed
Baseline reproduces the Standard Model structural class	Disproved ($Y_{d^{(0)}} = Y_{e^{(0)}}$, $Y_{u^{(0)}} = Y_{\nu^{(0)}}$, one common spectrum, single real $1 \leftrightarrow 3$ rotation, no CP)
Conditional charged-sector representatives Y_e^C , Y_d^C , Y_u^C under a declared Hermitian convention	Assembled and quarantined (§37)
Conditional relative-frame candidates V_{CKM}^C , U_{PMNS}^C	Assembled and quarantined; z_{C3} and the assumed lepton-frame premises remain underived
Matching-point audit branch MCHY-1C(α_0): full charged matrices, mass table, rank hierarchy certificates	Assembled; consumes two measured non-flavour anchors ($\alpha(0)$, $v_\star$), zero flavour observables
Neutral texture as explicit linear forms in (m_1, m_2, m_3)	Returned; masses, Majorana phases, M_D , M_R , M_L^{irr} remain open
Inverse universal-closure consistency test	First execution complete (§38)
One universal nonfactorising address rule carries all nine charged singular values	Passes exactly; lepton baseline $\equiv$ CLAS operator, quark fold $\equiv$ exact half-lazy profile
Binary rank-six positive compression of the charged operators	Disproved analytically (Proposition 3; positivity violated by factor 37.37)
Complex C_3 trine rank-six compression	Excluded for the declared representatives ($\lambda_{\min}^{\max} = -0.533$, all six assignments, numerically converged across three methods; the pole-pinned mixed variant is a different problem at -0.401 , §38.3)
Frame-invariant binary no-go	Proved (Proposition 3'): violation ≥ 37.20 from singular values alone — no frame, sign or representative choice rescues it
Three-address carrier $C^3 \otimes C^3$ with fixed doublet-locked selectors	Realised analytically (Proposition 4) — universally, for any targets under the declared convention; architectural, not evidential (vacuity note, §38.5)
Rank-six positive universal witness inside the nine-dimensional carrier	Constructed exactly (Proposition 5); upgraded to global minimum under colour superselection (Proposition 7)
Colour-Sector Equivariance (Premise CSE) forces lepton–quark block separation	Named premise about the microscopic realisation, with falsifier — not channel-level representation theory ($\mathbf{3} \otimes \mathbf{3} \supset \mathbf{1}$)

Claim	MCHY-1 status
Diagonal magnitude kernel identified exactly with named corpus coefficients (Proposition 6)	Closed: coefficients named, not derived — the constants themselves remain underived
S_{cl}, T_{cl} derived without consuming Y_u or Y_d as primitive inputs	Closed conditionally (consumes $A_\ell, B_{col}, D_\chi, V_q$)
Complete positive-completion classification $R = T^\dagger S^{-1} T + Z, Z \succcurlyeq 0$	Closed (Proposition 8): no unidentified freedom remains
Charged inverse universal-closure calculation	Closed conditionally under the inherited corpus premises
Forward discharge of the three remaining steps	Attempted (§38.7): targets uniquely fixed, values provably non-identifiable from current inputs
C_3 support democracy of the quark frame	Conditionally discharged from functional C_3 invariance (Proposition 10)
B_{col} from J_C, J_R ; blind V_q ; $CY3'$ residue	Each proved non-identifiable without new constitutive information (Propositions 9, 11, 12)
Completeness fork	If the present action is complete, $Y_{f^{mixed}} = 0$ and the Yukawa route fails; the required chiral bilinear is absent
Conditioning of the strictly positive representative	$\kappa \approx 7.79 \times 10^5$, within $1.39\times$ the unavoidable top-to-electron hierarchy bound
Third inverse-test round (derive S, T from the operator algebra)	Conditionally closed on MCHY-1C: S_{cl}, T_{cl} are exact functions of $A_\ell, B_{col}, D_\chi, V_q$; the positive universal-operator family is classified by one residual $Z \succcurlyeq 0$; colour-equivariant minimum rank is six
Jarlskog sign of V_{CKM}^C under the standard quartet	Returns -3.00×10^{-5} ; orientation must be frozen before comparison
Numerical 6×6 neutral block	Not returnable — B_N^{frozen} typed only; $Y_v, R_N, K_N, M_\star$ open
$CY3'$ residues $\Delta_{f^{CY3'}}$	Unevaluated — not zero
Interface conjugation alone yields distinct sector spectra on the baseline branch	Disproved (baseline isospectrality, Proposition 1)
Completion-interface representations for u, d, e, ν	Inherited
Three neutral right carriers	Inherited conditionally from Paper 2
Occupancy addresses determine numerical embedding amplitudes by themselves	Disproved
Canonical polar construction of embeddings from address maps	Closed

Claim	MCHY-1 status
Embeddings are fixed by a unique admissible intertwiner, if one exists	Closed as theorem
Full-rank certificate $\sigma_{\min}(Y_f) > 0$	Closed
Four <i>physical</i> numerical Yukawa matrices presently available	No — only the structurally insufficient baseline branch
Numerical Hessian Jacobians J_C , $D_\star$, J_R presently available	No
Numerical scalar-profile projector presently available	No
Numerical generation embeddings presently available	No
Blind publication protocol	Closed
Current numerical closure test	Not yet passed
Empirical flavour confirmation	Not claimed

Abstract

The Closure-Operator Yukawa Construction Theorem reduces the Standard Model flavour problem to the evaluation of a common scalar-readout closure Hessian between independently fixed chiral carrier embeddings,

$$Y_{f_c} = \mathcal{E}_{fL}^\dagger P_Y K_{cl}^{(-1/2)} \mathcal{H}_{cl} K_{cl}^{(-1/2)} P_Y \mathcal{E}_{fR}.$$

Its principal open debt is microscopic: derive the fermion-dressed substrate functional, evaluate the physical closure Hessian, derive the scalar-compatible readout projector and calculate the chiral embeddings without flavour data. This paper establishes the complete microscopic evaluation theorem and executes its first exactly calculable components.

First, it corrects a load-bearing type ambiguity. The stable closure Hessian is a positive self-adjoint operator on the physical fluctuation carrier, whereas the refinement transport generator is an algebra-valued generator. A nonzero Hermitian operator cannot be both positive and traceless. The positive Yukawa mediator is therefore denoted $\mathcal{H}_{cl}$, while T_{cl} denotes the transport generator; they may be dynamically related but may not be identified as one typed object.

Second, the paper fixes the second-variation convention (real physical second variation, giving radial curvature $m_{\text{rad}}^2 = 2\lambda_{\text{sub}} \rho_c^2$) and derives the exact potential Hessian $H_V = 2\lambda_{\text{sub}} \rho_c^2 P_{\text{rad}}$ together with the universal residual-Hessian formula: at any admissible zero of $A_X = \|R_X\|^2 \{W_X\}$, the contribution is $H_X = 2J_X^\dagger W_X J_X \geq 0$. The complete bare positive closure Hessian is

$$H_{\text{bare}} = 2\lambda_{\text{sub}} \rho_{\text{c}}^2 P_{\text{rad}} + 2J_{\text{C}}^\dagger W_{\text{C}} J_{\text{C}} + 2J_{\text{T}}^\dagger W_{\text{T}} J_{\text{T}} + 2J_{\text{R}}^\dagger W_{\text{R}} J_{\text{R}}.$$

Third, on the symmetric four-event diamond background $\rho_{\star, \text{e}} = \rho_{\text{c}} n_{\text{e}}$, the minimal covariant graph-incidence transport residual with trivial holonomy is evaluated exactly: $H_{\text{T}} \sim 2L_{\diamond} \otimes I_8$ with $\text{spec } L_{\diamond} = \{0, 2, 2, 4\}$. The result is positive and explicit, and it is blind to the $\mathbb{C}^8$ internal address: it supplies no intrinsic generation hierarchy, sector splitting, complex phase or flavour orientation of its own. Nontrivial flavour information must reside in J_{C} , J_{R} , refinement-dependent transport weights, the radial/profile sector, the embeddings, or the fermion-dressed mixed blocks.

Fourth, the full physical fluctuation carrier is decomposed as $\mathbb{T}_{\star} = \mathcal{H}_{\text{rad}} \oplus \mathcal{H}_{\text{prof}} \oplus \mathcal{H}_{\text{aux}}$, and the physical reduced Hessian is the positive Schur complement $\mathcal{H}_{\text{red}} = H_{\text{cc}} - H_{\text{c}\eta} H_{\eta\eta}^{-1} H_{\eta\text{c}}$.

Fifth, the paper proves the Scalar-Field/Readout-Rank Separation Theorem, $\text{rank } Y_{\text{f}} \leq \text{rank } P_{\text{Y}}$, and strengthens it with the Common-Spectrum No-Go: $\text{rank } P_{\text{Y}} = 3$ with isometric canonical embeddings forces all sectors to share one singular spectrum, so four distinct full-rank spectra require $\text{rank } P_{\text{Y}} > 3$. The no-go establishes only that lower bound; the minimal generation-by-completion *factorised* candidate is $\mathcal{H}_{\text{Y}}^{(3 \times 2)} \simeq \mathbb{C}_{\text{generation}}^3 \otimes \mathbb{C}_{\text{completion}}^2$ of rank six, selected only if the readout carrier is independently proved to factorise.

Sixth, the exact relaxed mixed-vertex formula is derived. With $\mathcal{D}_{\text{Q}} = \partial_{\text{Q}} - (C^{-1})^{\text{ab}} U_{\text{bQ}} \partial_{\eta^{\text{a}}}$,

$$Y_{\text{f}\text{c}} = K_{\text{fL}}^{(-1/2)} [\mathcal{D}_{\text{Q}} \mathcal{B}_{\text{f}}]_{\star} K_{\text{fR}}^{(-1/2)},$$

and the inherited completion–Hessian identity CY3' becomes the exact microscopic equality

$$[\mathcal{D}_{\text{Q}} \mathcal{B}_{\text{f}}]_{\star} \stackrel{?}{=} K_{\text{fL}}^{(1/2)} \mathcal{E}_{\text{fL}}^\dagger P_{\text{Y}} \hat{H}_{\text{red}} P_{\text{Y}} \mathcal{E}_{\text{fR}} K_{\text{fR}}^{(1/2)}.$$

Seventh, the chiral embeddings are characterised as derived equivariant partial isometries and given a canonical polar construction from raw address maps, $\mathcal{E}_{\text{fX}} = B_{\text{fX}} (B_{\text{fX}}^\dagger B_{\text{fX}})^{(-1/2)}$ with $B_{\text{fX}} = K_{\text{cl}}^{(1/2)} A_{\text{fX}} K_{\text{fX}}^{(-1/2)}$. Their uniqueness is a quotient statement over the unitary commutant of the readout operator.

Eighth, the paper executes the first blind baseline evaluation, branch MCHY-1(0): $H_{\text{cl}}^{(0)} = 2\kappa_{\text{T}} L_{\diamond} \otimes I_8 + 2\lambda_{\text{sub}} \rho_{\text{c}}^2 P_{\text{rad}}(\theta)$ on the diamond, the minimal rank-six factorised readout $P_{\text{Y}}^{(0)} = P_{\text{evt}} \otimes P_{12}$, and maximally symmetric interface-conjugate embeddings. The branch returns four explicit full-rank matrices with common singular spectrum $\simeq (5.8176, 3.0109, 2.9642)$ at the frozen trial values ($\kappa_{\text{T}} = 1$, $r = (32/63)^2$, $\theta = \pi/6$), the exact degeneracies $Y_{\text{d}}^{(0)} = Y_{\text{e}}^{(0)}$ and $Y_{\text{u}}^{(0)} = S Y_{\text{d}}^{(0)} S$, a single real $1 \leftrightarrow 3$ relative rotation $\approx 2.24^\circ$, and no CP phase. The baseline therefore fails the Standard Model structural class while succeeding as a blind pipeline test, and it localises the missing physics: J_{C} or J_{R} must break quark–lepton universality, refinement-dependent transport must generate hierarchy, holonomy or complex record structure must generate CP, and derived address maps must generate the $1 \leftrightarrow 2$ and $2 \leftrightarrow 3$ frames.

Ninth, a corpus-assembly audit sharpens the unknown ledger. The occupancy map supplies exact address selectors S_{fX} with $S_{fX}^\dagger S_{fX} = I_3$ (the sterile S_{NR} inherited separately), coordinatising the eight raw embeddings as $A_{fX} = J_\Omega S_{fX}$ through one address-realisation bridge — whose genuine universality is registered as the named Premise UAR, an open theorem target rather than a result. The census and projected weak-doublet gates support a six-dimensional generation-by-completion *pre-readout* address carrier, $P_\Omega^{(3 \times 2)} = P_{\text{gen}} \otimes P_{\text{comp}}$ — subsequently reduced by the inverse test to a completion subcarrier of the nine-dimensional charged address carrier $P_\Omega^{(3 \times 3)}$, with the physical readout $P_Y^{\text{phys}} = \text{Proj}(\text{im}(P_Y^{\text{micro}} J_\Omega P_\Omega^{(3 \times 3)}))$ of rank at most nine, to be calculated. The substrate gate reduces the unknown Hessian, conditional on multiplicity-freeness, to a three-isotype normal form $H_{cl} = rP_{\text{rad}} + \beta_A P_A + \beta_L P_L + \beta_\perp P_\perp$. A conditional corpus branch MCHY-1C is assembled and quarantined: the explicit operator candidate $Y_e^{\text{CLAS}} \simeq \text{diag}(2.992 \times 10^{-6}, 6.197 \times 10^{-4}, 1.0696 \times 10^{-2})$ on the frozen $\alpha_{\text{int}} = 7960$ branch; pre-RG quark coefficient ledgers $D_d^{\text{grid}}, D_u^{\text{grid}}$ and their declared positive-Hermitian representatives ($Y_u^{\text{C}} = V_{\text{CKM}}^{\text{C}\dagger} D_u^{\text{grid}} V_{\text{CKM}}^{\text{C}}$ displayed and verified); the relative-frame candidates $V_{\text{CKM}}^{\text{C}}$ ($|J_q| \simeq 3.00 \times 10^{-5}$, with a binding sign-orientation caveat) and $U_{\text{PMNS}}^{\text{C}}$ ($\theta_{12} \simeq 33.40^\circ$, $\theta_{13} \simeq 8.59^\circ$, $\theta_{23} \simeq 48.37^\circ$, $\delta_{\text{CP}} \simeq 227.24^\circ$, $J_{\text{lep}} \simeq -0.02448$); and the typed but numerically unevaluated neutral block B_N^{frozen} on the Route-A branch $v_{cl} = 243.722807638$ GeV. A second, empirically anchored audit branch MCHY-1C(α_0) is then assembled on $\alpha_{\text{int}} = \alpha(0)$ and $v_\star = 246.21965$ GeV — two measured non-flavour anchors, zero flavour observables — returning full charged matrices with certificates ($\kappa(Y_u) \simeq 1.35 \times 10^5$, all determinants nonzero), an unmatched tree-level mass table (e.g. τ at 1.865 GeV, b at 4.863 GeV, top matching-point at 291.8 GeV), and the complete neutral texture as explicit linear forms in the unknown light masses. All corpus branches are inherited under their source gates' open premises and firewalled from the blind evaluation; every CY3' residue remains unevaluated — not zero; and the inverse universal-closure consistency test receives its first constrained execution. The result is a partial pass and a sharp analytic failure: one universal nonfactorising address rule $L(q, t) = A + qB + qtD$ reproduces all nine charged singular values exactly, with A identically the CLAS operator and the role-odd fold exactly half-lazy (increments $\frac{1}{2} \ln(77/3)$ and $\frac{1}{4} \ln(77/3)$, an algebraic identity of the grid) — while the natural binary rank-six positive compression fails analytically (third-generation positivity violated by a factor of 37.37, unrepairable by the residual Hermitian freedom), the complex trine fails for the declared representatives with $\lambda_{\text{min}}^{\text{max}} = -0.533$ (numerically converged across three methods, all six assignments) — while the binary no-go is upgraded to frame invariance from singular values alone (Proposition 3', factor ≥ 37.20) — and free address directions rescue positivity only through near-collinear solutions at condition number $\simeq 6.4 \times 10^9$. The generation-by-completion binary carrier $\mathbb{C}^3 \otimes \mathbb{C}^2$ is therefore conditionally excluded. The second round then realises the three-address carrier $\mathbb{C}^3_{\text{generation}} \otimes \mathbb{C}^3_{\text{address}}$ with fixed, doublet-locked, non-collinear selectors — a construction that succeeds for any positive targets under the declared convention (vacuity note, §38.5), its content architectural: the canonical Schur family $C_Y^{(9)}(\mu)$, with $S = (Y_d + Y_u)/\sqrt{2}$ and $T = (Y_d - Y_u)/\sqrt{2}$, is positive for all $\mu \geq 0$ and recovers all three charged operators exactly, a positive rank-six universal witness is constructed exactly inside the nine-dimensional carrier (global rank minimality not claimed), and the strictly positive representative sits within a factor 1.39 of the unavoidable top-to-electron conditioning bound. The reconciled architecture is a nine-dimensional address-generation carrier housing an explicit rank-six readout witness, with the physical readout rank an open calculation bounded by nine. The third round

then closes the charged inverse calculation conditionally: the diagonal magnitude kernel admits exact corpus-coefficient identification in $\text{span}\{I, G, \Pi_3\}$ — the CLAS localisation operator, the exact half-lazy fold and a colour-centroid operator, every coefficient a named frozen constant (Proposition 6) — from which S_{cl} and T_{cl} are assembled from named conditional corpus operators — a forward reassembly within the conditional corpus algebra, consuming A_{ℓ} , B_{col} , D_{χ} and the conditional frame V_q , never Y_u or Y_d as primitive matrices, but not a derivation from the substrate. Colour-sector equivariance (Premise CSE — a named premise about the microscopic address realisation, with its own falsifier) forces lepton–quark block separation and upgrades the rank-six witness to a genuine global minimum (Proposition 7), and the complete positive-completion classification (Proposition 8) leaves no unidentified freedom. The charged inverse universal-closure calculation is thereby closed under the inherited corpus premises. Finally, the three remaining steps are attempted *forwards* from the substrate, yielding the Threefold Non-Identifiability Certificate: the corpus uniquely fixes their targets but provably not their values — infinitely many (J_C, J_R) pairs realise any colour operator, every unitary frame is generable until $P_W H_{\text{cl}} P_W$ is evaluated, and every $CY3'$ residue is attainable from the unspecified dressed bilinear — while the C_3 support democracy of the quark frame *is* conditionally discharged from functional C_3 invariance. The irreducible new inputs are exact: $U_{ij}(\rho)$, $\Phi_{ij}(\rho)$, W_C , W_R , $P_{Y^{\text{micro}}}$, and a derived chiral bilinear $\bar{\psi}_{\text{fL}} \mathcal{K}_{\text{f}}[\rho, \eta] \psi_{\text{fR}}$ in the dressed functional.

The charged outputs are Y_u, Y_d, Y_e . The neutral output is the complete symmetric block

$$B_N = \begin{pmatrix} M_L^{\text{irr}} (v_{\text{cl}} \sqrt{2}) Y_v & (v_{\text{cl}} \sqrt{2}) Y_v^T M_R \end{pmatrix},$$

with its Takagi spectrum and light Schur operator. Every singular value, rank, spectral projector, determinant phase, CKM relative frame, charged-lepton frame, PMNS-compatible frame and CP-odd invariant is required as a blind output.

The paper supplies an executable freeze, computation, conditioning and publication protocol, and defines the staged programme: explicit substrate residual maps $\rightarrow \hat{H}_{\text{cl}} \rightarrow P_Y \rightarrow \mathcal{E}_{\text{fL/R}} \rightarrow Y_{\text{f}}$, with the symmetric-diamond 32-dimensional Hessian as Stage A. It issues a numerical non-closure certificate for the current corpus: the bosonic parent functional is partly explicit, but its transport coefficients, trace-sector extension, fermion-dressed mixed variations, scalar-profile projector and numerical chiral intertwiners are not yet supplied. Accordingly, no numerical 3×3 matrix is claimed as the physical VERSF Yukawa output. Four numerical baseline matrices are nevertheless published as branch MCHY-1(0), and their failure defines the jobs of the missing microscopic terms. The remaining calculation is finite and precisely typed.

Notation and Binding Type Conventions

The substrate event carrier is

$$\mathcal{H}_{\text{sub}} = \ell^2(E) \otimes \mathbb{C}^8.$$

The full closure-normalised substrate coordinate is ρ . The gauge-invariant radial closure amplitude is ϱ . The electroweak interface is

$$H(x) = ((v_{cl} + \varrho(x))^{\sqrt{2}} U_{EW}(x) \chi_0, \chi_0 = (0, 1)^T.$$

The complete dressed field tuple is

$$\Xi = (\rho, \eta; Q_L, u_R, d_R, L_L, e_R, N_R; \text{conjugate carriers}),$$

with stable frozen background $\Xi_\star$. The positive substrate Hessian is

$$\mathbb{H}_\star = \delta^2 \Gamma_{\text{sub}}^{\text{dressed}} / \delta \Xi \delta \Xi |_{\{\Xi_\star\}}.$$

Second-variation convention (binding). All Hessians in this paper are real physical second variations of the action with respect to real fluctuation coordinates, not Wirtinger mixed derivatives. In this convention the radial curvature of the quartic well is

$$m_{\text{rad}}^2 = 2\lambda_{\text{sub}} \rho_{\text{c}}^2.$$

Any inherited text that quotes $\lambda_{\text{sub}} \rho_{\text{c}}^2$ for the same quantity is using the Wirtinger convention; the factor of two is hereby fixed before any Yukawa evaluation, and every downstream numerical statement is bound to the real convention.

Hypercharge convention (binding). This paper uses $Q = T_3 + Y$ with $Y(\Phi_{cl}) = +1/2$, i.e. $\Phi_{cl} \sim (1, 2, +1/2)$. Inherited papers that use $Q = T_3 + Y_{SM}/2$ with $Y_{SM}(\Phi) = +1$ are converted by $Y = Y_{SM}/2$ on import; no representation label is carried across without this conversion.

Real-versus-complex representation (binding). The physical second-variation operator on the diamond carrier is 64×64 and real symmetric, $H_{\mathbb{R}} \in \text{Sym}(64, \mathbb{R})$. The 32×32 complex Hermitian notation is an economical representation used only where the response is complex-linear and contains no independent anomalous $\delta\rho \delta\rho$ block. The distinction is load-bearing for the radial sector: the genuine potential Hessian acts on $\text{Re}(n^\dagger \delta\rho)$ and assigns zero stiffness to the phase partner $\text{Im}(n^\dagger \delta\rho)$, whereas the complex projector $|n\rangle\langle n|$ stiffens both. Wherever the complex representation is used, the phase directions must first be classified among the gauge, null and auxiliary modes.

The symbol $\mathcal{H}_{cl}$ denotes the positive physical closure-stiffness operator obtained from the bosonic block after gauge, null and auxiliary directions are treated correctly. The symbol T_{cl} denotes the closure-transport generator. They are not identified.

The canonical closure Hessian is

$$\hat{H}_{cl} = K_{cl}^{(-1/2)} \mathcal{H}_{cl} K_{cl}^{(-1/2)}.$$

Where an un-hatted Latin H_{cl} appears — in the isotype normal form, the baseline branch $H_{cl}^{(0)}$, and inherited projected-block expressions such as $P_W H_{cl} P_W$ — it denotes the bare

positive closure Hessian in the representation of its context, prior to auxiliary relaxation and canonicalisation; the script $\mathcal{H}_{\text{cl}}$ is reserved for the abstract positive stiffness operator and $\hat{H}_{\text{cl}}$ for its canonical form.

The scalar-compatible Yukawa projector satisfies $P_Y = P_Y^\dagger = P_Y^2$. The common Yukawa readout operator is

$$\hat{C}_Y = P_Y \hat{H}_{\text{cl}} P_Y.$$

The raw chiral embeddings are $\iota_{\text{fL}} : \mathcal{H}_{\text{fL}} \rightarrow \mathcal{H}_{\text{cl}}$ and $\iota_{\text{fR}} : \mathcal{H}_{\text{fR}} \rightarrow \mathcal{H}_{\text{cl}}$. The canonical embeddings are

$$\mathcal{E}_{\text{fL}} = K_{\text{cl}}^{(1/2)} \iota_{\text{fL}} K_{\text{fL}}^{(-1/2)}, \quad \mathcal{E}_{\text{fR}} = K_{\text{cl}}^{(1/2)} \iota_{\text{fR}} K_{\text{fR}}^{(-1/2)}.$$

“Full” means complete at the declared leading local, canonically normalised, single-radial-insertion, renormalisable closure-Hessian order, with all retained auxiliary modes relaxed. It does not mean all higher-dimensional, composite-channel and nonlocal corrections vanish.

Contents

Part I — The Frozen Dressed Substrate and the Explicit Positive Hessian

1. Aim and programme position
2. Inherited inputs
3. The dressed substrate functional
4. The frozen admissible background
5. Hessian–generator type separation
6. The exact radial potential Hessian
7. The universal residual-Hessian formula
8. The full physical fluctuation carrier
9. Exact auxiliary relaxation
10. Positivity and stability

Part II — Scalar Readout and the Microscopic Yukawa Vertex

11. The completion-interface direction
12. Scalar-field/readout-rank separation
13. The common-spectrum no-go and the minimal readout carrier
14. The fermion-bilinear block

- 15. The relaxed mixed-variation theorem
- 16. The microscopic CY3' matching theorem
- 17. Canonical units and vertex normalisation
- 18. Composite and higher-order channels

Part III — Chiral Embeddings and Carrier Determinacy

- 19. The frozen matter-carrier ledger
- 20. Neutral-carrier inheritance
- 21. Embeddings as equivariant partial isometries
- 22. The canonical polar construction
- 23. Embedding uniqueness modulo the commutant
- 24. Shared-doublet consistency
- 25. Transversality and full-rank certification
- 26. Carrier-separation firewalls

Part IV — Four-Sector Evaluation

- 27. Charged-sector operators
- 28. Neutral symmetric completion
- 29. Singular spectra and physical projectors
- 30. CKM relative frame
- 31. Charged-lepton, neutrino and PMNS-compatible frames
- 32. Determinant and phase invariants
- 33. Texture, rank and hierarchy audit

Part V — The First Executable Evaluation and Blind Protocol

- 34. The symmetric-diamond target
- 35. The flat-transport degeneracy theorem
- 36. The first blind baseline evaluation — branch MCHY-1(0)
- 37. The corpus-assembly branch — MCHY-1C
- 38. The inverse universal-closure consistency test — first execution
- 39. The staged executable programme
- 40. Frozen input manifest
- 41. Evaluation algorithm
- 42. Conditioning and numerical thresholds
- 43. Publication package
- 44. Non-insertion firewall
- 45. Current corpus evaluation
- 46. Falsification conditions
- 47. MCHY-1 theorem
- 48. Closure certificate

Part I — The Frozen Dressed Substrate and the Explicit Positive Hessian

1. Aim and programme position

The previous Yukawa gate proved that the leading fermion completion operators, if VERSF is correct, must be compressions of one scalar-compatible closure Hessian. It deliberately did not evaluate the Hessian, projector or embeddings.

The present paper has five aims:

1. to identify the correctly typed parent microscopic object;
2. to derive every component of the positive closure Hessian that the frozen functional already determines exactly, and to evaluate the first explicit background;
3. to derive the exact mixed variation that defines each Yukawa matrix;
4. to specify how the projector and embeddings are calculated without flavour data, and to prove the structural no-gos that constrain them;
5. to define the complete staged calculation and publish an honest closure verdict.

The correct order of execution is fixed once and for all:

explicit substrate residual maps $\rightarrow \hat{H}_{cl} \rightarrow P_Y \rightarrow \mathcal{E}_{fL/R} \rightarrow Y_f$.

No stage may be skipped, and no later stage may be back-fitted into an earlier one.

This paper consumes Paper 1's scalar-interface carrier and Paper 2's neutral-carrier ledger. Paper 1's numerical v_{cl} is not used in constructing any dimensionless Yukawa matrix. This separation is binding because the scalar paper's scale-locking route used a pre-existing electron-interface relation. The independently evaluated Y_e must therefore be compared with that relation only after its matrix is frozen, never used to define it.

2. Inherited inputs

2.1 Substrate functional

The strongest explicit parent functional presently available is

$$F[\rho] = \Sigma_e (\lambda_{sub}/4)(\rho_e^\dagger \rho_e - \rho_c^2)^2 + A_C[\rho] + A_T[\rho] + A_R[\rho],$$

written on a discrete event set, or with $\int d\mu$ on the coarse-grained carrier. The three admissibility terms represent closure, transport and record consistency. The admissible manifold is

$$\mathcal{M}_{\text{adm}} = \{ \rho : A = 0 \}, A = A_C + A_T + A_R,$$

and the gradient-flow selector is

$$\Pi_{\text{adm}}^{\text{GF}} \Psi = \lim_{\sigma \rightarrow \infty} \rho_{\sigma} \Psi(\sigma)$$

under the analytic, coercive and basin-uniqueness conditions of the substrate theorem.

2.2 Electroweak completion interface

The chiral electroweak skeleton requires the minimal interface $\Phi_{\text{cl}} \sim (\mathbf{1}, \mathbf{2}, +\frac{1}{2})$ and its conjugate $\tilde{\Phi}_{\text{cl}} = i\sigma_2 \Phi_{\text{cl}}^* \sim (\mathbf{1}, \mathbf{2}, -\frac{1}{2})$. The sector attachments are

$$\bar{L}_L \Phi_{\text{cl}} e_R, \bar{Q}_L \Phi_{\text{cl}} d_R, \bar{Q}_L \tilde{\Phi}_{\text{cl}} u_R, \bar{L}_L \tilde{\Phi}_{\text{cl}} N_R,$$

the last when the neutral right carrier exists.

2.3 Scalar carrier

The scalar gate supplies a unit interface embedding $E_H^\dagger E_H = 1$ and the exact physical-kernel reduction rule. It also makes clear that the radial amplitude and the traceless deformation algebra are distinct sectors.

2.4 Matter addresses

The occupancy programme fixes one VERSF address for each matter slot — generation key, colour support, weak role, chirality, hypercharge and completion route. It does not by itself supply numerical embedding amplitudes.

2.5 Fermionic carrier

The CAR reconstruction supplies the substrate Fock carrier and the free relativistic fermion fields under its stated coarse-graining and mode-function premises. The interaction and species-specific mixed substrate coupling are explicitly downstream debts.

2.6 Neutral branch

Paper 2 freezes three active neutral carriers and three terminal neutral singlet partners,

$$\mathcal{H}_{\text{vL}} \simeq \mathbb{C}^3, \mathcal{H}_{\text{NR}} \simeq \mathbb{C}^3,$$

with the selected type-I-dominant branch conditional on its microscopic NC-2 and NC-3 returns.

2.7 Explicit background

The substrate-dynamics paper supplies one explicit nonuniform admissible background: the symmetric four-event diamond,

$$\mathbf{n}_a = \mathbf{e}_1, \mathbf{n}_b = \cos \theta \mathbf{e}_1 + \sin \theta \mathbf{e}_2, \mathbf{n}_c = \cos \theta \mathbf{e}_1 - \sin \theta \mathbf{e}_2, \mathbf{n}_d = \mathbf{e}_1, \rho_{\star,e} = \rho_c \mathbf{n}_e, e \in \{a, b, c, d\}.$$

This is the immediate Hessian calculation target of Part V.

3. The dressed substrate functional

The bosonic substrate functional alone cannot produce a fermion Yukawa vertex. The required parent object is the complete fermion-dressed effective substrate functional

$$\Gamma_{\text{sub}}^{\text{dressed}}[\rho, \eta, \bar{Q}_L, Q_L, \bar{u}_R, u_R, \bar{d}_R, d_R, \bar{L}_L, L_L, \bar{e}_R, e_R, \bar{N}_R, N_R].$$

Here η denotes every non-radial closure, transport, record, gauge-fixed auxiliary and heavy mode retained at the declared order.

The functional is not defined by appending arbitrary Yukawa matrices. It is defined as the result of coarse-graining the same substrate action that generates the bosonic admissibility and fermionic CAR sectors.

For each sector define the raw left–right bilinear block

$$\mathcal{B}_f(\rho, \eta) = \delta^2 \Gamma_{\text{sub}}^{\text{dressed}} / \delta \bar{\psi}_f \delta \psi_f \big|_{\{\psi=0\}}.$$

Its representation type is fixed before its numerical value is known. For the charged sectors,

$$\mathcal{B}_u : \mathcal{H}_{uR} \rightarrow \mathcal{H}_{uL}, \mathcal{B}_d : \mathcal{H}_{dR} \rightarrow \mathcal{H}_{dL}, \mathcal{B}_e : \mathcal{H}_{eR} \rightarrow \mathcal{H}_{eL}.$$

For the neutral Dirac block,

$$\mathcal{B}_v : \mathcal{H}_{NR} \rightarrow \mathcal{H}_{vL}.$$

The Majorana blocks are separately typed symmetric forms and are not part of $\mathcal{B}_v$.

4. The frozen admissible background

The background is determined before any flavour data are opened:

$$\Xi_{\star} = \text{stationary point of } \Gamma_{\text{sub}}^{\text{dressed}} \text{ on } \mathcal{M}_{\text{adm}},$$

with

$$\delta\Gamma/\delta\rho|_{\star} = 0, \delta I/\delta\eta|_{\star} = 0, \psi_{-\star} = 0$$

for the vacuum fermion background. A local minimum must satisfy positivity on the physical bosonic tangent carrier after gauge and null directions are removed.

A negative physical eigenvalue is not a small Yukawa coefficient. It is an instability of the chosen background.

5. Hessian–generator type separation

Theorem 1 — Positive-Hessian/Traceless-Generator No-Go

Let A be a finite-dimensional Hermitian operator. If $A \geq 0$ and $\text{Tr } A = 0$, then $A = 0$.

Proof. Every eigenvalue of A is nonnegative. Their sum is the trace. A sum of nonnegative numbers vanishes only when every number vanishes. Therefore all eigenvalues are zero and $A = 0$. ■

Consequence

A nonzero positive closure Hessian cannot be a traceless Hermitian $\mathfrak{su}(8)$ generator. The binding distinction is

$$\mathcal{H}_{\text{cl}} \in \text{Herm}^+(\mathbb{T}_{\text{cl}}^{\text{phys}}) \text{ versus } T_{\text{cl}} \in \mathfrak{su}(8)$$

in the Hermitian-generator convention. The substrate update $e^{(i\epsilon T_{\text{cl}})}$ may be constructed from a dynamical response derived from $\mathcal{H}_{\text{cl}}$, but the exponentiated generator and the positive stiffness operator are not silently identified.

In particular, any inherited text that names a nonzero traceless matrix such as L_{12} "the closure Hessian" is retyped by this theorem: such an object can serve as a transport generator or as an eigen-direction label; it cannot itself be the nonzero positive stiffness Hessian.

Representation-theoretic refinement

The phrase "Hessian restricted to $\mathfrak{su}(8)$ " must mean a positive bilinear form on traceless orbit directions,

$$\mathcal{H}_{\text{ang}} : \mathfrak{su}(8)_{\mathfrak{h}_{-\star}} \rightarrow \mathfrak{su}(8)_{\mathfrak{h}_{-\star}},$$

not an element of $\mathfrak{su}(8)$ with positive spectrum. This type correction is binding throughout MCHY-1.

6. The exact radial potential Hessian

For one event, write the potential

$$V(\rho) = (\lambda_{\text{sub}}/4)(|\rho|^2 - \rho_c^2)^2.$$

Theorem 2 — Radial Potential Hessian

At any admissible background $\rho_\star = \rho_c n$ with $|n| = 1$, the real physical second variation of V is

$$\delta^2 V = 2\lambda_{\text{sub}} \rho_c^2 (\text{Re } n^\dagger \delta \rho)^2,$$

so the exact potential contribution to the physical Hessian is

$$H_V = 2\lambda_{\text{sub}} \rho_c^2 P_{\text{rad}}, P_{\text{rad}} = \bigoplus_e |n_e\rangle_{\mathbb{R}} \langle n_e|_{\mathbb{R}}.$$

Proof. Write $|\rho|^2 = \rho_c^2 + 2\rho_c \text{Re } n^\dagger \delta \rho + |\delta \rho|^2$. Then $|\rho|^2 - \rho_c^2 = 2\rho_c \text{Re } n^\dagger \delta \rho + O(\delta \rho^2)$, and $V = (\lambda_{\text{sub}}/4)(2\rho_c \text{Re } n^\dagger \delta \rho)^2 + O(\delta \rho^3) = \lambda_{\text{sub}} \rho_c^2 (\text{Re } n^\dagger \delta \rho)^2 + O(\delta \rho^3)$. The second variation $\delta^2 V$, defined so that $V = V_\star + \frac{1}{2}\delta^2 V + O(\delta \rho^3)$, is therefore $2\lambda_{\text{sub}} \rho_c^2 (\text{Re } n^\dagger \delta \rho)^2$. ■

Consequences

1. The quartic potential supplies exactly one positive radial stiffness direction per event, with curvature $m_{\text{rad}}^2 = 2\lambda_{\text{sub}} \rho_c^2$ in the declared real convention.
2. Every direction tangent to the fixed-norm sphere receives zero stiffness from V . The generation, flavour and orientation structure of the Yukawa problem **cannot** come from the potential term. It must come from the closure, transport and record admissibility terms, or from the fermion-dressed mixed blocks.
3. The convention declaration of the Notation section removes the factor-of-two ambiguity between the real and Wirtinger conventions before any Yukawa evaluation is attempted.

7. The universal residual-Hessian formula

Write the three admissibility residuals schematically as $R_C(\rho)$, $R_T(\rho)$, $R_R(\rho)$, with

$$A_X = \|R_X\|^2_{\{W_X\}}, W_X > 0.$$

At an admissible background, $R_X(\rho_\star) = 0$. Let

$$J_X = \partial R_X / \partial \rho |_{\{\rho_\star\}}.$$

Theorem 3 — Universal Residual Hessian

The second variation of a weighted squared residual at a zero of that residual is

$$\delta^2 A_X = 2\langle J_X \delta \rho, W_X J_X \delta \rho \rangle,$$

so each admissibility contribution to the physical Hessian is

$$H_X = 2 J_X^\dagger W_X J_X \geq 0.$$

Proof. Expand $R_X(\rho_\star + \delta\rho) = J_X \delta\rho + O(\delta\rho^2)$. Then $A_X = \langle J_X \delta\rho + O(\delta\rho^2), W_X(J_X \delta\rho + O(\delta\rho^2)) \rangle = \langle J_X \delta\rho, W_X J_X \delta\rho \rangle + O(\delta\rho^3)$, since the would-be linear term vanishes with $R_X(\rho_\star) = 0$. Positivity follows from $W_X > 0$. ■

The complete bare positive Hessian

Combining Theorems 2 and 3,

$$H_{\text{bare}} = 2\lambda_{\text{sub}} \rho_c^2 P_{\text{rad}} + 2J_C^\dagger W_C J_C + 2J_T^\dagger W_T J_T + 2J_R^\dagger W_R J_R.$$

Every term is manifestly positive semi-definite; the sum is the positive parent object that Theorem 1 requires and that the construction theorem consumes. For the transport term, since J_ρ is declared linear, $R_T(\rho) = D_\star \rho$ and

$$H_T = 2D_\star^\dagger D_\star.$$

This formula is presently exact but not yet numerical: the substrate paper itself records that J_ρ 's explicit linear operator, coefficient and momentum normalisation remain unspecified. The matrices J_C , $D_\star$, J_R are the outstanding microscopic inputs; every flavour-relevant question about the bosonic Hessian is now a question about them and about nothing else.

8. The full physical fluctuation carrier

At the frozen background, decompose the retained physical fluctuation carrier as

$$T_\star = \mathcal{H}_{\text{rad}} \oplus \mathcal{H}_{\text{prof}} \oplus \mathcal{H}_{\text{aux}},$$

where $\mathcal{H}_{\text{rad}}$ carries the one-dimensional radial closure-amplitude line, $\mathcal{H}_{\text{prof}}$ contains scalar-compatible internal closure profiles capable of distinguishing generation and sector support, and $\mathcal{H}_{\text{aux}}$ contains every orthogonal mode to be relaxed.

The full bosonic Hessian has block form

$$H_{\text{bos}} = \begin{pmatrix} H_{cc} & H_{c\eta} \\ H_{\eta c} & H_{\eta\eta} \end{pmatrix}, \mathcal{H}_c = \mathcal{H}_{\text{rad}} \oplus \mathcal{H}_{\text{prof}}.$$

The carrier $\mathcal{H}_c$ must not be replaced by the radial line alone.

9. Exact auxiliary relaxation

Theorem 4 — Reduced Closure-Hessian Theorem

Assume $\mathbb{H}_{\text{bos}} \geq 0$ and $H_{\eta\eta} > 0$ on the auxiliary physical support. After the auxiliary modes are relaxed, the physical closure Hessian on $\mathcal{H}_{\text{c}}$ is

$$\mathcal{H}_{\text{red}} = H_{\text{cc}} - H_{\text{c}\eta} H_{\eta\eta}^{-1} H_{\eta\text{c}},$$

and $\mathcal{H}_{\text{red}} \geq 0$.

Proof. For fixed $x \in \mathcal{H}_{\text{c}}$, minimise $(x, y)^\dagger \mathbb{H}_{\text{bos}}(x, y)$ over $y \in \mathcal{H}_{\text{aux}}$. The stationary point is $y_\star = -H_{\eta\eta}^{-1} H_{\eta\text{c}} x$. Substitution gives $x^\dagger (H_{\text{cc}} - H_{\text{c}\eta} H_{\eta\eta}^{-1} H_{\eta\text{c}}) x$. The parent quadratic form is nonnegative, so the minimised form is nonnegative. ■

Canonical form

With positive physical kinetic metric K_{cl} ,

$$\hat{H}_{\text{red}} = K_{\text{cl}}^{(-1/2)} \mathcal{H}_{\text{red}} K_{\text{cl}}^{(-1/2)}.$$

No raw Hessian eigenvalue is interpreted physically before this canonicalisation.

10. Positivity and stability

The physical spectrum is

$$\hat{H}_{\text{red}} = \sum_{\alpha} \lambda_{\alpha} \Pi_{\alpha}, \lambda_{\alpha} \geq 0.$$

Degenerate levels are represented by projectors Π_{α} , not arbitrarily chosen eigenvectors. The following are distinct outcomes:

- $\lambda_{\alpha} < 0$: background instability;
- $\lambda_{\alpha} = 0$: exact flat, gauge or protected null direction requiring classification;
- $0 < \lambda_{\alpha} \ll 1$: soft but physical closure direction;
- large λ_{α} : stiff closure direction.

A negative mode may not be reinterpreted as a phase, texture suppression or light fermion channel.

Part II — Scalar Readout and the Microscopic Yukawa Vertex

11. The completion-interface direction

Write the electroweak interface as

$$H(x) = ((v_{cl} + q(x))^{\wedge/2}) U_{EW}(x) \chi_0, \chi_0 = (0, 1)^T.$$

The radial amplitude q is a scalar. The orientation $U_{EW} \chi_0$ carries the weak-doublet representation. The radial field derivative is therefore distinct from projection onto the full internal closure-profile carrier.

12. Scalar-field/readout-rank separation

Theorem 5 — Scalar-Field/Readout-Rank Separation

Let P_Y be the total scalar-compatible closure-readout projector and

$$Y_f = E_{fL}^\dagger P_Y A P_Y E_{fR}$$

for any operator A . Then

$$\text{rank } Y_f \leq \text{rank } P_Y.$$

Consequently, a predicted rank-three Yukawa matrix requires $\text{rank } P_Y \geq 3$.

Proof. By the product-rank inequality, $\text{rank}(E_{fL}^\dagger P_Y A P_Y E_{fR}) \leq \text{rank } P_Y$. The second claim follows immediately. ■

Corollary — Rank-One Total Projector No-Go

If $P_Y = E_H E_H^\dagger$ on the complete closure carrier, then $\text{rank } Y_f \leq 1$. That identification is therefore incompatible with any full-rank three-generation sector.

Correct structural reading

The one-dimensional radial line supplies the scalar insertion. The total readout support must additionally contain closure-profile directions:

$$\mathcal{H}_Y \subseteq \mathcal{H}_{\text{rad}} \otimes \mathcal{H}_{\text{prof}}.$$

Under a factorised interface construction,

$$P_Y = P_{\text{rad}} \otimes P_{\text{prof}}, P_{\text{rad}} = E_H E_H^\dagger,$$

with $\text{rank } P_{\text{prof}} \geq 3$ for every full-rank three-generation sector. The factorised form is a calculation target, not assumed where the microscopic carrier fails to factorise. A natural general form is

$$P_Y = \sum_{\alpha} |\chi_{\alpha} \otimes h\rangle \langle \chi_{\alpha} \otimes h|, \alpha = 1, \dots, r_Y,$$

where h is the one-dimensional radial/interface direction and the χ_{α} are independent closure or refinement profiles.

13. The common-spectrum no-go and the minimal readout carrier

The inequality $\text{rank } P_Y \geq 3$ is necessary but not sufficient. There is a stronger constraint.

Theorem 6 — Common-Spectrum No-Go

Suppose $\text{rank } P_Y = 3$ and every full-rank canonical embedding

$$\mathcal{E}_{fL/R} : \mathbb{C}^3 \rightarrow \text{im } P_Y$$

is a partial isometry, as the canonical discipline requires. Then each embedding is unitary onto $\text{im } P_Y$, and for every sector

$$Y_f = U_{fL} \dagger \hat{C}_Y U_{fR}$$

has the same singular values as $\hat{C}_Y$ restricted to $\text{im } P_Y$. Hence all four sectors possess **the same three Yukawa singular values**.

Proof. A full-rank partial isometry between spaces of equal finite dimension is unitary. Singular values are invariant under left and right multiplication by unitaries. ■

Consequence

$\text{rank } P_Y = 3$ is insufficient for four distinct full-rank spectra

under the canonical partial-isometry discipline. The admissible resolutions are exactly three:

1. **rank $P_Y > 3$** , so that different sectors can occupy genuinely different subspaces of the readout support;
2. non-isometric embeddings, which would require a derived, non-canonical kinetic contraction and must then be typed and justified as such;
3. genuinely sector-dependent higher-order operators, which belong to the composite channels of §18 and cannot be smuggled into the leading vertex.

The minimal generation-by-completion factorised candidate carrier

The common-spectrum no-go alone establishes only $\text{rank } P_Y > 3$; it does not establish that the minimum rank is six. A rank-four common carrier can already contain several inequivalent three-dimensional subspaces and is not excluded by Theorem 6.

The cleanest *factorised* candidate consistent with resolution 1 is

$$\mathcal{H}_{Y^{(3 \times 2)}} \simeq \mathbb{C}^3_{\text{generation}} \otimes \mathbb{C}^2_{\text{completion}}, \text{rank } P_Y = 6.$$

The two-dimensional completion factor can carry the independently typed left/right (equivalently forward/reverse) completion geometry, while the three-dimensional factor carries the three frozen generation profiles. Rank six is selected only if the readout carrier is independently proved to factorise into a three-level generation profile and a two-level completion orientation. This is a candidate to **derive** from the scalar-compatible eigenprofiles of $\hat{H}_{\text{red}}$; it is not a result to assume, and its rank may not be adjusted retrospectively because a spectrum looks wrong. Where the baseline branch uses it, it is a frozen factorised ansatz.

Executed-test update. The inverse universal-closure consistency test of §38, run against the conditional MCHY-1C(α_0) charged operators, *excludes* the binary $\mathbb{C}^3 \otimes \mathbb{C}^2$ carrier: the natural fold assignment fails positivity analytically (Proposition 3, violation factor 37.37), the exclusion is frame-invariant from the singular values alone (Proposition 3', factor ≥ 37.20), the complex trine fails for the declared representatives with $\lambda_{\min}^{\max} = -0.533$ (numerically converged across three methods), and free directions rescue positivity only through near-collinear, pathologically conditioned solutions. The successor — the three-address carrier $\mathbb{C}^3_{\text{generation}} \otimes \mathbb{C}^3_{\text{address}}$ — is then *realised analytically* (§38.5, Propositions 4–5): one positive universal operator on the nine-dimensional carrier reproduces all three charged operators exactly through fixed, well-separated, doublet-locked selectors, and a rank-six positive witness is constructed exactly within it. Per the binding vacuity note of §38.5, this realisation succeeds for any positive targets under the declared convention — its content is architectural, and the data-discriminating results are the exclusions. The reconciled picture is $\dim \mathcal{H}_\Omega = 9$ with an existing rank-six witness, while $\text{rank } P_{Y^{\text{phys}}}$ itself remains an open calculation bounded by nine: the rank-six readout intuition acquires an explicit realisation, but not yet a minimality or derivation proof. The binary exclusion depends only on the conditional singular spectra; the trine exclusion additionally depends on the declared representatives; a physical evaluation returning different charged spectra would reopen both.

Address-level support: the pre-readout carrier

The factorised candidate is not merely aesthetic — but its corpus support must be typed precisely. The Generation and Species Census supplies three independent generation-depth return layers, and the Projected Weak-Doublet Closure Hamiltonian independently decomposes the relevant carrier through $P_W = P_Y P_{\text{role}} P_C$, where P_C selects the three generation-depth directions, P_{role} the two weak-role readings, and P_Y is the *additional* mass/readout projector, treated by that gate as a distinct third factor. The occupancy map identifies the two completion readings H and $\tilde{H}$. An address-level projector is therefore already defined on $\mathcal{H}_\Omega$:

$$P_{\text{gen}} = \sum_{i=1}^3 |i\rangle\langle i|, P_{\text{comp}} = |H\rangle\langle H| + |\tilde{H}\rangle\langle \tilde{H}|, P_{\Omega^{(3 \times 2)}} = P_{\text{gen}} \otimes P_{\text{comp}}, \text{rank } P_{\Omega^{(3 \times 2)}} = 6.$$

This is a six-dimensional *pre-readout address carrier*, $\mathcal{H}_{\Omega^{(3 \times 2)}} \simeq \mathbb{C}^3_{\text{generation}} \otimes \mathbb{C}^2_{\text{role/completion}}$. It is not itself the physical scalar-readout projector: the source gate explicitly

places P_Y alongside P_{role} and P_C , so declaring the still-unknown readout derived from these projectors would be circular.

Carrier status after the inverse test. The executed test of §38 conditionally excludes this binary carrier as the *complete* charged-sector address carrier on the MCHY-1C targets: it remains a valid completion-orientation subcarrier and the frozen carrier of the blind baseline, but it does not separately resolve the colourless charged-lepton channel from the coloured $H/\tilde{H}$ plane. The successor address carrier is

$$\mathcal{H}_{\Omega}^{(3 \times 3)} = \mathbb{C}^3_{\text{generation}} \otimes \text{span}\{|E_c\rangle, |D\rangle, |U\rangle\}, P_{\Omega}^{(3 \times 3)} = P_{\text{gen}} \otimes P_{\text{addr}}, \text{rank } P_{\Omega}^{(3 \times 3)} = 9.$$

The physical readout is

$$\mathcal{H}_{Y^{\text{phys}}} = \text{im}(P_Y^{\text{micro}} J_{\Omega} P_{\Omega}^{(3 \times 3)}), P_Y^{\text{phys}} = \text{Proj}(\mathcal{H}_{Y^{\text{phys}}}), \text{rank } P_Y^{\text{phys}} \leq 9,$$

where P_Y^{micro} is the microscopically derived scalar-compatible projector of Stage C, and the rank must be *calculated*, not declared. The inverse construction of §38.5 proves the *existence* of a positive rank-six readout witness inside this nine-dimensional carrier, and §38.6 upgrades that witness: under colour-sector equivariance (Premise CSE), rank six is the *global minimum* over all positive universal operators compatible with the fixed completion-channel selectors (Proposition 7). What remains unproved is that the blindly derived physical projector realises this minimum — the physical rank is bounded by nine, minimised at six under CSE, and not yet calculated. The binary 3×2 object retains exactly three roles in this paper: the blind-baseline ansatz of §36, the quark $H/\tilde{H}$ completion subcarrier inside the 3×3 space, and a conditionally excluded reduction of the full charged address carrier.

14. The fermion-bilinear block

For each sector,

$$\mathcal{B}_f(q, \eta) = \delta^2 \Gamma_{\text{sub}^{\text{dressed}}} / \delta \bar{\psi}_{fL} \delta \psi_{fR} \big|_{\{\psi=0\}},$$

a map $\mathcal{B}_f : \mathcal{H}_{fR} \rightarrow \mathcal{H}_{fL}$. The canonical mass block at the vacuum is

$$M_{f_c} = K_{fL}^{(-1/2)} \mathcal{B}_f(\bar{q}, \eta_{\star}) K_{fR}^{(-1/2)}.$$

The canonical Yukawa vertex is the derivative with respect to the canonically normalised radial field. A matrix inserted directly into $\mathcal{B}_f$ is not a derivation. The functional must return it.

15. The relaxed mixed-variation theorem

Let the constant-background effective potential be $U(q, \eta)$, with $U_a(q, \eta_{\star}(q)) = 0$ defining the relaxed branch, and $C_{ab} = U_{ab}$. Implicit differentiation gives

$$d\eta_\star/dq = -(C^{-1})^{ab} U_{bq}.$$

Define

$$\mathcal{D}_q = \partial_q - (C^{-1})^{ab} U_{bq} \partial_{\eta^a}.$$

Theorem 7 — Relaxed Mixed-Yukawa Vertex

The complete declared-order canonical Yukawa operator is

$$Y_{f_c} = K_{fL}^{(-1/2)} [\mathcal{D}_q \mathcal{B}_f(q, \eta)]_\star K_{fR}^{(-1/2)},$$

equivalently

$$Y_{f_c} = K_{fL}^{(-1/2)} [\partial_q \mathcal{B}_f - (C^{-1})^{ab} U_{bq} \partial_{\eta^a} \mathcal{B}_f]_\star K_{fR}^{(-1/2)}.$$

Proof. The reduced bilinear is $\mathcal{B}_f^{\text{red}}(q) = \mathcal{B}_f(q, \eta_\star(q))$. Applying the chain rule and the implicit derivative of $\eta_\star$ yields $d\mathcal{B}_f^{\text{red}}/dq = \partial_q \mathcal{B}_f - (C^{-1})^{ab} U_{bq} \partial_{\eta^a} \mathcal{B}_f$. Canonical fermion normalisation gives the displayed result. ■

Significance

The second term is compulsory whenever orthogonal modes respond to the radial displacement. Omitting it is equivalent to holding part of the substrate artificially fixed. No isolated term in the parent action may be called the complete Yukawa vertex unless every omitted term is proved inert under $\mathcal{D}_q$.

Definitional requirements for $[\mathcal{D}_q \mathcal{B}_f]_\star$ (binding)

For the eventual comparison with Standard Model Yukawa matrices, the physical evaluation must declare, before extraction: (i) that $\Gamma_{\text{sub}}^{\text{dressed}}$ is the *renormalised IPI effective action*; (ii) the Euclidean or Lorentzian convention; (iii) the external momentum configuration at which the local vertex is extracted (the default being the zero-momentum limit after amputation); (iv) the renormalisation scheme and matching scale $\mu_\star$; (v) the wavefunction-residue normalisation entering K_{fL} , K_{fR} ; and (vi) the prescription separating nonlocal form factors from the leading local coupling. Two otherwise correct calculations can return different matrices purely through different subtraction points or residue conventions; these six declarations are therefore manifest items (§40), frozen before any output is seen.

16. The microscopic CY3' matching theorem

CMSY-1 proposes the closure compression

$$Y_{f_c} = \mathcal{E}_{fL}^\dagger P_Y \hat{H}_{\text{red}} P_Y \mathcal{E}_{fR}.$$

MCHY-1 converts this from an architectural premise into a direct microscopic test.

Theorem 8 — Completion–Hessian Matching Criterion

The CY3' identity holds for sector f if and only if

$$[\mathcal{D}_Q \mathcal{B}_f]_\star = K_{fL}^{(1/2)} \mathcal{E}_{fL}^\dagger P_Y \hat{H}_{red} P_Y \mathcal{E}_{fR} K_{fR}^{(1/2)}.$$

When this equality holds, $Y_{f_c} = \mathcal{E}_{fL}^\dagger P_Y \hat{H}_{red} P_Y \mathcal{E}_{fR}$. When it fails, the difference

$$\Delta_f^{CY3'} = K_{fL}^{(-1/2)} [\mathcal{D}_Q \mathcal{B}_f]_\star K_{fR}^{(-1/2)} - \mathcal{E}_{fL}^\dagger P_Y \hat{H}_{red} P_Y \mathcal{E}_{fR}$$

is a physical matching residue and must be published. It may not be set to zero by definition.

Universal matching norm

Define

$$\varepsilon_f^{\text{match}} = \|\Delta_f^{CY3'}\|_2 / \max \{ \|Y_f^{\text{mixed}}\|_2, \|Y_f^{\text{Hess}}\|_2, \varepsilon_{\text{num}} \}.$$

Exact leading-order closure requires $\varepsilon_f^{\text{match}} = 0$ within certified numerical error.

17. Canonical units and vertex normalisation

If the canonical substrate units make $P_Y \hat{H}_{red} P_Y$ dimensionless, no extra universal vertex coefficient appears. Otherwise the correct formula is

$$Y_{f_c} = \mathcal{N}_Y \mathcal{E}_{fL}^\dagger P_Y \hat{H}_{red} P_Y \mathcal{E}_{fR},$$

where $\mathcal{N}_Y$ is one independently calculated universal conversion factor. The following is forbidden:

$\mathcal{N}_Y$ chosen so that one fermion mass agrees with observation.

A universal coefficient fixed from an observed top, electron or down-quark Yukawa is a calibration, not a derivation.

18. Composite and higher-order channels

The leading result applies to the linear scalar-compatible profile carrier. Composite channels require an extended carrier,

$$\mathcal{H}_{eff}^{(N)} = \mathcal{H}_{cl} \oplus \mathcal{H}_{cl}^{\wedge \otimes 2} \oplus \cdots \oplus \mathcal{H}_{\nabla} \oplus \mathcal{H}_{\{\Phi\text{-dress}\}},$$

and the full declared higher-order vertex has the form

$$Y_{-f}^{\wedge(N)} = \sum_{n=1}^N \mathcal{E}_{-fL}^{\wedge(n)\dagger} \mathcal{F}_{-f}^{\wedge(n)} \mathcal{E}_{-fR}^{\wedge(n)}.$$

A sector-specific confinement operator, accessibility correction or frame-curvature operator cannot be inserted into the common leading Hessian merely because it improves one matrix. Such effects must be typed as higher-order sector dressing and evaluated separately. In particular, resolution 3 of the common-spectrum no-go lives here, and only here.

Part III — Chiral Embeddings and Carrier Determinacy

19. The frozen matter-carrier ledger

For each generation $a = 1, 2, 3$, the charged carriers are

$$Q_{-L}^{\wedge(a)}, u_{-R}^{\wedge(a)}, d_{-R}^{\wedge(a)}, L_{-L}^{\wedge(a)}, e_{-R}^{\wedge(a)}.$$

Their representation addresses are fixed before any mass ordering is known. The left up and down carriers descend from one shared doublet carrier:

$$\mathcal{H}_{-uL} \oplus \mathcal{H}_{-dL} \subset \mathcal{H}_{-QL} \otimes \mathbb{C}^2_{-W},$$

and likewise

$$\mathcal{H}_{-vL} \oplus \mathcal{H}_{-eL} \subset \mathcal{H}_{-LL} \otimes \mathbb{C}^2_{-W}.$$

The right carriers are independently typed terminal singlets.

20. Neutral-carrier inheritance

Paper 2 freezes $\dim \mathcal{H}_{-vL} = 3$ and $\dim \mathcal{H}_{-NR} = 3$. The neutral Dirac Yukawa therefore has type

$$Y_{-v} : \mathbb{C}^3_{-NR} \rightarrow \mathbb{C}^3_{-vL}.$$

No sterile dimension may be added or removed because it improves oscillation phenomenology. The selected neutral branch further requires separate evaluation of

$$M_{-R} = M_{-R}^T \text{ and } M_{-L}^{\wedge \text{irr}} = M_{-L}^{\wedge \text{irr},T}.$$

These are not entries of Y_{-v} .

21. Embeddings as equivariant partial isometries

Let G_{exact} denote the complete exact symmetry and address structure. Define the admissible embedding set

$$\mathfrak{I}_{\text{fX}} = \{ \iota : \mathcal{H}_{\text{fX}} \rightarrow \mathcal{H}_{\text{cl}} \}$$

subject to:

1. **Representation equivariance.** $\iota R_{\text{fX}}(g) = R_{\text{cl}}(g) \iota$ for all $g \in G_{\text{exact}}$.
2. **Address support.** $P_{\Omega}(\text{fX}) \iota = \iota$.
3. **Kinetic compatibility.** $\iota^\dagger K_{\text{cl}} \iota = K_{\text{fX}}$ for a partial isometry, or the corresponding contraction inequality.
4. **Chirality.** $P_L \iota_{\text{fL}} = \iota_{\text{fL}}$ and $P_R \iota_{\text{fR}} = \iota_{\text{fR}}$.
5. **Interface typing.** Φ_{cl} or $\tilde{\Phi}_{\text{cl}}$ as fixed by the sector.
6. **Refinement stability.** The image survives every admissible refinement used in the matching construction.
7. **No target-data dependence.** No mass, angle or phase enters the constraints.

The numerical embedding problem is to calculate the admissible intertwiners from the substrate operator, not to select matrices with the desired singular-value decomposition.

22. The canonical polar construction

Constraints 1–7 fix the allowed support of each embedding but leave column normalisation and internal orientation apparently open. Both are removed by the following construction.

Theorem 9 — Canonical Polar Embedding

Suppose the occupancy and representation programme supplies a raw injective address map

$$A_{\text{fX}} : \mathbb{C}^3 \rightarrow \text{im } P_Y, X = L, R,$$

and — for the equivariance of the result — that **K_{cl} and K_{fX} are positive symmetry-invariant metrics and A_{fX} intertwines the corresponding representations.** Then $B_{\text{fX}}^\dagger B_{\text{fX}}$ lies in the domain commutant, its positive inverse square root does too, and the polar factor inherits equivariance; without metric invariance, the construction still returns a canonical partial isometry, but not necessarily an admissible equivariant one in the sense of §21.

When both metrics are trivial, the unique partial isometry with the same column space and the same address orientation is the polar factor

$$\mathcal{E}_{\text{fX}} = A_{\text{fX}} (A_{\text{fX}}^\dagger A_{\text{fX}})^{(-1/2)}.$$

The fully metric-compatible construction canonicalises the domain and codomain metrics together: form the doubly-normalised map

$$\mathbf{B}_{fX} = \mathbf{K}_{cl}^{(1/2)} \mathbf{A}_{fX} \mathbf{K}_{fX}^{(-1/2)},$$

and take its polar factor,

$$\mathcal{E}_{fX} = \mathbf{B}_{fX} (\mathbf{B}_{fX}^\dagger \mathbf{B}_{fX})^{(-1/2)}.$$

When $\mathbf{K}_{fX} = \mathbf{I}$ this reduces to $\mathcal{E}_{fX} = \mathbf{K}_{cl}^{(1/2)} \mathbf{A}_{fX} (\mathbf{A}_{fX}^\dagger \mathbf{K}_{cl} \mathbf{A}_{fX})^{(-1/2)}$, and when additionally $\mathbf{K}_{cl} = \mathbf{I}$ it reduces to the trivial-metric expression.

Proof. $\mathbf{B}_{fX}$ is injective because $\mathbf{K}_{cl}^{(1/2)}$, $\mathbf{K}_{fX}^{(-1/2)}$ are invertible and $\mathbf{A}_{fX}$ is injective; hence $\mathbf{B}_{fX}^\dagger \mathbf{B}_{fX}$ is positive definite and its inverse square root exists and is unique among positive operators. The displayed map satisfies $\mathcal{E}_{fX}^\dagger \mathcal{E}_{fX} = 1$ on $\mathbb{C}^3$ and has range $\mathbf{K}_{cl}^{(1/2)} \cdot \text{range}(\mathbf{A}_{fX})$. The polar factor is the unique partial isometry $\mathbf{P}$ with $\text{range}(\mathbf{P}) = \text{range}(\mathbf{B}_{fX})$ and $\mathbf{P}^\dagger \mathbf{B}_{fX} > 0$, which fixes the residual right unitary to the identity. ■

Remark. The double normalisation matters: canonicalising only the codomain metric while leaving the fermionic domain metric $\mathbf{K}_{fX}$ untreated would silently reintroduce the very column-scale freedom the construction exists to remove.

The address-selector reduction

The occupancy programme supplies more than support constraints. The World-to-Framework Occupancy Map assigns every chiral component a unique full address

$$\Omega(f) = (i, \sigma, C, W, T_3, \chi, Y, Q, H, A, E),$$

including generation key, colour, weak role, chirality and the H-versus- $\tilde{H}$ completion route, and it proves component-address injectivity, with the charged closure routes $u_R : \tilde{H}$, $d_R : H$, $e_R : H$. Introduce the address carrier

$$\mathcal{H}_\Omega = \text{span} \{ |\Omega(f_X^i)\rangle \}.$$

Then the raw address maps are not eight unrelated matrices. Each sector's selector is the sparse incidence map

$$\mathbf{S}_{fX} |i\rangle = |\Omega(f_X^i)\rangle, f = u, d, e, v; X = L, R,$$

and component-address injectivity gives the exact result

$$\mathbf{S}_{fX}^\dagger \mathbf{S}_{fX} = \mathbf{I}_3.$$

The remaining map is the *address-to-closure bridge*

$$J_{\Omega} : \mathcal{H}_{\Omega} \rightarrow \mathcal{H}_{cl},$$

through which the raw embeddings are coordinatised as

$$A_{fX} = J_{\Omega} S_{fX},$$

and the canonical embeddings of Theorem 9 become

$$B_{fX} = K_{cl}^{(1/2)} J_{\Omega} S_{fX} K_{fX}^{(-1/2)}, \mathcal{E}_{fX} = B_{fX} (B_{fX}^{\dagger} B_{fX})^{(-1/2)}.$$

This coordinatisation is not yet a physical reduction. Any collection of eight maps can always be packaged into one map on their direct sum; that is bookkeeping compression, not dynamical unification. The genuine reduction of eight unknowns to one requires a named premise:

Premise UAR — Universal Address-Realisation. There exists one sector-independent intertwiner $J_{\Omega} : \mathcal{H}_{\Omega} \rightarrow \mathcal{H}_{cl}$ such that

$$J_{\Omega} R_{\Omega}(g) = R_{cl}(g) J_{\Omega}$$

for every exact representation transformation g , and such that its restriction to each occupied address sector supplies the corresponding raw physical embedding.

Falsifier. If the microscopic evaluation returns sector embeddings that admit no common equivariant J_{Ω} — that is, if the substrate realises the addresses through genuinely sector-dependent maps not related by the exact symmetry — Premise UAR fails, and the unknown ledger reverts to eight sector maps rather than one bridge.

Until UAR is proved, the binding statement is: the eight embeddings are coordinatised by one formally unified address-realisation map J_{Ω} ; whether this is a genuinely universal substrate operator rather than a direct-sum packaging of sector maps is an open theorem target. What *is* exact and unconditional are the selectors S_{fX} and their isometry $S_{fX}^{\dagger} S_{fX} = I_3$ — every sector distinction downstream of the bridge is carried by them, which also structurally enforces the shared-doublet constraints of §24, since the left selectors of one doublet descend from one address block.

Neutral exception. The sterile carrier N_R is not contained in the minimal occupancy census; that map explicitly excludes a sterile right-handed neutrino. Its selector S_{NR} must be inherited separately from Paper 2's neutral extension and carries that gate's conditionality.

Consequence

Arbitrary column normalisation is eliminated. The remaining embedding problem is exactly:

derive the raw address matrices A_{fX} from the frozen occupancy, chirality, weak-role, colour, hypercharge and refinement projectors.

The representation and completion-interface papers fix their allowed support; they do not yet provide their numerical coordinates. That derivation — not any normalisation convention — is the outstanding debt.

23. Embedding uniqueness modulo the commutant

Let the *unitary commutant* be

$$UComm(\hat{C}_Y) = \{ U : U^\dagger U = I, U \hat{C}_Y = \hat{C}_Y U \}.$$

Unitarity is required and not decorative: the invariance argument below uses $U^\dagger \hat{C}_Y U = \hat{C}_Y$, which holds for commuting U only when U is unitary.

Theorem 10 — Physical Embedding Determinacy

Suppose the admissible pair set $\mathfrak{Z}_{fL} \times \mathfrak{Z}_{fR}$ is one orbit under the simultaneous action

$$(\mathfrak{z}_{fL}, \mathfrak{z}_{fR}) \mapsto (U \mathfrak{z}_{fL}, U \mathfrak{z}_{fR}), U \in UComm(\hat{C}_Y).$$

Then

$$Y_{fc} = \mathcal{E}_{fL}^\dagger \hat{C}_Y \mathcal{E}_{fR}$$

is unique, although individual embedding representatives are not.

Proof. For U in the unitary commutant, $U^\dagger \hat{C}_Y U = U^\dagger U \hat{C}_Y = \hat{C}_Y$, hence

$$(U \mathcal{E}_{fL})^\dagger \hat{C}_Y (U \mathcal{E}_{fR}) = \mathcal{E}_{fL}^\dagger U^\dagger \hat{C}_Y U \mathcal{E}_{fR} = \mathcal{E}_{fL}^\dagger \hat{C}_Y \mathcal{E}_{fR}.$$

Every admissible pair lies on that orbit, so every admissible representative returns the same operator. ■

Failure condition

If the admissible set contains two inequivalent unitary-commutant orbits giving different Y_f , the substrate has not uniquely determined the flavour operator. A further upstream selection principle is required. Choosing the orbit that best fits flavour data is prohibited.

24. Shared-doublet consistency

The left quark embeddings must descend from one doublet embedding:

$$\mathfrak{z}_{uL} = P_u \mathfrak{z}_{QL}, \mathfrak{z}_{dL} = P_d \mathfrak{z}_{QL},$$

and similarly

$$\iota_v L = P_v \iota_{LL}, \iota_e L = P_e \iota_{LL}.$$

Independent arbitrary left embeddings would manufacture CKM and PMNS freedom. The relative frames must arise from the common left carrier interacting differently with the two right-terminal channels and with the interface conjugation, not from separately fitted left bases.

25. Transversality and full-rank certification

Full-rank embeddings and positive closure support do not guarantee a full-rank cross-compression. Define

$$\tau_f = \sigma_{\min}(\mathcal{E}_f L^\dagger \hat{C}_Y \mathcal{E}_f R).$$

Theorem 11 — Yukawa Transversality Certificate

For a square three-generation sector,

$$\text{rank } Y_f = 3 \Leftrightarrow \tau_f > 0.$$

Proof. A square matrix is invertible exactly when its smallest singular value is positive. ■

The paper must publish $\tau_u, \tau_d, \tau_e, \tau_v$. Nonzero individual entries and pairwise overlaps do not substitute for this certificate.

26. Carrier-separation firewalls

26.1 Vertex versus cycle carrier

The W_7 vertex Laplacian, its cycle space $H^1(W_7)$, the $\mathbb{C}^8$ closure carrier and the 63-dimensional traceless deformation algebra are distinct spaces. An eigenvalue calculated on one may not be transferred to another without an explicit chain map or induced operator.

26.2 Laplacian versus Hessian

A graph Laplacian may contribute to a closure Hessian — §35 evaluates exactly such a contribution. It is not automatically the complete substrate Hessian.

26.3 Realisation operator versus Yukawa mediator

A finite closure-state evolution operator can supply admissible modes and completion densities. It is not automatically the positive scalar-readout Hessian.

26.4 Flavour diagnostic versus input

The following are downstream diagnostics, not microscopic inputs: charged-lepton ratio laws; quark mass grids; 1/12 participation factors; CKM curvature ansätze; PMNS attachment branches; observed mass hierarchies. They may be opened only after the four operator outputs are frozen.

Part IV — Four-Sector Evaluation

27. Charged-sector operators

At the frozen matching scale $\mu_\star$,

$$Y_u = \mathcal{E}_u L^\dagger \hat{C}_Y \mathcal{E}_u R, Y_d = \mathcal{E}_d L^\dagger \hat{C}_Y \mathcal{E}_d R, Y_e = \mathcal{E}_e L^\dagger \hat{C}_Y \mathcal{E}_e R.$$

The matrices are dimensionless. The corresponding completion-boundary mass matrices are

$$M_f(\mu_\star) = (v_{cl}/2) Y_f(\mu_\star).$$

The conditional v_{cl} from Paper 1 is applied only at this downstream step.

28. Neutral symmetric completion

The Dirac block is

$$Y_\nu = \mathcal{E}_\nu L^\dagger \hat{C}_Y \mathcal{E}_\nu R, M_D = (v_{cl}/2) Y_\nu.$$

The full neutral operator is

$$\mathbf{B}_N = \begin{pmatrix} M_L^{\text{irr}} & M_D \end{pmatrix}^{**} \begin{pmatrix} M_D^T & M_R \end{pmatrix}^{**}$$

where $M_R = M_R^T$ must be obtained from the sterile symmetric completion channel and $M_L^{\text{irr}} = M_L^{\text{irr},T}$ from any irreducible direct active Majorana channel. The seesaw light operator is

$$M_{\text{light}} = M_L^{\text{irr}} - M_D M_R^{-1} M_D^T + O(\Theta^3 M_R), \Theta = M_D M_R^{-1}.$$

Paper 3 must return

$$\text{rank } M_D, \det M_R, \sigma_{\min}(M_R)/\sigma_{\max}(M_D), \rho_L = \|M_L^{\text{irr}}\|/\|M_D M_R^{-1} M_D^T\|.$$

These are the microscopic arbiters of Paper 2's conditional branch.

29. Singular spectra and physical projectors

For $f = u, d, e$,

$$U_{fL}^\dagger Y_f U_{fR} = D_f = \text{diag}(y_{f1}, y_{f2}, y_{f3}), y_{fi} \geq 0.$$

The physical left and right magnitude operators are

$$H_f^L = Y_f Y_f^\dagger, H_f^R = Y_f^\dagger Y_f,$$

published in spectral form

$$H_f^L = \sum_i y_{fi}^2 P_{fi}^L, H_f^R = \sum_i y_{fi}^2 P_{fi}^R.$$

The projectors are physical in nondegenerate sectors. Within a degenerate level, only the total level projector is physical. The neutral symmetric operator is Takagi-factorised:

$$U_N^T B_N U_N = \text{diag}(m_{N1}, \dots, m_{N6}), m_{Ni} \geq 0.$$

30. CKM relative frame

The quark relative frame is

$$V_{\text{CKM}} = U_{uL}^\dagger U_{dL}.$$

The paper must publish $|V_{ij}|$, the Jarlskog invariant

$$J_q = \text{Im}(V_{us} V_{cb} V_{ub}^* V_{cs}^*),$$

and the basis-invariant commutator

$$J_q^{\text{op}} = \text{Im Tr}([H_u^L, H_d^L]^3).$$

No CKM angle or phase is used to choose the embeddings. Any resemblance to a previously proposed transport generator is assessed only after this matrix is frozen.

31. Charged-lepton, neutrino and PMNS-compatible frames

The charged-lepton frame is U_{eL} . Let U_ν denote the active light-neutrino Takagi frame after canonical active-carrier projection and controlled seesaw reduction. The neutrino frame must be calculated from B_N or M_{light} , not supplied by the PMNS matrix.

With the canonical charged-current transporter $C_W P_A$,

$$U_{\text{PMNS}} = U_{eL}^\dagger C_W P_A U_\nu.$$

The exact convention used for C_W , active projection and any heavy–light normalisation must be frozen before diagonalisation. The full active charged-current matrix need not be exactly unitary when heavy states mix:

$$N_A N_A^\dagger = I - \Theta \Theta^\dagger + O(\Theta^4).$$

Both the full and effective active matrices must be published.

32. Determinant and phase invariants

For every charged sector publish $\det Y_f$ and $\arg \det Y_f$ in the declared canonical representative, together with rephasing-invariant cycle phases. The determinant phase of one matrix is not by itself a physical CKM phase. The physical quark CP obstruction is relative:

$$\text{Im Tr}([H_u^L, H_d^L]^3).$$

For the neutral sector publish the appropriate Majorana-sensitive invariants in addition to the Dirac-type leptonic Jarlskog invariant.

All complex structure must be traced to: (1) complex closure projectors; (2) oriented carrier embeddings; (3) interface contractions; (4) transported kinetic geometry; (5) symmetric neutral completion blocks. No independent complex Yukawa coefficient is allowed.

33. Texture, rank and hierarchy audit

For every matrix entry,

$$(Y_f)_{ij} = \sum_\alpha \lambda_\alpha \langle \mathcal{E}_f^L e_{Li}, \Pi_\alpha \mathcal{E}_f^R e_{Rj} \rangle.$$

A physical texture zero requires projector-level vanishing anchored to a derived invariant subspace. Publish the spectral invariants:

$$I_{f1} = \text{Tr } H_f^L, I_{f2} = \frac{1}{2}[(\text{Tr } H_f^L)^2 - \text{Tr}(H_f^L)^2], I_{f3} = \det H_f^L.$$

The hierarchy-transfer inequality remains

$$\sigma_{\{k+1\}}(Y_f) \leq \|\mathcal{E}_f^L\|_2 \|\mathcal{E}_f^R\|_2 \lambda_{\{k+1\}},$$

but the actual hierarchy is not closed until the evaluated λ_α , projectors and embeddings return it.

Part V — The First Executable Evaluation and Blind Protocol

34. The symmetric-diamond target

The first Hessian that the frozen corpus permits to be evaluated is the physical 64×64 real symmetric block — represented, per the binding convention of the Notation section, as a complex 32-dimensional block on

$$\ell^2(\{a, b, c, d\}) \otimes \mathbb{C}^8,$$

at the symmetric-diamond background

$$n_a = e_1, n_b = \cos \theta e_1 + \sin \theta e_2, n_c = \cos \theta e_1 - \sin \theta e_2, n_d = e_1, \rho_{\star,e} = \rho_c n_e,$$

with edges a–b, a–c, b–d, c–d. Theorem 2 already supplies the exact potential block: $2\lambda_{\text{sub}} \rho_c^2 P_{\text{rad}}$, one radial line per event. Theorem 3 supplies the exact functional form of the remaining blocks once the explicit inputs

$$U_{ij}(\rho), J_\rho = L_J \rho, \Phi_{ij}(\rho)$$

are fixed. The required outputs of Stage A are $J_C, D_\star, J_R$, and hence H_{bare} on the diamond.

The isotype-resolved normal form

The substrate paper reduces Stage A further. It identifies three relevant diamond isotypes on the internal carrier:

R_A : the L_{12} -antisymmetric directions; R_L : the (e_1, e_2) -plane reflection-even longitudinal directions; $R_\perp$: the collective e_3 – e_8 directions;

and states that each admissibility Hessian $G_X, X \in \{C, T, R\}$, acts as a scalar on each irreducible isotype:

$$G_X|_R = \alpha_X(R) I_R.$$

One condition governs the validity of this reduction. Conditional on the diamond representation being multiplicity-free at the retained order — or on the admissibility Hessian being proportional to the identity on every multiplicity space — the reduction is exact; otherwise Schur's lemma returns a *matrix* on each multiplicity space, not a single scalar coefficient, and symmetry-allowed mixing between equivalent copies survives. The auxiliary Schur reduction of Theorem 4 must also respect the same decomposition for the normal form to pass through relaxation intact.

Under that condition, the open Hessian calculation reduces to roughly nine coefficients $\alpha_{\text{C}}^{\wedge}(\mathbf{R})$, $\alpha_{\text{T}}^{\wedge}(\mathbf{R})$, $\alpha_{\text{R}}^{\wedge}(\mathbf{R})$ over the three isotypes, and the physical positive Hessian can already be written in the normal form

$$\mathbf{H}_{\text{cl}} = r \mathbf{P}_{\text{rad}} + \beta_{\text{A}} \mathbf{P}_{\text{A}} + \beta_{\text{L}} \mathbf{P}_{\text{L}} + \beta_{\text{I}} \mathbf{P}_{\text{I}}, \quad \beta_{\text{R}} = \alpha_{\text{C}}^{\wedge}(\mathbf{R}) + \alpha_{\text{T}}^{\wedge}(\mathbf{R}) + \alpha_{\text{R}}^{\wedge}(\mathbf{R}),$$

followed by the auxiliary relaxation and canonicalisation of Theorems 4–5,

$$\hat{\mathbf{H}}_{\text{cl}} = \mathbf{K}_{\text{cl}}^{(-1/2)} (\mathbf{H}_{\text{cl}} - \mathbf{H}_{\text{c}\eta} \mathbf{H}_{\eta\eta}^{(-1)} \mathbf{H}_{\eta\text{c}}) \mathbf{K}_{\text{cl}}^{(-1/2)}.$$

This does not evaluate the β_{R} , but it collapses the unknown from a general element of $\text{Sym}(64, \mathbb{R})$ to a small isotype ledger — three stiffness coefficients plus the exactly known radial term.

Binding limitation. The corpus defines the closure operator $\mathbf{C}$ and the record morphism Φ , but does not yet specify how $\mathbf{U}_{ij}$ and Φ_{ij} vary with ρ . Without that constitutive dependence, the closure and record Jacobians $\mathbf{J}_{\text{C}}$ and $\mathbf{J}_{\text{R}}$ cannot receive numerical values merely by differentiating their displayed norms. The $\alpha_{\text{X}}^{\wedge}(\mathbf{R})$ remain named unknowns, not implicit results.

35. The flat-transport degeneracy theorem

One transport implementation can already be evaluated exactly, and its result is a binding negative constraint.

Take the covariant graph-incidence residual

$$\mathbf{R}_{\text{T}}(\rho)_{ij} = \rho_j - \mathbf{U}_{ij} \rho_i$$

on the four diamond edges, so that $\mathbf{D}_{\star} = \mathbf{B}_{\text{U}}$ is the covariant incidence matrix and

$$\mathbf{H}_{\text{T}} = 2 \mathbf{B}_{\text{U}}^{\dagger} \mathbf{B}_{\text{U}}.$$

Theorem 12 — Flat-Transport Degeneracy

For the symmetric flat transport of the substrate paper's Appendix F, the connection has trivial path holonomy and is gauge-equivalent to the ordinary diamond graph Laplacian:

$$\mathbf{B}_{\text{U}}^{\dagger} \mathbf{B}_{\text{U}} \sim \mathbf{L}_{\diamond} \otimes \mathbf{I}_8.$$

The diamond graph is $\mathbf{K}_{2,2}$, with

$$\text{spec } \mathbf{L}_{\diamond} = \{0, 2, 2, 4\},$$

hence

$\text{spec } H_T = \{0, 4, 4, 8\},$

each level repeated across all eight internal closure directions.

Proof. Trivial holonomy means every U_{ij} can be gauged to the identity by vertex-local unitaries, under which $B_U^\dagger B_U$ is conjugated to $B^\dagger B \otimes I_8$ with B the ordinary incidence matrix of the diamond. $B^\dagger B$ is the graph Laplacian of $K_{2,2}$, whose spectrum $\{0, 2, 2, 4\}$ is standard: the complete bipartite graph $K_{\{m,n\}}$ has Laplacian eigenvalues 0, m ($n-1$ times), n ($m-1$ times), $m+n$. The factor 2 from Theorem 3 rescales to $\{0, 4, 4, 8\}$. ■

Consequence — the useful negative result

This minimal transport Hessian is positive and fully explicit, and it is far too degenerate. It produces:

- no intrinsic generation hierarchy — every internal direction sees the same graph spectrum;
- no intrinsic distinction between sectors — H_T is blind to the $\mathbb{C}^8$ internal address;
- no complex phase — trivial holonomy admits a real representative;
- no intrinsic flavour orientation — under common event-profile embeddings its compressions are unitarily equivalent.

Precisely: flat covariant transport supplies no internal-sector splitting, complex phase or flavour orientation of its own. Any sector distinction obtained while retaining this flat transport term — such as the nonzero relative rotation of the baseline branch in §36.6 — must already reside in the radial/profile term or in the embeddings; it is not generated by the transport Hessian itself.

Ordinary flat covariant transport therefore **cannot** carry the Yukawa answer by itself. The nontrivial flavour information must reside in at least one of:

J_C, J_R , refinement-dependent transport weights, the fermion-dressed mixed blocks.

This theorem converts a vague hope ("perhaps transport explains flavour") into a sharp exclusion, and it directs the remaining microscopic effort to exactly the terms where structure can still live.

36. The first blind baseline evaluation — branch MCHY-1(0)

The strongest honest numerical return the current corpus permits is a first blind baseline: the exactly derived radial Hessian of Theorem 2, the flat covariant transport of Theorem 12 on the symmetric diamond, the minimal rank-six factorised readout of §13 as a frozen ansatz, and the maximally symmetric interface-conjugate embeddings. Every input below is frozen before any output is inspected, in accordance with §40. No measured mass, mixing angle or phase is consumed anywhere in the branch.

This is not the physical VERSF result: J_C, J_R , the fermion-dressed mixed blocks and the numerical address maps are absent, and only the Hessian-compression side of Theorem 8 can be evaluated — the mixed-vertex side, and hence the CY3' residue, remains open. The branch is nevertheless the correct first calculational object: a fixed, fully explicit ledger against which every missing Jacobian can be tested.

36.1 The frozen baseline branch

Freeze the minimal branch

$$K_{cl} = I, H_{cl^{(0)}} = 2\kappa_T L_{\diamond} \otimes I_8 + r P_{rad}(\theta), r = 2\lambda_{sub} \rho_c^2,$$

with $\kappa_T > 0$ the transport-stiffness coefficient. The diamond Laplacian in event order (a, b, c, d) is

$$L_{\diamond} = \begin{pmatrix} 2 & -1 & -1 & 0 \\ -1 & 2 & 0 & -1 \\ -1 & 0 & 2 & -1 \\ 0 & -1 & -1 & 2 \end{pmatrix},$$

and the radial projector at the diamond background is

$$P_{rad}(\theta) = \bigoplus_{x=a,b,c,d} |n_x\rangle\langle n_x|, n_a = e_1, n_b = \cos \theta e_1 + \sin \theta e_2, n_c = \cos \theta e_1 - \sin \theta e_2, n_d = e_1.$$

This is a 32×32 complex-linear representation; its real physical second-variation form is 64×64 . Three coefficient premises are declared and hashed:

- $\kappa_T = 1$ — a unit convention on the transport sector;
- $r = (32/63)^2 \approx 0.2579995$ — inherited from the conditional HR-1 radial-curvature ratio of the scalar-completion stack; a cross-paper conditional premise, not a flavour datum;
- $\theta = \pi/6$ — a new illustrative frozen branch, *not* a derived diamond result, registered as a premise under the manifest rule of §40.

36.2 The scalar-compatible readout projector

The event eigenprofiles of $L_{\diamond}$ are

$$\chi_0 = \frac{1}{2}(1, 1, 1, 1)^T, \chi_1 = (1/\sqrt{2})(0, 1, -1, 0)^T, \chi_2 = (1/\sqrt{2})(1, 0, 0, -1)^T, \chi_3 = \frac{1}{2}(1, -1, -1, 1)^T,$$

with

$$L_{\diamond} \chi_0 = 0, L_{\diamond} \chi_1 = 2\chi_1, L_{\diamond} \chi_2 = 2\chi_2, L_{\diamond} \chi_3 = 4\chi_3.$$

The minimal factorised scalar-compatible readout retains the three nonconstant profiles and the (e_1, e_2) completion plane:

$$P_{\text{evt}} = I_4 - |\chi_0\rangle\langle\chi_0| = \frac{1}{4} \begin{pmatrix} 3 & -1 & -1 & -1 \\ -1 & 3 & -1 & -1 \\ -1 & -1 & 3 & -1 \\ -1 & -1 & -1 & 3 \end{pmatrix}, P_{12} = \text{diag}(1, 1, 0, 0, 0, 0, 0, 0),$$

$$P_{\text{Y}^{(0)}} = P_{\text{evt}} \otimes P_{12}, (P_{\text{Y}^{(0)}})^2 = P_{\text{Y}^{(0)}} = (P_{\text{Y}^{(0)}})^\dagger, \text{rank } P_{\text{Y}^{(0)}} = 6.$$

This realises the generation-by-completion factorised candidate $\mathcal{H}_{\text{Y}}^{(3 \times 2)}$ of §13 — as a frozen factorised ansatz, not yet an independently derived substrate projector, and without any claim that six is the minimal admissible rank. Its rank exceeds three, so the common-spectrum no-go of Theorem 6 does not bind it in principle; §36.5 shows how the baseline nevertheless returns a common spectrum through its embeddings.

36.3 The readout closure Hessian

Use the readout basis

$$\mathcal{B}_{\text{Y}} = (\chi_1 \otimes e_1, \chi_1 \otimes e_2, \chi_2 \otimes e_1, \chi_2 \otimes e_2, \chi_3 \otimes e_1, \chi_3 \otimes e_2),$$

and define $c = \cos \theta$, $s = \sin \theta$, $q = (\sqrt{2} r/4) \sin 2\theta$. Then $\hat{C}_{\text{Y}^{(0)}} = P_{\text{Y}^{(0)}} H_{\text{cl}^{(0)}} P_{\text{Y}^{(0)}}$ is represented on $\mathcal{B}_{\text{Y}}$ by

$$\hat{C}_{\text{Y}^{(0)}} = \begin{pmatrix} 4\kappa_{\text{T}} + r c^2 & 0 & 0 & 0 & 0 & -q \\ 0 & 0 & 0 & 4\kappa_{\text{T}} & 0 & 0 \\ 0 & 0 & 0 & 0 & 4\kappa_{\text{T}} & 0 \\ 0 & 0 & 0 & 0 & 0 & 4\kappa_{\text{T}} \\ 0 & -q & 0 & 0 & 8\kappa_{\text{T}} & + (r/2)(1+c^2) \\ 0 & 0 & 0 & 0 & 0 & 8\kappa_{\text{T}} + (r/2)s^2 \end{pmatrix} \begin{matrix} ** \\ ** \\ ** \\ ** \\ ** \\ ** \end{matrix} \begin{matrix} | & 0 & 4\kappa_{\text{T}} & + r s^2 & 0 & 0 & -q & 0 \\ | & 0 & 0 & 4\kappa_{\text{T}} & + r & 0 & 0 & 0 \\ | & 0 & 0 & 0 & 4\kappa_{\text{T}} & + r & 0 & 0 \\ | & 0 & 0 & 0 & 0 & 4\kappa_{\text{T}} & + r & 0 \\ | & -q & 0 & 0 & 0 & 0 & 8\kappa_{\text{T}} & + (r/2)(1+c^2) \\ | & 0 & 0 & 0 & 0 & 0 & 8\kappa_{\text{T}} & + (r/2)s^2 \end{matrix} \begin{matrix} ** \\ ** \\ ** \\ ** \\ ** \\ ** \end{matrix}$$

It is positive for $\kappa_{\text{T}}, r \geq 0$. Every entry has been verified by direct construction of $P_{\text{Y}^{(0)}} H_{\text{cl}^{(0)}} P_{\text{Y}^{(0)}}$ on the completion plane.

36.4 Baseline embeddings

The corpus fixes the distinction between the interface Φ_{cl} and its conjugate $\tilde{\Phi}_{\text{cl}}$ but does not yet return numerical quark and lepton address maps. The maximally symmetric no-extra-structure choice uses the completion-plane vectors

$$v_0 = (1, 0)^T, v_{\pm} = (1/\sqrt{2})(1, \pm 1)^T,$$

giving the common left embedding and the two right completion embeddings (columns orthonormal, $E_0^\dagger E_0 = E_+^\dagger E_+ = E_-^\dagger E_- = I_3$):

$$E_0 = \text{generation } g \mapsto \chi_{\text{g-profile}} \otimes e_1, E_+ = \text{generation } g \mapsto \chi_{\text{g-profile}} \otimes v_+, E_- = \text{generation } g \mapsto \chi_{\text{g-profile}} \otimes v_-,$$

explicitly, in the basis $\mathcal{B}_{\text{Y}}$,

$$E_0 = \begin{pmatrix} 1 & 0 & 0 \\ 0 & 1 & 0 \\ 0 & 0 & -1 \end{pmatrix}, E_+ = (1/\sqrt{2}) \begin{pmatrix} 1 & 0 & 0 \\ 0 & 1 & 0 \\ 0 & 0 & 1 \end{pmatrix}, E_- = (1/\sqrt{2}) \begin{pmatrix} 1 & 0 & 0 \\ 0 & 1 & 0 \\ 0 & 0 & 1 \end{pmatrix} \begin{matrix} | & 0 & 0 & 0 \\ | & 0 & 0 & 0 \\ | & 0 & 0 & 0 \\ | & 0 & 0 & 0 \\ | & 0 & 0 & 0 \\ | & 0 & 0 & 0 \end{matrix} \begin{matrix} | & 1 & 0 & 0 \\ | & 0 & 1 & 0 \\ | & 0 & 0 & 1 \\ | & 0 & 0 & 1 \\ | & 0 & 0 & 1 \\ | & 0 & 0 & 1 \end{matrix} \begin{matrix} | & 1 & 0 & 0 \\ | & 0 & 1 & 0 \\ | & 0 & 0 & 1 \\ | & 0 & 0 & 1 \\ | & 0 & 0 & 1 \\ | & 0 & 0 & 1 \end{matrix} \begin{matrix} | & 0 & 0 & 0 \\ | & 0 & 0 & 0 \\ | & 0 & 0 & 0 \\ | & 0 & 0 & 0 \\ | & 0 & 0 & 0 \\ | & 0 & 0 & 0 \end{matrix}$$

The baseline assignment tracks the only sector distinction the explicit scalar structure presently returns — Φ_{cl} versus $\tilde{\Phi}_{\text{cl}}$:

$$\mathcal{E}_{\text{uL}}^{(0)} = E_0, \mathcal{E}_{\text{uR}}^{(0)} = E_-; \mathcal{E}_{\text{dL}}^{(0)} = E_0, \mathcal{E}_{\text{dR}}^{(0)} = E_+; \mathcal{E}_{\text{eL}}^{(0)} = E_0, \mathcal{E}_{\text{eR}}^{(0)} = E_+; \mathcal{E}_{\text{vL}}^{(0)} = E_0, \mathcal{E}_{\text{NR}}^{(0)} = E_-.$$

It contains no independently calculated colour-versus-lepton address profiles; that absence is exactly what §36.5 measures.

36.5 The four returned matrices

Set $\mathcal{N}_{\text{Y}} = 1$ on the canonical-unit branch and define

$$A = (4\kappa_{\text{T}} + r \cos^2\theta)^{1/2}, B = (4\kappa_{\text{T}} + r)^{1/2}, C = (16\kappa_{\text{T}} + r(1 + \cos^2\theta))(2\sqrt{2}), D = (r/4) \sin 2\theta.$$

The direct compressions give

$$\mathbf{Y}_{\text{d}}^{(0)} = \mathbf{Y}_{\text{e}}^{(0)} = \mathcal{N}_{\text{Y}} \begin{pmatrix} A & 0 & -D \\ 0 & 0 & C \end{pmatrix}, \mathbf{Y}_{\text{u}}^{(0)} = \mathbf{Y}_{\text{v}}^{(0)} = \mathcal{N}_{\text{Y}} \begin{pmatrix} A & 0 & D \\ 0 & 0 & C \end{pmatrix} \quad \text{**} \quad \begin{pmatrix} 0 & B & 0 \\ 0 & 0 & 0 \end{pmatrix}$$

These are the first actual four-matrix return of the programme. Each baseline matrix has three nonzero singular values, so the branch passes the rank test; these are structural singular values of a frozen compression under $\mathcal{N}_{\text{Y}} = 1$ and an unevaluated matching scale — they are not physical fermion masses or physical Yukawa magnitudes. The returns further contain: a real $1 \leftrightarrow 3$ mixing channel; no $1 \leftrightarrow 2$ mixing; no $2 \leftrightarrow 3$ mixing; no CP phase; exact equality of the down and charged-lepton sectors; and exact equality of the up and Dirac-neutrino sectors. Moreover

$$\mathbf{Y}_{\text{u}}^{(0)} = \mathbf{S} \mathbf{Y}_{\text{d}}^{(0)} \mathbf{S}, \mathbf{S} = \text{diag}(1, 1, -1),$$

so all four matrices share one singular spectrum.

Proposition 1 — Baseline Signature-Conjugate Isospectrality

On the baseline branch, the up-type and down-type returns are exactly signature-conjugate:

$$\mathbf{S} \begin{pmatrix} A & 0 & -D \\ 0 & 0 & C \end{pmatrix} \mathbf{S} = \begin{pmatrix} A & 0 & D \\ 0 & 0 & C \end{pmatrix} \quad \text{**} \quad \begin{pmatrix} 0 & B & 0 \\ 0 & 0 & 0 \end{pmatrix} \quad \text{**} \quad \begin{pmatrix} -D & 0 & C \\ 0 & 0 & C \end{pmatrix}, \mathbf{S} = \text{diag}(1, 1, -1),$$

that is, $\mathbf{Y}_{\text{u}}^{(0)} = \mathbf{S} \mathbf{Y}_{\text{d}}^{(0)} \mathbf{S}$. Since left and right multiplication by unitary signature matrices preserves singular values,

$$\sigma(\mathbf{Y}_{\text{u}}^{(0)}) = \sigma(\mathbf{Y}_{\text{d}}^{(0)}).$$

Proof. Direct verification of the displayed conjugation, entry by entry: $\mathbf{S}$ flips the sign of the (1,3) and (3,1) entries and fixes all others. Signature matrices are unitary, and singular values are invariant under unitary equivalence. ■

Scope. This is a statement about the evaluated branch, not a general theorem. A general signature-isospectrality result would require stronger intertwining hypotheses relating the completion-plane involution Σ , the readout operator $\hat{C}_Y$, the left embedding and the generation-space signature maps; merely requiring Σ to commute with the left projection is not sufficient in general.

Consequence (baseline-level). On this branch, the interface conjugation $\Phi_{cl} \leftrightarrow \tilde{\Phi}_{cl}$ alone does not separate the up-type from the down-type spectrum. Distinct sector spectra therefore require address maps that differ by more than a completion-plane orientation — at least on any branch sharing the baseline's structure. This sharpens the target of Stage D beyond what Theorem 6 already required.

36.6 The numerical trial branch

At the frozen values $\kappa_T = 1$, $r = (32/63)^2 = 0.2579995$, $\theta = \pi/6$:

$$\hat{C}_Y^{(0)} \simeq \begin{pmatrix} 4.193500 & 0 & 0 & 0 & 0 & -0.078996 \\ 0 & 0 & 0 & 4.000000 & 0 & 0 \\ 0 & 0 & 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 & 0 & 0 \end{pmatrix} \begin{pmatrix} 0 & 4.064500 & 0 & 0 & -0.078996 & 0 \\ 0 & 0 & 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 & 0 & 0 \end{pmatrix} \begin{pmatrix} 0 & 4.258000 & 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 & 0 & 0 \end{pmatrix},$$

and with $\mathcal{N}_Y = 1$,

$$Y_{d^{(0)}} = Y_{e^{(0)}} \simeq \begin{pmatrix} 2.965252 & 0 & -0.055859 \\ 5.816483 & . & . \end{pmatrix}^{**} \begin{pmatrix} 0 & 3.010860 & 0 \\ 0 & 0 & 0 \\ 0 & 0 & 0 \end{pmatrix}^{**} \begin{pmatrix} -0.055859 & 0 & 0 \\ 0 & 0 & 0 \\ 0 & 0 & 0 \end{pmatrix}$$

$$Y_{u^{(0)}} = Y_{v^{(0)}} \simeq \begin{pmatrix} 2.965252 & 0 & 0.055859 \\ . & . & . \end{pmatrix}^{**} \begin{pmatrix} 0 & 3.010860 & 0 \\ 0 & 0 & 0 \\ 0 & 0 & 0 \end{pmatrix}^{**} \begin{pmatrix} 0.055859 & 0 & 5.816483 \\ 0 & 0 & 0 \\ 0 & 0 & 0 \end{pmatrix}$$

The common singular spectrum is

$$\sigma(Y_{f^{(0)}}) \simeq (5.817577, 3.010860, 2.964158) \text{ for every sector.}$$

The relative left frame is a real $1 \leftrightarrow 3$ rotation of angle 2φ with $\tan 2\varphi = 2D/(C - A)$:

$$V_{CKM}^{(0)} = R_{13}(2\varphi), 2\varphi \simeq 0.039162 \text{ rad} \simeq 2.244^\circ,$$

with $J_q = 0$ and $\mathcal{J}_q^{\text{op}} = \text{Im Tr}([H_u^L, H_d^L]^3) = 0$. There is no physical CP phase; the entire returned operator system admits a simultaneous real representative, in agreement with F11 and Theorem 12. The transversality certificates are $\tau_f = \sigma_{\min}(Y_{f^{(0)}}) \simeq 2.964158 > 0$ for all four sectors.

36.7 What the baseline establishes, and referee notes

The baseline is a genuine result, and it is a negative one at the structural level:

radial potential + flat diamond transport + minimal interface conjugation returns four full-rank matrices in the wrong structural class: $Y_d = Y_e$, $Y_u = Y_v$, one common spectrum, a single real $1 \leftrightarrow 3$ rotation, no CP.

The missing terms therefore have sharply defined jobs:

- **J_C or J_R** must break quark–lepton universality — nothing else in the baseline can separate Y_d from Y_e ;
- **refinement-dependent transport weights** must generate the strong three-level hierarchy — the flat spectrum (5.82, 3.01, 2.96) is nearly degenerate where nature spans five orders of magnitude;
- **nontrivial holonomy or complex record structure** must generate physical CP phases — the flat branch is exactly real;
- **independently derived address maps** must generate the $1 \leftrightarrow 2$ and $2 \leftrightarrow 3$ relative frames — on this branch, the completion-plane sign alone yields only the $1 \leftrightarrow 3$ channel and, by Proposition 1, cannot split the spectra.

Four referee notes bind the branch:

1. **Radial-projector typing.** Per the binding real-versus-complex convention of the Notation section: the complex representation $P_{\text{rad}} = \bigoplus |n_x\rangle\langle n_x|$ assigns stiffness r to the imaginary partners of the n_x directions, whereas Theorem 2's real second variation gives those phase directions zero potential stiffness. The compression onto the real readout basis $\mathcal{B}_Y$ is insensitive to the difference, so every matrix above is unaffected — but the 32-dimensional complex P_{rad} may not be reused inside the Schur reduction of Theorem 4 without first classifying the phase directions among the gauge, null and auxiliary modes.
2. **Inherited scale.** $r = (32/63)^2$ is a conditional cross-paper input from HR-1; it is frozen, hashed and flagged, and its falsifier is inherited with it.
3. **Illustrative angle.** $\theta = \pi/6$ is not a derived diamond result. A future substrate determination of θ replaces this premise; until then every baseline number is conditional on it.
4. **One-sided CY3'.** Only Y_f^{Hess} exists on this branch. The mixed-vertex side Y_f^{mixed} requires the dressed functional; $\Delta_f^{\text{CY3'}}$ therefore remains unevaluated, not zero.

The published branch objects are $H_{\text{cl}}^{(0)}$, $P_Y^{(0)}$, $\mathcal{E}_{\text{fL}}/R^{(0)}$, $Y_u^{(0)}$, $Y_d^{(0)}$, $Y_e^{(0)}$, $Y_v^{(0)}$ — returned without flavour fitting. The branch does not close Paper 3. It is the first calculational baseline against which the missing closure, record and dressed-fermion Jacobians are to be tested, one frozen premise at a time.

37. The corpus-assembly branch — MCHY-1C

The baseline branch of §36 is blind but structurally insufficient. The prior VERSF corpus, however, already contains conditional charged-sector returns that are far stronger than the baseline — derived under their own frozen premises in earlier gates, before this paper's

evaluation standard existed. Honesty requires recording them; the firewall requires quarantining them. This section does both.

Standing of the branch. Everything in this section is a *conditional corpus-returned* object: it inherits the open premises of its source gate (localisation, one-twelfth participation, interface projection, geometric coupling) and it is **not** a blind microscopic return of the present evaluation theorem. Under the non-insertion firewall of §43, no object in this section may be used to select J_Ω , the isotype coefficients β_R , the readout projector or any embedding. The branch is a comparison ledger for the physical evaluation, never an input to it.

37.1 The conditional charged-lepton operator

The Charged-Lepton Absolute Scale gate returns an explicit canonical Yukawa operator on the physical branch basis:

$$Y_e^{\text{CLAS}} = c_\ell \text{diag}(1, e^{(16/3)}, e^{(32/3)/12}), c_\ell = \sqrt{2} \alpha_{\text{int}}^2 / (8\pi).$$

On the frozen geometric-coupling branch $\alpha_{\text{int}} = 7/960$,

$$c_\ell \simeq 2.9918 \times 10^{-6},$$

$$Y_e^{\text{CLAS}} \simeq \text{diag}(2.992 \times 10^{-6}, 6.197 \times 10^{-4}, 1.0696 \times 10^{-2}).$$

(Coefficient convention: the branch numbers fix $c_\ell = \sqrt{2} \alpha_{\text{int}}^2 / (8\pi)$; the alternative reading $\alpha_{\text{int}}^2 / (8\pi\sqrt{2})$ is smaller by a factor of two and does not reproduce them. The convention is frozen here.)

This is an explicit conditional operator *candidate* whose numerical entries are fixed once the CLAS premise stack is granted — and that stack is substantial: the source gate itself records that none of its assumptions 3–11 is independently discharged, with the tau-channel structure and the one-twelfth participation factor among the sharpest gaps. Within that conditionality it is nevertheless much stronger than the structurally insufficient baseline $Y_e^{(0)}$: its entries are numbers, not placeholders.

37.2 Conditional pre-RG quark coefficient ledgers

The quark current-mass projection gate gives the dimensionless down coefficient

$$y_d^{(0)} = 9\sqrt{2} \alpha_{\text{int}}^2 / (8\pi) = 9 c_\ell \simeq 2.6926 \times 10^{-5},$$

and the inherited quark grid gives the coefficient ledgers

$$D_d^{\text{grid}} = y_d^{(0)} \text{diag}(1, 20, 5e^{(16/3)}) \simeq \text{diag}(2.693 \times 10^{-5}, 5.385 \times 10^{-4}, 2.7886 \times 10^{-2}),$$

$$D_u^{\text{grid}} = y_d^{(0)} \text{diag}(6/13, 3080/13, 62154) \simeq \text{diag}(1.243 \times 10^{-5}, 6.379 \times 10^{-3}, 1.67356).$$

These are pre-RG projection/grid coefficient ledgers, not physical singular spectra at one scale. The source gate itself records that its projected anchor is not automatically a pole mass or a high-scale Yukawa: the common-scale spectra require the scheme, running and threshold matching

$$D_d(\mu_\star) = R_d(\mu_\star \leftarrow \mu_{cl})[D_d^{\text{grid}}], D_u(\mu_\star) = R_u(\mu_\star \leftarrow \mu_{cl})[D_u^{\text{grid}}],$$

where R_u, R_d need not act as one universal scalar multiplier across the three generations. The top coefficient 1.67356 is a *top-compression grid target*, explicitly scheme-sensitive in the source gate, and is not promoted to a returned top Yukawa here.

The full matrices in a common weak basis would be

$$Y_d = U_d L D_d(\mu_\star) U_d R^\dagger, Y_u = U_u L D_u(\mu_\star) U_u R^\dagger, V_{CKM} = U_u L^\dagger U_d L,$$

but the relative frames may not be filled from the existing CKM curvature gate: that gate itself records that its crucial complex component z_{C3} must still be derived from the projected closure Hamiltonian. The frames are therefore open blind targets, and no corpus CKM candidate is promoted here.

37.3 Declared matrix convention and conditional charged-sector representatives

Even granting the grid ledgers and a CKM candidate, the full matrices are not unique. Singular spectra fix D_u, D_d ; a CKM matrix fixes only the relative frame $U_u L^\dagger U_d L$. It does not fix the two right frames or the common weak-basis orientation:

$$Y_u = U_u L D_u U_u R^\dagger, Y_d = U_d L D_d U_d R^\dagger.$$

Any displayed full matrix therefore requires a declared convention. The convention frozen for this branch is the positive Hermitian representative:

$$U_d L = U_d R = I, U_e L = U_e R = I, U_u L = U_u R = V_{CKM}^{C\dagger},$$

so that

$$Y_d^C = D_d^{\text{grid}}, Y_e^C = Y_e^{\text{CLAS}}, Y_u^C = V_{CKM}^{C\dagger} D_u^{\text{grid}} V_{CKM}^C.$$

Numerically, the conditional up-sector representative is

$$Y_u^C \simeq \begin{pmatrix} 4.55227 \times 10^{-4} - 1.91262 \times 10^{-3} + 2.29783 \times 10^{-4} i & 1.33187 \times 10^{-2} - 5.67604 \times 10^{-3} i & -1.91262 \times 10^{-3} - 2.29783 \times 10^{-4} i \\ -1.91262 \times 10^{-3} - 2.29783 \times 10^{-4} i & 8.69744 \times 10^{-3} - 6.62222 \times 10^{-2} + 6.01599 \times 10^{-4} i & 1.33187 \times 10^{-2} + 5.67604 \times 10^{-3} i \\ 1.33187 \times 10^{-2} + 5.67604 \times 10^{-3} i & -6.62222 \times 10^{-2} + 6.01599 \times 10^{-4} i & 1.6707955 \end{pmatrix}.$$

Its eigenvalues equal D_u^{grid} (verified to displayed precision), and it generates the conditional CKM frame of §37.4. It is one canonical representative, not a uniquely substrate-selected full matrix, and it inherits every conditionality of the grid ledgers, including the absent RG matching.

37.4 Conditional relative frames

Quark frame. Using the inherited relative generator and the C_3 curvature-residue candidate,

$$V_CKM^C \simeq \begin{pmatrix} 0.9747958 + 0.0000049 i & 0.2230707 - 0.0000013 i & 0.0010709 + 0.0033918 i \\ -0.2229298 + 0.0000297 i & 0.9739912 + 0.0004777 i & 0.0405359 - 0.0003894 i \\ + 0.0033911 i & -0.0397530 + 0.0003767 i & 0.9991716 - 0.0004635 i \end{pmatrix},$$

with magnitude matrix

$$|V_CKM^C| \simeq \begin{pmatrix} 0.974796 & 0.223071 & 0.003557 \\ 0.222930 & 0.973991 & 0.040538 \\ 0.008690 & 0.039755 & 0.999172 \end{pmatrix},$$

and Jarlskog magnitude $|J_q| \simeq 3.00 \times 10^{-5}$.

Binding referee note (sign). Under the standard quartet convention $J_q = \text{Im}(V_us V_cb V_ub^* V_cs^*)$, the displayed matrix returns $J_q \simeq -3.00 \times 10^{-5}$, with the sign stable across all rephasing-invariant quartets. The physical CKM Jarlskog is positive in that convention. Either the branch frame is oriented as the conjugate of the standard assignment or its source uses a different quartet convention; the orientation must be frozen and declared before any physical comparison, per §33.

Conditionality. This is not a blind substrate return: the curvature input $z = (b\sqrt{3}) e^{(5\pi i/6)}$ is explicitly described by its source gate as a geometry-motivated candidate whose phase and amplitude have not been derived from the projected closure Hamiltonian (the open z_C3 debt).

Lepton frame. On the electron–tau, upper-octant, negative- J_lep branch selected by NFCC-1R, the conditional oscillation-frame representative is

$$U_PMNS^C \simeq \begin{pmatrix} 0.825459 & 0.544323 & -0.101462 + 0.109703 i \\ -0.302405 + 0.068453 i & 0.596352 + 0.045139 i & 0.739048 \\ 0.467740 + 0.060842 i & -0.586881 + 0.040120 i & 0.656871 \end{pmatrix},$$

with magnitude matrix

$$|U_PMNS^C| \simeq \begin{pmatrix} 0.825459 & 0.544323 & 0.149430 \\ 0.310056 & 0.598058 & 0.739048 \\ 0.471680 & 0.588251 & 0.656871 \end{pmatrix},$$

and extracted invariants

$$\theta_{12} \simeq 33.40^\circ, \theta_{13} \simeq 8.59^\circ, \theta_{23} \simeq 48.37^\circ, \delta_CP \simeq 227.24^\circ, J_lep \simeq -0.02448$$

(all independently reverified from the displayed matrix). **Precision note.** The displayed matrix is unitary only to $\approx 7 \times 10^{-7}$ at the quoted digits; the neutral-texture coefficients of §37.7 inherit that truncation. The publication package must carry either more digits or the generating parameters ($\theta_{12}, \theta_{13}, \theta_{23}, \delta_CP$) from which the frame is reconstructed exactly. The frame remains

conditional because the source calculation *assumes* the charged-lepton frame, the canonical weak bridge, Frame–Takagi compatibility and active completeness rather than deriving them numerically. Under the declared convention, $U_{eL} = I$ and $U_{\nu}^{\text{oscillation}} = U_{\text{PMNS}}^C$. The Majorana column phases are not returned.

37.5 The frozen neutral block

The strongest neutral return from the frozen branch is the *typed but numerically unevaluated* operator

$$B_{**}^{\text{N}^{\text{frozen}}} = K_N^{(-T/2)} \begin{pmatrix} 0 & (v_{cl}/2) Y_\nu \end{pmatrix} K_N^{(-1/2)}, ** \begin{pmatrix} (v_{cl}/2) Y_\nu^T M_\star R_N \end{pmatrix}$$

with, on the Route-A scalar branch, $v_{cl} = 243.722807638$ GeV. The substrate branch fixes the structural facts:

$$\text{rank } Y_\nu = 3, R_N = R_N^T, \det R_N \neq 0, M_{L^{(0)}} = 0,$$

three light plus three heavy Majorana modes, type-I-dominant seesaw, normal ordering $m_1 < m_2 < m_3$, $\epsilon\tau$ attachment, $\theta_{23} > 45^\circ$, $J_{\text{lep}} < 0$.

But the numerical matrices Y_ν , R_N , K_N and the heavy scale $M_\star$ are explicitly left to the physical evaluation. **No numerical 6×6 neutral block can be returned without inventing those entries**, and the conditional PMNS frame does not remove the degeneracy: infinitely many choices of light masses, heavy masses, Majorana phases and active–sterile factorisations share one oscillation frame.

37.6 CY3' on the corpus branch

The microscopic certificate requires two independently evaluated sides,

$$Y_f^{\text{mixed}} = K_{fL}^{(-1/2)} [\mathcal{D}_Q \delta^2 \Gamma_{\text{sub}}^{\text{dressed}} / \delta \bar{\psi}_{fL} \delta \psi_{fR}]_\star K_{fR}^{(-1/2)} \text{ and } Y_f^{\text{Hess}} = \mathcal{E}_{fL}^\dagger P_Y \hat{H}_{cl} P_Y \mathcal{E}_{fR},$$

with $\Delta_f^{\text{CY3}'} = Y_f^{\text{mixed}} - Y_f^{\text{Hess}}$. The corpus supplies neither a numerical fermion-dressed functional nor two independent sides. The only honest result is therefore

$$\Delta_u^{\text{CY3}'}, \Delta_d^{\text{CY3}'}, \Delta_e^{\text{CY3}'}, \Delta_\nu^{\text{CY3}'} \text{ are unevaluated — not zero.}$$

In particular, *defining* the mixed vertex to equal the assembled corpus representatives would make CY3' true by construction and destroy the intended test. The corpus branch carries no CY3' credit whatsoever.

37.7 The matching-point audit branch — MCHY-1C(α_0)

The strongest corpus-assembled calculation to date replaces the geometric coupling branch with a *matching-point audit branch*, frozen as:

$$\alpha_{\text{int}} = \alpha(0) = 1/137.035999084, v_{\star} = 246.21965 \text{ GeV}, c_{\ell} = \sqrt{2} \alpha_{\text{int}}^2 / (8\pi) = 2.996441458 \times 10^{-6}.$$

Firewall typing (binding). This branch consumes exactly two *measured non-flavour* interface anchors — the fine-structure constant $\alpha(0)$ and the electroweak scale $v_{\star}$ — and zero flavour observables: no fermion mass, mass ratio, mixing angle or CP phase enters. It is therefore an empirically anchored audit branch, distinct in kind from both the fully geometric $\alpha_{\text{int}} = 7/960$ branch of §37.1 (0.156% smaller kernel) and the blind baseline of §36. Its status is conditional corpus assembly, one rung closer to experiment, one rung further from a pure substrate derivation.

Returned singular spectra

With the down anchor $y_{\text{d}}^{(0)} = 9c_{\ell} = 2.696797312 \times 10^{-5}$ from the quark-current projection kernel, and the exact closed-form top coefficient

$$5e^{(16/3)} \cdot (6/13) \cdot (77/3)^{(3/2)} = 62154.04$$

(of which the earlier integer grid entry 62154 was a rounding),

$$\mathbf{D}_{\text{e}} = \text{diag}(2.9964415 \times 10^{-6}, 6.2064468 \times 10^{-4}, 1.0712702 \times 10^{-2}), \mathbf{D}_{\text{d}} = \text{diag}(2.6967973 \times 10^{-5}, 5.3935946 \times 10^{-4}, 2.7929010 \times 10^{-2}), \mathbf{D}_{\text{u}} = \text{diag}(1.2446757 \times 10^{-5}, 6.3893352 \times 10^{-3}, 1.6761684).$$

The ratio closure is algebraically exact once the inherited factors are granted; the $77/3$ charm hinge and the $e^{(16/3)4}$ bottom step remain the most exposed ingredients.

Matching-point mass table

Under $M_{\text{f}} = v_{\star} Y_{\text{f}} / \sqrt{2}$, the *unmatched tree-level* matching-point masses are:

Sector	Matching-point masses
e, μ , τ	0.521691 MeV, 108.056 MeV, 1.86512 GeV
d, s, b	4.69522 MeV, 93.9044 MeV, 4.86254 GeV
u, c, t	2.16702 MeV, 1.11241 GeV, 291.827 GeV

These are not pole masses or common low-energy running masses; radiative evolution, thresholds and scheme matching remain unapplied, exactly as CLAS-1 and QCMP-1 firewall. The identification of $m_{\text{d}}^{(0)} = 4.6952 \text{ MeV}$ with the $\overline{\text{MS}}$ value at 2 GeV is explicitly not yet proved.

Explicit charged matrices and certificates

Under the declared Hermitian convention of §37.3 (the equality of the up left and right frames is a representative choice, not a substrate derivation of the right frame),

$$Y_d^C = D_d, Y_e^C = D_e,$$

$$Y_u^C \approx \begin{pmatrix} 4.559372 \times 10^{-4} - 1.915604 \times 10^{-3} + 2.301418 \times 10^{-4} i & 1.333946 \times 10^{-2} - 5.684899 \times 10^{-3} i & -1.915604 \times 10^{-3} - 2.301418 \times 10^{-4} i \\ -1.915604 \times 10^{-3} - 2.301418 \times 10^{-4} i & 8.711010 \times 10^{-3} - 6.632557 \times 10^{-2} - 6.025382 \times 10^{-4} i & 1.333946 \times 10^{-2} + 5.684899 \times 10^{-3} i \\ 1.333946 \times 10^{-2} + 5.684899 \times 10^{-3} i & -6.632557 \times 10^{-2} + 6.025382 \times 10^{-4} i & 1.673403 \end{pmatrix}, **$$

built from the V_{CKM}^C of §37.4 (unitary to numerical precision; $|J_q| = 2.997 \times 10^{-5}$; the sign-orientation note of §37.4 stands). All three matrices are full rank, with certificates

$$\kappa(Y_e) = 3.575 \times 10^3, \kappa(Y_d) = 1.036 \times 10^3, \kappa(Y_u) = 1.347 \times 10^5, \det Y_e \simeq 1.9923 \times 10^{-11}, \det Y_d \simeq 4.0624 \times 10^{-10}, \det Y_u \simeq 1.3330 \times 10^{-7}.$$

This branch therefore passes the charged-sector rank and hierarchy tests that the blind symmetric-diamond baseline fails — under its inherited conditionality.

The neutral texture return

With the branch-aligned charged-lepton frame $U_e L = I$, zero displayed Majorana phases and unknown Takagi masses m_1, m_2, m_3 , the effective light Majorana operator must be

$$M_{\nu, eff} = U_{PMNS} \text{diag}(m_1, m_2, m_3) U_{PMNS}^\dagger, *$$

whose six independent entries are the explicit linear forms

$$\begin{aligned} M_{ee} &= 0.681422 m_1 + 0.296269 m_2 + (-0.001743 + 0.022241 i) m_3, M_{e\mu} = (-0.249639 - 0.056488 i) m_1 + (0.324585 - 0.024560 i) m_2 + (-0.074946 - 0.081048 i) m_3, \\ M_{e\tau} &= (0.386076 - 0.050205 i) m_1 + (-0.319466 - 0.021828 i) m_2 + (-0.066611 - 0.072034 i) m_3, \\ M_{\mu\mu} &= (0.086772 + 0.041389 i) m_1 + (0.353572 - 0.053814 i) m_2 + 0.546218 m_3, \\ M_{\mu\tau} &= (-0.145601 - 0.013612 i) m_1 + (-0.351809 + 0.002568 i) m_2 + 0.485467 m_3, \\ M_{\tau\tau} &= (0.215042 - 0.056890 i) m_1 + (0.342871 + 0.047075 i) m_2 + 0.431472 m_3 \end{aligned}$$

(all eighteen coefficients independently reverified from the reconstructed frame). This is a genuine return: the entire neutral texture is fixed once three masses and two Majorana phases are supplied. But the corpus does not derive those masses or phases, nor $M_D, M_R, M_L^{\wedge irr}$ — they remain classified open, and the degeneracy statement of §37.5 stands.

The next decisive calculation

This branch fixes the target of the physical evaluation with new sharpness. The mathematically decisive next step is an *inverse universal-closure consistency test*:

determine whether one positive rank-six compression $P_Y \hat{H}_{cl} P_Y$ and one sector-independent address-realisation map J_Ω can reproduce the three charged operators Y_e^C, Y_d^C, Y_u^C simultaneously, without sector-dependent dressing.

A positive answer would show the corpus-assembled flavour data are *compatible* with the single-Hessian architecture and would return candidate $(\hat{H}_{cl}, J_\Omega)$ pairs for the microscopic calculation to hit. A negative answer would prove, ahead of the physical evaluation, that the leading single-insertion architecture cannot hold and that typed higher-order sector dressing (§18) is compulsory. Either answer is decisive, and neither requires the dressed functional.

37.8 What the corpus cannot yet assemble

The neutrino gates are explicit that the numerical M_L^{irr}, M_D, M_R and Y_ν remain uncalculated from substrate dynamics, and the PMNS gate records the full H_{cl} projection and the blind numerical PMNS derivation as open. The typed frozen block of §37.5 and the conditional oscillation frame of §37.4 are the entire neutral inheritance; no conditional numerical neutral matrix exists to record.

37.9 The branch ledger

The programme now carries four distinct strata, which must never be conflated:

MCHY-1(0) — blind, structurally refuted: the baseline of §36; wrong structural class, published as the fixed ruler. **MCHY-1C(geo)** — conditional, corpus-assembled on the geometric coupling $\alpha_{int} = 7/960$: the operator candidate Y_e^{CLAS} ; the pre-RG ledgers D_d^{grid}, D_u^{grid} ; the declared Hermitian representatives; the relative-frame candidates V_{CKM}^C and U_{PMNS}^C ; and the typed neutral block B_N^{frozen} . **MCHY-1C(α_0)** — conditional, empirically anchored: the matching-point audit branch of §37.7, consuming two measured non-flavour anchors ($\alpha(0), v_\star$) and zero flavour observables; full charged matrices, matching-point mass table, rank and hierarchy certificates, and the neutral texture as explicit linear forms in (m_1, m_2, m_3) . **MCHY-1 physical** — open: the blind single-substrate evaluation through J_Ω, β_R, P_Y and the dressed mixed vertex, with the CY3' certificate; not yet executed. Its gateway — the inverse universal-closure consistency test — has now received two executions in §38: the universal magnitude rule passes exactly; the binary rank-six carrier fails analytically; and the three-address rank-nine carrier is realised analytically with an exactly constructed rank-six positive witness (a construction that succeeds for any positive targets under the declared convention; §38.5). The third round has closed the charged inverse calculation conditionally: the magnitude kernel's coefficients are identified with named corpus constants, S_{cl} and T_{cl} are reassembled within the conditional corpus algebra, and rank six is the global minimum conditional on Premise CSE. The physical evaluation and the $\mathcal{H}_{15}$ five-slot test are the remaining steps.

The decisive future test is already fixed by this ledger: when the physical evaluation returns, its Y_e must be compared against Y_e^{CLAS} , and its quark spectra against the RG-matched $D_u(\mu_\star)$ and $D_d(\mu_\star)$ descended from the grid ledgers, *after* freezing — agreement would consolidate the conditional gates, disagreement would falsify one side at a named premise.

Pre-registered post-freeze comparison targets. The α_0 matrices being frozen and published, post-freeze comparison against literature running masses is licit under Stage 13. A post-freeze audit of the implied per-generation dressing factors $R_f = m^{\text{lit}}(M_Z)/m^{\text{match}}$ yields two named targets, pre-registered here for the future matching gate:

The literature inputs are standard $\overline{\text{MS}}$ running masses evaluated at $\mu = M_Z$, of the class compiled in the standard running-mass tables: $m_e \approx 0.4866$ MeV, $m_\mu \approx 102.7$ MeV, $m_\tau \approx 1746$ MeV; $m_u \approx 1.3$ MeV, $m_d \approx 2.7$ MeV, $m_s \approx 55$ MeV, $m_c \approx 0.62$ GeV, $m_b \approx 2.86$ GeV, $m_t \approx 168.3$ GeV. Without these the quoted factors are unreproducible; they are recorded here for exactly that reason.

1. **The quark sector-universal dressing factor.** All six quark factors cluster tightly ($\approx 0.580 \pm 0.013$) across both sectors: the entire quark discrepancy is absorbable by a single sector-universal multiplicative dressing — precisely the structure a QCD-dominated $R_u \approx R_d$ would produce. The grid ratios survive contact with data; the debt is one common factor. (Its numerical proximity to $1/\sqrt{3} \approx 0.5774$ is recorded with explicit suspicion, given the programme's history with triality factors: proximity is not derivation.)
Uncertainty caveat: $m_u(M_Z)$ and $m_d(M_Z)$ carry $\approx 30\text{--}40\%$ uncertainties and $m_s(M_Z) \approx \pm 15\%$, so the clustering claim rests genuinely on the charm, bottom and top factors (0.557, 0.588, 0.577), with u and d consistent but not constraining.
2. **The lepton residual — the sharpest exposure.** The three lepton factors (e: 0.933, μ : 0.951, τ : 0.936) cluster at $\approx 0.940 \pm 0.008$ with a *non-monotone* residual pattern across generations. Lepton masses are known to ten digits, so this target carries no input-uncertainty caveat; and Standard Model lepton-Yukawa running is a percent-level, monotone-in-generation effect that cannot supply a common $\approx 6\%$ suppression with non-monotone residuals. This target is RG-immune in a way the quark discrepancy is not, and it is the sharpest quantitative exposure of the CLAS/ α_0 stack. The future matching gate must either derive it or fail on it.

38. The inverse universal-closure consistency test — first execution

The gateway test pre-registered in §37.7 has now been executed in its properly constrained form. An unconstrained inverse test is too weak — if the embeddings may be reconstructed freely from the target matrices, almost any spectra can be "explained". The executed version therefore imposes the VERSF structure as hard constraints:

- sector dependence must come from *fixed address coordinates*, not per-sector free matrices;
- the closure-side operator is *universal* — one operator for all sectors;
- no independently chosen 3×3 embedding is allowed for any sector;
- generation, quark/lepton status and upper/lower completion may interact.

The targets are the MCHY-1C(α_0) charged operators of §37.7. The execution returns a partial pass and a sharp analytic failure, and both are informative.

38.1 The structured nonfactorising magnitude test — pass, exactly

Assign each charged sector two address bits: $q = 0$ (charged lepton) or 1 (quark), and $t = +1$ ($u\tilde{H}$) or -1 (dH). Consider one universal log-magnitude kernel on generation space,

$$\mathbf{L}(q, t) = \mathbf{A} + q\mathbf{B} + qt\mathbf{D},$$

which is deliberately *not* factorised: the $qt\mathbf{D}$ term is a genuine interaction between quark status and completion orientation. The three sectors read

$$\log D_e = A, \log D_d = A + B - D, \log D_u = A + B + D.$$

Solving against the matching-point spectra gives, exactly,

$$A = \text{diag}(-12.718085, -7.384752, -4.536325), B = \text{diag}(1.810630, 1.095625, 3.005536), D = \text{diag}(-0.386595, 1.236002, 2.047300),$$

and exponentiating the three sector readings reproduces all nine charged singular values to machine precision.

Proposition 2 — Exact Structure of the Magnitude Kernel

On the MCHY-1C(α_0) targets:

(i) The charged-lepton baseline is *identically* the CLAS operator,

$$\mathbf{A} = \ln \mathbf{c}_\ell \cdot \mathbf{I} + (16/3) \mathbf{G} - \ln 12 \cdot \mathbf{\Pi}_3, \mathbf{G} = \text{diag}(0, 1, 2), \mathbf{\Pi}_3 = \text{diag}(0, 0, 1),$$

holding to machine precision (residual $< 10^{-15}$). **Scope of (i):** because the α_0 targets of §37.7 were themselves constructed as $\mathbf{c}_\ell$ times the CLAS ladder, this identity is a consistency audit of the test pipeline — true by construction of its inputs — not a discovered cross-paper interlock. It certifies that the inverse solve returns its inputs faithfully; it does not add evidence for CLAS.

(ii) The role-odd operator $\mathbf{D}$ has *exactly* the half-lazy profile: its increments are the algebraic identities

$$\mathbf{D}_2 - \mathbf{D}_1 = \frac{1}{2} \ln(77/3) = 1.622597, \mathbf{D}_3 - \mathbf{D}_2 = \frac{1}{4} \ln(77/3) = 0.811298,$$

so $(\mathbf{D}_2 - \mathbf{D}_1)(\mathbf{D}_3 - \mathbf{D}_2) = 2$ exactly — not approximately. Equivalently $\mathbf{D}_g = \mathbf{D}_0 + 2\Delta(1 - 2^{(-g)})$ with $\Delta = \frac{1}{2} \ln(77/3)$.

Proof. (i) is direct substitution of the CLAS ladder into $\log D_e$. (ii): $\mathbf{D}_g = \frac{1}{2} \ln(\mathbf{D}_{u,g}/\mathbf{D}_{d,g})$; the grid ratios are $6/13, (3080/13)20 = (6/13)(77/3)^2/(77/3)\dots$ explicitly, $\mathbf{D}_2 - \mathbf{D}_1 = \frac{1}{2} \ln[(3080/260) \cdot (13/6)] = \frac{1}{2} \ln(77/3)$ and $\mathbf{D}_3 - \mathbf{D}_2 = \frac{1}{2} \ln[(6/13)(77/3)^{(3/2)}/(3080/260) \cdot \dots] = \frac{1}{4} \ln(77/3)$, both exact consequences of the $77/3$ charm hinge in the inherited grid. ■

The half-lazy profile is therefore not a numerical coincidence discovered by fitting: it is an exact algebraic identity of the grid — a previously unstated exact fact of QCMP-1's frozen constants — and it is the genuinely new discovery of this subsection, alongside the positivity architecture of §§38.2–38.5. Unlike (i), it *could* have failed: nothing in the construction of the targets forces the fold increments into an exact 2:1 ratio.

The remaining unexplained magnitude object is concentrated entirely in $B = \text{diag}(1.810630, 1.095625, 3.005536)$, naturally interpreted as the generation-dependent colour/confinement/commitment correction. The occupancy map permits exactly this interaction structure, because its addresses independently retain generation, species, colour, weak component, chirality and $H/\tilde{H}$ route — it does not reduce matter to generation plus one binary label.

Standing. The test is reconstructive: B is inferred from the conditional spectra, not derived from H_{cl} . What it establishes is *mutual algebraic compatibility*: one universal nonfactorising address rule can carry all three charged spectra, with its lepton baseline and quark fold profile matching the independently proposed VERSF mechanisms identically.

38.2 The strict binary rank-six compression test — analytic fail

The strong version imposes $\mathcal{H}_Y = \mathbb{C}^2_{\text{address}} \otimes \mathbb{C}^3_{\text{generation}}$ with one universal positive operator $C_Y \succcurlyeq 0$ of rank at most six and fixed embeddings

$$E_f = v_f \otimes I_3, Y_f = (v_f^\dagger \otimes I_3) C_Y (v_f \otimes I_3),$$

with the natural binary fold assignment $v_e = (1, 0)^T$, $v_u = (1, 1)^T/\sqrt{2}$, $v_d = (1, -1)^T/\sqrt{2}$.

Proposition 3 — Binary Rank-Six Compression No-Go (conditional on the MCHY-1C(α_0) targets)

No Hermitian $C_Y \succcurlyeq 0$ on $\mathbb{C}^2 \otimes \mathbb{C}^3$ reproduces the three charged operators under the binary fold assignment.

Proof. The three compressions fix the block structure completely up to one Hermitian freedom: writing C_Y in 3×3 blocks, $Y_e = C_{11}$, $Y_u + Y_d = C_{11} + C_{22}$ and $Y_u - Y_d = C_{12} + C_{12}^\dagger$, so

$$C_Y = \begin{pmatrix} Y_e & H + iK \\ Y_e H + iK & H - iK Y_u + Y_d - Y_e \end{pmatrix}, \quad H = \frac{1}{2}(Y_u - Y_d), \quad K = K^\dagger$$

with K the only remaining freedom, invisible to the three real-address compressions. Positivity requires every principal minor to be nonnegative; on the third-generation subspace $\text{span}\{e_1 \otimes g_3, e_2 \otimes g_3\}$ it requires

$$|H_{33} + iK_{33}|^2 \leq (Y_e)_{33} (Y_u + Y_d - Y_e)_{33}.$$

Numerically $H_{33} = \frac{1}{2}(1.673403 - 0.027929) = 0.822737$, so $H_{33}^2 = 0.676896$, while the right-hand side is $(0.0107127)(1.690619) = 0.0181111$. The inequality is violated by the factor

$$\mathbf{H}_{33}^2 / [(\mathbf{Y}_{\mathbf{e}})_{33} (\mathbf{Y}_{\mathbf{u}} + \mathbf{Y}_{\mathbf{d}} - \mathbf{Y}_{\mathbf{e}})_{33}] = 37.37,$$

and since K_{33} is real, adding iK_{33} only *increases* the modulus. No K repairs it. ■

The failure has a transparent physical reading: the top–bottom splitting (carried by H_{33}) is far too large relative to the geometric mean of the tau and the up+down third-generation weight. A single binary address coordinate cannot hold the top quark and the tau lepton in one positive rank-six operator.

Proposition 3' — Frame-Invariant Binary No-Go

Proposition 3 as stated uses the declared positive-Hermitian representatives, and the physical data do not fix the right-handed frames. The no-go nevertheless survives invariantly.

Lemma (forced Hermiticity). With the symmetric binary selectors $\mathbf{E}_{\mathbf{f}} = \mathbf{v}_{\mathbf{f}} \otimes \mathbf{I}_3$ applied on both sides, and $\mathbf{v}_{\mathbf{e}} = (1, 0)^T$, $\mathbf{v}_{\mathbf{u}} = (1, 1)^T/\sqrt{2}$, $\mathbf{v}_{\mathbf{d}} = (1, -1)^T/\sqrt{2}$, the compressions are explicitly

$$\mathbf{Y}_{\mathbf{e}} = C_{11}, \mathbf{Y}_{\mathbf{u}} = \frac{1}{2}(C_{11} + C_{12} + C_{21} + C_{22}), \mathbf{Y}_{\mathbf{d}} = \frac{1}{2}(C_{11} - C_{12} - C_{21} + C_{22}),$$

so that

$$\mathbf{Y}_{\mathbf{u}} + \mathbf{Y}_{\mathbf{d}} = C_{11} + C_{22}, \mathbf{Y}_{\mathbf{u}} - \mathbf{Y}_{\mathbf{d}} = C_{12} + C_{21} = C_{12} + C_{12}^\dagger,$$

both manifestly Hermitian for Hermitian C . The binary architecture can therefore only ever host *Hermitian* representatives of the physical equivalence classes, and the invariant question is whether positivity holds for *any* Hermitian representative — any frames, any eigen-sign choices — compatible with the physical singular values.

Claim. It does not, with violation bounded below using singular values alone:

$$\text{violation} \geq [\frac{1}{2}(\sigma_{\max}(\mathbf{Y}_{\mathbf{u}}) - \sigma_{\max}(\mathbf{Y}_{\mathbf{d}}))]^2 / [\sigma_{\max}(\mathbf{Y}_{\mathbf{e}})(\sigma_{\max}(\mathbf{Y}_{\mathbf{u}}) + \sigma_{\max}(\mathbf{Y}_{\mathbf{d}}))] = 37.20.$$

Proof. (i) $\mathbf{Y}_{\mathbf{e}} = C_{11} \geq 0$ forces all-positive lepton eigenvalues, and since the physical $\sigma_{\min}(\mathbf{Y}_{\mathbf{e}}) > 0$, in fact $\mathbf{Y}_{\mathbf{e}} > 0$. Then $C_{22} = \mathbf{Y}_{\mathbf{u}} + \mathbf{Y}_{\mathbf{d}} - \mathbf{Y}_{\mathbf{e}} \geq 0$ forces the extreme eigenvalue of $\mathbf{Y}_{\mathbf{u}}$ to be $+\sigma_{\max}(\mathbf{Y}_{\mathbf{u}})$: were it $-\sigma_{\max}(\mathbf{Y}_{\mathbf{u}})$, its unit eigenvector $\mathbf{v}$ would give

$$\langle \mathbf{v}, C_{22} \mathbf{v} \rangle = -\sigma_{\max}(\mathbf{Y}_{\mathbf{u}}) + \langle \mathbf{v}, \mathbf{Y}_{\mathbf{d}} \mathbf{v} \rangle - \langle \mathbf{v}, \mathbf{Y}_{\mathbf{e}} \mathbf{v} \rangle \leq -\sigma_{\max}(\mathbf{Y}_{\mathbf{u}}) + \lambda_{\max}(\mathbf{Y}_{\mathbf{d}}) - \lambda_{\min}(\mathbf{Y}_{\mathbf{e}}) < 0,$$

the final strict inequality holding because $\lambda_{\min}(\mathbf{Y}_{\mathbf{e}}) > 0$ makes the bound strictly worse than $-\sigma_{\max}(\mathbf{Y}_{\mathbf{u}}) + \sigma_{\max}(\mathbf{Y}_{\mathbf{d}}) < 0$ already is. (ii) By Weyl, $\|\mathbf{H}\| = \|\frac{1}{2}(\mathbf{Y}_{\mathbf{u}} - \mathbf{Y}_{\mathbf{d}})\| \geq \frac{1}{2}(\lambda_{\max}(\mathbf{Y}_{\mathbf{u}}) - \lambda_{\max}(\mathbf{Y}_{\mathbf{d}})) \geq \frac{1}{2}(\sigma_{\max}(\mathbf{Y}_{\mathbf{u}}) - \sigma_{\max}(\mathbf{Y}_{\mathbf{d}})) = 0.82412$. (iii) At $\mathbf{x}$ the extreme eigenvector of $\mathbf{H}$, the principal 2×2 minor on $\text{span}\{\mathbf{e}_1 \otimes \mathbf{x}, \mathbf{e}_2 \otimes \mathbf{x}\}$ requires $\mathbf{H}_{\mathbf{xx}^2} + \mathbf{K}_{\mathbf{xx}^2} \leq (\mathbf{Y}_{\mathbf{e}})_{\mathbf{xx}} (C_{22})_{\mathbf{xx}} \leq \sigma_{\max}(\mathbf{Y}_{\mathbf{e}})(\sigma_{\max}(\mathbf{Y}_{\mathbf{u}}) + \sigma_{\max}(\mathbf{Y}_{\mathbf{d}})) = 0.018256$, while

$H_{xx}^2 = \|H\|^2 \geq 0.67917$. The ratio is ≥ 37.20 for every admissible frame and sign assignment (Monte-Carlo over random frames and signs confirms, worst case ≈ 38.5). ■

The binary exclusion is therefore a property of the physical singular spectra themselves — no representative choice, frame convention or sign assignment rescues it.

Scope remark. The invariance of Proposition 3' is over Hermitian representatives, which the forced-Hermiticity Lemma shows is the complete equivalence class the *symmetric-selector* binary architecture can host. An architecture with independent left and right address vectors per sector could host non-Hermitian representatives; that generalisation is the floating-directions family of §38.3, explored there numerically (near-collinear escapes only, at $\kappa \approx 6.4 \times 10^9$) but not yet given a fully invariant analytic treatment. A frame-invariant no-go for asymmetric binary selectors remains an open, well-posed problem. The trine bound of §38.3 likewise was computed for the declared conditional representatives and has not been made frame-invariant; its scope is stated accordingly there.

38.3 Complex trine and floating directions — fail and pathological escape

Complex trine. Assigning the three sectors to the $\mathbb{C}_3$ -separated directions $v_k = (1, e^{(2\pi i k/3)})^T / \sqrt{2}$ — attractive because VERSF already uses C_3 holonomy in the quark frame — also fails positivity *for the declared conditional representatives* (unlike the binary case, this exclusion has not been made frame-invariant; see Proposition 3'). Concave maximisation of $\lambda_{\min}$ over the full nine-real-parameter Hermitian completion family, numerically converged across three independent optimisation methods — subgradient ascent from multiple starts, annealed soft-min optimisation, and convex feasibility by alternating projection between the affine constraint family and the shifted PSD cone — gives

$$\lambda_{\min}^{\max}(C_Y) = -0.53305,$$

identical for all six sector-to-trine assignments, with feasibility failing for every threshold above -0.5331 . A formal certificate (the exact affine family, the optimisation formulation, primal and dual solutions with a duality-gap bound, and code and input hashes) is a publication-package requirement, per §43.

Caveat (binding). The *mixed* problem in which the electron sector is pinned to the pole direction $(1, 0)^T$, with the two quark sectors on trine directions, has $\lambda_{\min}^{\max} \approx -0.401$ for every quark assignment. It is a different problem and is not the trine bound; the trine bound is -0.53305 as stated, and the qualitative verdict — strictly negative, both geometries excluded — holds throughout.

Floating directions. Allowing the two-dimensional address directions to vary freely produces a positive exact solution *only* in a pathological corner: the fitted electron, down and up address lines collapse to projective separations of 0.210° , 1.040° and 1.250° — essentially one line — with universal-operator spectrum spanning $(6.9 \times 10^{-7}, \dots, 4422.8)$ and condition number $\kappa(C_Y) \approx 6.4 \times 10^9$. Constraining the lines to remain separated, the best numerical searches

return $\lambda_{\min} \simeq -0.0101$ at 15° minimum separation, -0.0998 at 30° , and -0.265 at 45° (reported numerical searches, not exhaustive no-gos). The pattern is unambiguous:

a well-separated two-dimensional address factor does not support the charged operators.

38.4 Verdict and the revised carrier hypothesis

Test	Verdict
Universal nonfactorising address rule for singular magnitudes	Pass, exactly
CLAS baseline plus exact half-lazy quark fold profile	Pass, exactly (algebraic identity)
Natural real binary rank-six positive compression	Fail, analytically (factor 37.37)
Complex C_3 trine rank-six compression	Fail for the declared representatives ($\lambda_{\min}^{\max} = -0.533$, all assignments)
Frame-invariant binary no-go	Proved from singular values alone (Proposition 3', factor ≥ 37.20)
Arbitrary two-dimensional address directions	Pass only via near-collinearity at $\kappa \simeq 6.4 \times 10^9$
Fully derived universal H_{cl} compression	Still open

The viable possibility survives, but its form has sharpened:

one universal rule is viable; one generation $\otimes$ one binary-address rank-six carrier is (conditionally) not.

The corpus itself points to why. The occupancy map distinguishes five matter-slot types per generation — Q_i, U_i, D_i, L_i, E_i — not a single binary $H/\tilde{H}$ coordinate. The projected weak-doublet programme already uses separate generation, weak-role and readout projections, with $\Omega_W = \Omega_{\text{even}} \otimes I + \Omega_{\text{odd}} \otimes \tau_3 + \Omega_{\text{mix}}$, and attributes CKM structure to the noncommutation of two projections of one closure Hamiltonian. And CMSY-1 explicitly allows the primitive rank-six linear carrier to be insufficient, with composite and interface-dressed channels living on an enlarged effective carrier.

The revised physical hypothesis is therefore:

$\mathcal{H}_Y \neq \mathbb{C}^3_{\text{generation}} \otimes \mathbb{C}^2_{\text{binary}}$ by itself.

A viable minimal carrier likely requires at least one additional address direction distinguishing colourless-H, coloured-H and coloured- $\tilde{H}$ completion routes:

$\mathcal{H}_Y^{\text{next}} \simeq \mathbb{C}^3_{\text{generation}} \otimes \mathbb{C}^3_{\text{address}}$,

to be tested *before* any reduction of its physical rank is attempted. This hypothesis was registered as the frozen target of the second inverse-test round — which has now been executed, with the results of §38.5.

38.5 The rank-nine three-address carrier — second execution: pass

The registered successor hypothesis has now been tested, and it passes — with an analytic positivity proof rather than a numerical search.

Carrier and fixed selectors

Take $\mathcal{H} = \mathbb{C}^3_{\text{address}} \otimes \mathbb{C}^3_{\text{generation}}$ with orthonormal address basis $|E_c\rangle, |D\rangle, |U\rangle$, where $|D\rangle$ and $|U\rangle$ represent the right-terminal quark closure classes d_R, u_R and $|E_c\rangle$ is the *completed charged-lepton channel* — a reduced three-address *quotient* of the full five-slot occupancy space, whose charged closure routes $u_R : \tilde{H}, d_R : H, e_R : H$ the occupancy map fixes. Explicitly, the quotient map is

$$\pi_{\text{charged}} : \{|Q\rangle, |U\rangle, |D\rangle, |L\rangle, |E\rangle\} \rightarrow \{|Q\rangle, |R\rangle, |E_c\rangle\},$$

and the reduced line $|E_c\rangle$ is the image of the *distinct* L_L and e_R occupancy addresses after completion-interface contraction. It is not an identification of their pre-completion occupancy slots — in the full ledger, e_L belongs to the L_L slot and e_R to the E_R slot, and they carry different chiral addresses. Deriving this quotient from J_Q remains open.

The right selectors are

$$\mathcal{E}_{eR} = |E_c\rangle \otimes I_3, \mathcal{E}_{dR} = |D\rangle \otimes I_3, \mathcal{E}_{uR} = |U\rangle \otimes I_3.$$

The left selectors impose the quark-doublet lock through one common quark line and its orthogonal completion contrast,

$$|Q\rangle = (|D\rangle + |U\rangle)/\sqrt{2}, |R\rangle = (|D\rangle - |U\rangle)/\sqrt{2}, \mathcal{E}_{dL} = \mathcal{E}_{uL} = |Q\rangle \otimes I_3, \mathcal{E}_{eL} = |E_c\rangle \otimes I_3,$$

with $\langle E_c|Q\rangle = \langle E_c|R\rangle = \langle Q|R\rangle = 0$: the address directions are cleanly separated, and no selector is fitted.

Proposition 4 — Three-Address Schur Realisation

Define, in the address basis $(|E_c\rangle, |Q\rangle, |R\rangle)$, the one-parameter family

$$\begin{pmatrix} C_{Y^q} \\ R_\mu \end{pmatrix}(\mu) = \begin{pmatrix} Y_e & 0 & 0 \\ 0 & 1 & 0 \\ 0 & 0 & 1 \end{pmatrix}, S = (Y_d + Y_u)/\sqrt{2}, T = (Y_d - Y_u)/\sqrt{2}, ** \begin{vmatrix} 0 & S & T \\ * & * & * \\ * & * & * \end{vmatrix} \begin{pmatrix} 0 \\ T^\dagger \end{pmatrix}$$

$$R_\mu = T^\dagger S^{-1} T + \mu I_3, \mu \geq 0. **$$

Then, on the MCHY-1C(α_0) targets:

(i) $S > 0$, with $\text{eig}(S) \simeq (4.4844 \times 10^{-5}, 4.9139 \times 10^{-3}, 1.204947)$; (ii) $C_Y^{(9)}(\mu) \geq 0$ for all $\mu \geq 0$, and $C_Y^{(9)}(\mu) > 0$ for all $\mu > 0$; (iii) the fixed selectors recover all three charged operators exactly:

$$\mathcal{E}_{\text{eL}}^\dagger C_Y^{(9)} \mathcal{E}_{\text{eR}} = Y_{\text{e}}, \mathcal{E}_{\text{dL}}^\dagger C_Y^{(9)} \mathcal{E}_{\text{dR}} = (S + T)^{1/2} = Y_{\text{d}}, \mathcal{E}_{\text{uL}}^\dagger C_Y^{(9)} \mathcal{E}_{\text{uR}} = (S - T)^{1/2} = Y_{\text{u}}.$$

Proof. (i) is direct computation. (ii): the quark block has Schur complement $R_\mu - T^\dagger S^{-1} T = \mu I_3 \geq 0$ with $S > 0$, so the block is positive semidefinite, strictly positive for $\mu > 0$; the lepton block $Y_{\text{e}} > 0$ is decoupled. (iii): $|D\rangle = (|Q\rangle + |R\rangle)^{1/2}$ gives $\langle Q| C |D\rangle = (C_{\text{QQ}} + C_{\text{QR}})^{1/2} = (S + T)^{1/2} = Y_{\text{d}}$, and similarly for $|U\rangle$ with the sign flipped; the lepton compression is trivial. Numerical reconstruction residuals: $0, \lesssim 10^{-16}, \lesssim 10^{-16}$ in Frobenius norm. ■

This is the analytic converse of Proposition 3: what one binary address plane cannot hold positively, one well-separated three-address carrier holds exactly.

Binding vacuity note. Under the declared Hermitian-representative convention, Proposition 4 holds for *arbitrary* targets: any three positive spectra and any unitary relative frame whatsoever. $S = (Y_{\text{d}} + Y_{\text{u}})^{1/2}$ is a sum of two positive-definite Hermitian matrices and is therefore automatically positive; the Schur family then always exists, the $\mu = 0$ rank-six witness always exists, and Proposition 8's classification is likewise target-independent. The three-address construction therefore *cannot fail on any conceivable charged flavour data* — it is compatible-by-construction, not compatible-by-test. Its content is architectural: the existence and complete classification of the minimal carrier the data do not exclude. The data-discriminating content of this section resides entirely in §§38.2–38.3 — the binary no-go, which Proposition 3' makes frame-invariant (the physical singular values themselves falsify the six-dimensional carrier by a factor ≥ 37.20), and the trine exclusion for the declared representatives. This note mirror-images the Standing note of §38.1 and is binding on every downstream citation of the "pass".

Proposition 5 — Rank-Six Canonical Schur Witness

At $\mu = 0$ the quark block factorises as

$$\begin{pmatrix} S & T \end{pmatrix} = \begin{pmatrix} S^{1/2} \end{pmatrix} \begin{pmatrix} S^{1/2} S^{-1/2} T \end{pmatrix}, \quad \begin{pmatrix} T^\dagger & T^\dagger S^{-1} T \end{pmatrix} \begin{pmatrix} T^\dagger S^{-1/2} \end{pmatrix}$$

so it has rank three, and with the full-rank lepton block,

$$\text{rank } C_Y^{(9)}(0) = 3 + 3 = 6.$$

Scope. At $\mu = 0$ the canonical Schur completion constructed in Proposition 4 has rank exactly six, and it is the minimum-rank member of the displayed μI_3 Schur family. Global minimality over *every* positive universal operator compatible with the fixed selectors is not proved at this stage — that upgrade arrives in §38.6, where colour-sector equivariance (Premise CSE) makes six a genuine global minimum (Proposition 7).

Consequence — the reconciliation. The carrier needs nine address-generation dimensions, and a rank-six positive readout witness exists exactly inside it. The earlier rank-six intuition was right that a rank-six readout can carry the charged sector, and wrong about placing it inside a six-dimensional carrier with a binary address coordinate. The corrected picture is

$$\dim \mathcal{H}_\Omega = 9, \exists \hat{C}_Y \geq 0 \text{ with rank 6, rank } P_Y^{\text{physical open}} (\leq 9).$$

The strictly positive representative and its conditioning

Within the canonical Schur family, the best-conditioned strictly positive choice sets the smallest quark-sector eigenvalue equal to the electron floor,

$$\mu_\star = 6.86154 \times 10^{-5},$$

giving the full spectrum

$$\text{eig } C_Y^{(9)}(\mu_\star) \simeq (2.99644 \times 10^{-6}, 2.99644 \times 10^{-6}, 3.57727 \times 10^{-5}, 6.74767 \times 10^{-5}, 5.74038 \times 10^{-4}, 6.20645 \times 10^{-4}, 8.85608 \times 10^{-3}, 1.07127 \times 10^{-2}, 2.33277),$$

hence rank 9 and $\kappa(C_Y^{(9)}) \simeq 7.785 \times 10^5$. That conditioning is essentially forced: for any universal operator compressed by partial isometries, $\lambda_{\max}(C_Y) \geq \|Y_{\text{ul}}\|_2$ and $\lambda_{\min}(C_Y) \leq \lambda_{\min}(Y_{\text{e}})$, so

$$\kappa(C_Y) \geq \|Y_{\text{ul}}\|_2 / \lambda_{\min}(Y_{\text{e}}) \simeq 5.594 \times 10^5,$$

and the constructed operator sits at only 1.39 times this absolute hierarchy bound — its conditioning is the physical top-to-electron hierarchy, not the pathological near-collinear geometry of §38.3, which it improves by a factor of $\approx 8.2 \times 10^3$.

What the extra address direction accomplished

The failed binary model forced all charged sectors through one two-dimensional address plane. The successful model separates the colourless completed charged-lepton channel $|E_{\text{c}}\rangle$ from the coloured quark completion plane $\text{span}\{|D\rangle, |U\rangle\}$, inside which $|Q\rangle$ is the common left-doublet attachment and $|R\rangle$ the $H/\tilde{H}$ completion contrast. That is a natural VERSF reading — common quark attachment + completion contrast + separate colourless lepton channel — and it remains a quotient of the complete five-slot address ledger rather than its final realisation.

The registered successor test. The stricter round descending from the exact occupancy ledger is the fifteen-dimensional fixed-selector test on

$$\mathcal{H}_{15} = \mathbb{C}^5_{\text{occupancy}} \otimes \mathbb{C}^3_{\text{generation}}, \text{ address basis } \{|Q\rangle, |U\rangle, |D\rangle, |L\rangle, |E\rangle\},$$

with distinct L_L and e_R chiral addresses. It would determine whether the successful three-address quotient genuinely descends from the exact occupancy ledger, or only from its charged contraction.

Scorecard

Question	Result
Fixed, non-collinear selectors	Pass
Shared Q_L selector for up and down	Pass
Distinct u_R, d_R, e_R selectors	Pass
One universal Hermitian operator	Pass
Positive semidefinite operator	Pass, analytically
Strictly positive rank-nine operator	Pass
Exact recovery of all charged matrices	Pass — unconditionally, for any targets (vacuity note)
Extreme selector collapse	Avoided
Conditioning beyond the physical hierarchy bound	Avoided (factor 1.39)
Full five-slot occupancy realisation ($\mathcal{H}_{15}$ test, distinct L_L/e_R addresses)	Not yet — registered
Right-handed physical frame derivation	Not yet
Operator derived from substrate rather than reconstructed	Not yet
CY3'	Not yet

The precise remaining warning, and the third round

The operator above was *reconstructed* from the target matrices; the test proves compatibility, not derivation. The next non-vacuous round is sharply posed:

derive the blocks $S = (Y_d + Y_u)/\sqrt{2}$ and $T = (Y_d - Y_u)/\sqrt{2}$ from the small VERSF operator algebra — generation depth, colour participation, completion orientation, saturation and projected C_3 curvature — without using Y_u or Y_d as input.

The magnitude kernel of §38.1 already shows the diagonal content of that algebra (the CLAS baseline, the exact half-lazy fold, the colour correction B); what remains is to produce S and T themselves from it. This round has now been executed — see §38.6.

38.6 The third round — conditional closure of the charged inverse calculation

The registered third round has now been executed, and it closes the charged inverse calculation at the conditional corpus level: fixed three-channel carrier, one colour-equivariant positive operator, all charged matrices recovered analytically, no unidentified freedom remaining. The physical substrate calculation stays open — J_Ω , P_Y^{micro} , the dressed mixed vertex and CY3' are untouched by this round.

The completion-channel carrier

The nine-dimensional carrier acquires its correct representation-theoretic definition, replacing the quotient description. Define three *completed interaction channels*,

$$C_e = \text{Hom}_{\{G_{EW}\}}(E_R, L_L \otimes \Phi_{cl}), C_d = \text{Hom}_{\{G_{EW}\}}(D_R, Q_L \otimes \Phi_{cl}), C_u = \text{Hom}_{\{G_{EW}\}}(U_R, Q_L \otimes \tilde{\Phi}_{cl}),$$

each carrying three generation slots, so that

$$\mathcal{H}_{ch} = C_e \oplus C_d \oplus C_u \simeq \mathbb{C}^3_{channel} \otimes \mathbb{C}^3_{generation}.$$

This is fully consistent with the occupancy ledger and the derived $H/\tilde{H}$ completion routes, and it does *not* collapse the distinct pre-completion chiral addresses: $|E_c\rangle$ is now, by definition, the electroweak-invariant completion channel built from the distinct L_L and E_R slots, not an identification of them. The quotient map $\pi_{charged}$ of §38.5 is thereby superseded by an intrinsic construction.

Premise CSE — Colour-Sector Equivariance (a premise about the microscopic realisation)

Premise CSE. The universal scalar-readout operator acts on a substrate carrier in which the quark completion channels occupy colour-nonsinglet directions with no colour-singlet component surviving into the readout support, so that the lepton–quark cross blocks vanish:

$$P_e C_Y P_q = 0.$$

Why this is a premise and not representation theory. Note first that G_{EW} here excludes $SU(3)_c$; the channel colour indices remain open. The colour content of a quark completion channel is $\mathbf{3} \otimes \mathbf{3} = \mathbf{1} \oplus \mathbf{8}$ — it *contains* a colour singlet, and it is precisely the singlet contraction that a gauge-invariant Yukawa term retains. At the post-contraction channel level all three channels are colour singlets, and Schur's lemma forbids nothing between them. The vanishing of the cross-blocks is therefore a genuine claim about the *microscopic* address realisation — about how J_Ω and P_Y^{micro} embed the quark completion content in the substrate — which is exactly the pair of objects §38.7 proves non-identifiable from present inputs.

Falsifier. If the substrate readout support of the quark completion channels contains a colour-singlet component overlapping the lepton channel, Premise CSE fails, and the rank-six minimum of Proposition 7 reverts to the rank-greater-than-three bound of Theorem 6.

Under Premise CSE, the block separation holds and upgrades the rank-six witness of Proposition 5 into a genuine minimum (Proposition 7 below). Proposition 7 itself stands as conditional mathematics either way.

Proposition 6 — Exact Corpus-Coefficient Identification of the Diagonal Magnitude Kernel

Let $G = \text{diag}(0, 1, 2)$, $\Pi_3 = \text{diag}(0, 0, 1)$, and set $L = 16/3$, $h = \ln(77/3)$, $r = \ln(6/13)$. Then the three log-kernel pieces of §38.1 are the exact elements of the commutative algebra $\text{span}\{I, G, \Pi_3\}$:

$$A_\ell = \ln c_\ell \cdot I + L G - \ln 12 \cdot \Pi_3, D_\chi = (r/2) I + (h/2) G - (h/4) \Pi_3, B_{\text{col}} = (\ln 9 + r/2) I + (\ln 20 + h/2 - L) G + (L + \ln(3/20) - h/4) \Pi_3,$$

with numerical values $B_{\text{col}} = \text{diag}(1.810630, 1.095625, 3.005536)$ and $D_\chi = \text{diag}(-0.386595, 1.236002, 2.047300)$, and the charged spectra follow exactly from the commuting exponentials

$$D_e = e^{A_\ell}, D_d = e^{A_\ell + B_{\text{col}} - D_\chi}, D_u = e^{A_\ell + B_{\text{col}} + D_\chi}.$$

Proof. Direct evaluation; commutativity of $\{I, G, \Pi_3\}$ makes each exponential exact, and each coefficient is a frozen corpus constant. All three spectra reproduce to machine precision. ■

Significance and term roles — stated precisely. Since $\text{span}\{I, G, \Pi_3\}$ is the space of *all* diagonal 3×3 matrices, the decomposition itself is trivially available to any diagonal kernel; the content of this proposition is *coefficient naming*: each coefficient group is identified with a specific frozen corpus constant carrying a physical reading — $\ln 9$ the first-generation quark/lepton colour-triality separation, $\ln 20$ the resolved strange-support multiplicity, $\ln(3/20)$ the heavy coloured participation correction, with h and L the fold and localisation terms expressing these on the common generation basis. The universal rule may be written compactly as

$$\Lambda(q, t) = \exp(A_\ell + q B_{\text{col}} + q t D_\chi),$$

with $q = 0/1$ for lepton/quark and $t = \mp 1$ for down/up-type completion. Every piece is a *named* constant; none is thereby *derived* — the constants 9, 20, $77/3$, $3/20$, 12, $16/3$ remain exactly as underived as before, per the binding caveat below.

Binding caveat (dynamical weight). The colour-participation theorem fixes the *form* of the coloured block by symmetry but does not fix the coloured-versus-singlet *dynamical weight*: that contrast still requires an independently derived response operator. B_{col} is therefore an exact corpus assembly, not yet a microscopic derivation — its coefficients must ultimately return from the pre-confinement closure/record response through J_C and J_R , with confinement and RG matching acting only as downstream filters (§38.7).

S_{cl} and T_{cl} without the target matrices

Let V_q be the frozen conditional quark relative frame of §37.4. Define

$$\tilde{D}_u = V_q^\dagger e^{A_\ell + B_{\text{col}} + D_\chi} V_q, \tilde{D}_d = e^{A_\ell + B_{\text{col}} - D_\chi},$$

and

$$S_{\text{cl}} = (\tilde{D}_d + \tilde{D}_u)^{1/2}, T_{\text{cl}} = (\tilde{D}_d - \tilde{D}_u)^{1/2}.$$

These consume only the frozen operator ingredients A_ℓ , B_{col} , D_χ and V_q — never Y_u or Y_d as primitive matrix inputs. Numerically,

$$S_{\text{cl}} \simeq \begin{pmatrix} 3.41465 \times 10^{-4} & -1.354534 \times 10^{-3} + 1.62737 \times 10^{-4} i & 9.432371 \times 10^{-3} - 4.019884 \times 10^{-3} i \\ -1.354534 \times 10^{-3} - 1.62737 \times 10^{-4} i & 6.540998 \times 10^{-3} - 4.6899249 \times 10^{-2} - 4.26062 \times 10^{-4} i \\ 9.432371 \times 10^{-3} + 4.019884 \times 10^{-3} i & -4.6899249 \times 10^{-2} + 4.26062 \times 10^{-4} i & 1.2030236 \end{pmatrix},$$

$$T_{\text{cl}} \simeq \begin{pmatrix} -3.03327 \times 10^{-4} & 1.354534 \times 10^{-3} - 1.62737 \times 10^{-4} i & -9.432371 \times 10^{-3} + 4.019884 \times 10^{-3} i \\ 1.354534 \times 10^{-3} + 1.62737 \times 10^{-4} i & -5.778229 \times 10^{-3} & 4.6899249 \times 10^{-2} + 4.26062 \times 10^{-4} i \\ -9.432371 \times 10^{-3} - 4.019884 \times 10^{-3} i & 4.6899249 \times 10^{-2} - 4.26062 \times 10^{-4} i & -1.1635261 \end{pmatrix}.$$

Both $\tilde{D}_d$ and $\tilde{D}_u$ are positive definite, so $S_{\text{cl}} > 0$ (eigenvalues 4.4844×10^{-5} , 4.9139×10^{-3} , 1.2049473 , matching §38.5 identically), and the universal readout is the same canonical Schur family — now written in its complete classified form $C_Y^{(9)}(Z)$ with quark completion $R = T_{\text{cl}}^\dagger S_{\text{cl}}^{-1} T_{\text{cl}} + Z$, $Z \geq 0$, per Proposition 8 — assembled from named conditional corpus operators — a forward reassembly within the conditional corpus algebra, not a derivation from the substrate — rather than backwards from the answers as matrices. The fixed selectors recover $Y_e = D_e$, $Y_d = (S_{\text{cl}} + T_{\text{cl}})^{1/2}$, $Y_u = (S_{\text{cl}} - T_{\text{cl}})^{1/2}$ with maximum reconstruction residual below 3×10^{-16} .

Proposition 7 — Colour-Superselected Minimal-Rank Theorem

Assume $Y_e > 0$, $S_{\text{cl}} > 0$, Premise CSE, and the fixed completion-channel selectors. Then every positive universal operator reproducing the three charged matrices satisfies

$$\text{rank } C_Y \geq \text{rank } Y_e + \text{rank } S_{\text{cl}} = 6,$$

and the canonical Schur completion at $\mu = 0$ attains the bound. Therefore

$$\min \text{rank } C_Y = 6$$

within the exact colour-equivariant three-channel architecture.

Proof. CSE block-diagonalises C_Y into the lepton channel and the coloured quark plane. The lepton principal block is $Y_e > 0$, of rank three. The quark block contains the principal 3×3 sub-block $\langle Q | C_Y | Q \rangle = S_{\text{cl}} > 0$, so its rank is at least three. Ranks add across the CSE direct sum, giving $\text{rank} \geq 6$; the $\mu = 0$ completion factorises the quark block at rank exactly three and saturates the bound. ■

The nonzero eigenvalues of the minimum-rank operator $C_Y^{(9)}(Z = 0)$ are

$$(2.99644 \times 10^{-6}, 5.38123 \times 10^{-4}, 6.20645 \times 10^{-4}, 8.82558 \times 10^{-3}, 1.071270 \times 10^{-2}, 2.332740),$$

with the other three eigenvalues vanishing exactly. The inherited representation-level statements — that an invariant operator has no intertwiner between an irreducible triplet and a singlet complement, and that every contributing contraction must contain an exact singlet — apply to the *pre-contraction substrate* carrier, which is where Premise CSE lives; they do not settle the channel level, per the premise's own typing above.

This discharges the Scope caveat of Proposition 5: rank six is a genuine global minimum — conditional on Premise CSE, a named microscopic-realisation premise with a stated falsifier, whose truth awaits exactly the objects $(J_Ω, P_Y^{\text{micro}})$ that §38.7 shows are not yet identifiable.

Proposition 8 — Complete Positive-Completion Classification

For fixed $S_cl > 0$ and T_cl , every positive quark-channel completion has the form

$$R = T_cl^\dagger S_cl^{-1} T_cl + Z, Z \succcurlyeq 0,$$

and conversely every $Z \succcurlyeq 0$ yields a positive universal operator. Hence $Z = 0$ gives the unique minimum-rank completion, $Z > 0$ a strictly positive rank-nine completion, and $Z = \mu I$ the canonical one-parameter family of the numerical audit.

Proof. With $S_cl > 0$, the block matrix is positive iff its Schur complement $R - T_cl^\dagger S_cl^{-1} T_cl$ is positive semidefinite; parametrise that complement as Z . ■

The inverse problem is thereby closed *completely*: no unidentified Hermitian freedom remains anywhere in the charged construction.

Revised claim level

The paper may now honestly state:

The charged inverse universal-closure calculation is closed conditionally.

In full: on the frozen MCHY-1C corpus branch, the three charged singular spectra are generated exactly by one nonfactorising log-kernel built from the CLAS localisation operator, the exact half-lazy fold operator and the colour-centroid operator; the conditional quark frame then determines S_cl and T_cl without consuming Y_u or Y_d as primitive matrix inputs; lepton–quark block separation holds under the named colour-supersélection premise (Premise CSE); and the complete positive-completion classification returns a universal nine-dimensional carrier with globally minimal rank-six readout.

The decisive unclosed steps are now exactly three:

B_col from the pre-confinement closure/record response through J_C, J_R — rather than from the inherited ratios 9, 20, 77/3, 3/20; confinement and RG are downstream filters, not the source; **V_q from the blind projected Hessian** — rather than the conditional C_3 curvature

branch; $Y_f^{\text{mixed}} = Y_f^{\text{Hess}}$ through an explicit dressed substrate functional — the CY3' certificate itself.

The exact paired verdict is therefore:

Charged inverse universal-closure calculation: conditionally closed. Blind microscopic Yukawa calculation: still open.

The inherited premises remain the CLAS stack, the grid constants, the conditional frame V_q (with its underived z_C3), and the absent RG matching; the physical evaluation — J_Q , P_Y^{micro} , the dressed mixed vertex, CY3', the neutral sector, and the registered $\mathcal{H}_{15}$ five-slot test — is untouched by this closure. The forward status of the three unclosed steps is established in §38.7.

38.7 Forward discharge attempt and the Threefold Non-Identifiability Certificate

The three steps of §38.6 have now been attempted *forwards* — from the substrate objects rather than from the charged matrices. The result is sharp: §38.6 makes the problem maximally localised, but not yet determined. One part of the V_q mechanism is conditionally discharged; the full B_col , the full V_q and CY3' cannot be calculated from the presently specified corpus — and, more strongly, they are provably *non-identifiable* until the missing constitutive maps and dressed interaction are supplied. That is not "more work is needed"; it is an exact mathematical result about what the current inputs do and do not determine.

Step 1 — B_col : type separation and non-identifiability

B_col is a *logarithmic spectral contrast*, $B_col = \frac{1}{2}(\log D_u + \log D_d) - \log D_e$, whereas the closure and record Jacobians produce a *positive stiffness operator* $G_CR = 2J_C^\dagger W_C J_C + 2J_R^\dagger W_R J_R$. The corpus supplies no independently derived map $L : G_CR \mapsto B_col$ — it does not even state whether the relation is linear, logarithmic, determinant-based, entropic, spectral, or generated only after auxiliary relaxation. And with the constitutive dependence of $U_ij(\rho)$ and $\Phi_ij(\rho)$ absent (§34), J_C and J_R cannot be evaluated at all.

Proposition 9 — Colour-Operator Non-Identifiability. Even granting, more strongly than the corpus permits, that $G_CR = B_col$: for any operator X with $0 \preceq X \preceq B_col$, the choice $W_C = W_R = I$, $J_C = X^{(1/2)\wedge/2}$, $J_R = (B_col - X)^{(1/2)\wedge/2}$ satisfies $2J_C^\dagger J_C + 2J_R^\dagger J_R = X + (B_col - X) = B_col$. Infinitely many such X exist; the target therefore determines neither J_C , nor J_R , nor their division into closure and record response. ■

The extreme member $J_R = 0$, $J_C = (B_col/2)^{(1/2)\wedge/2} = \text{diag}(0.951480, 0.740144, 1.225874)$ reproduces B_col exactly — and is transparently reverse engineering, not substrate physics. This is the same obstruction the coloured-participation gate established abstractly: symmetry compresses the response to a relative sector weight but cannot determine that weight numerically without a canonical response map and response metric.

Correction to §38.6 — confinement is downstream. The strange-confinement theorem separates address $\rightarrow$ Yukawa spectral line $\rightarrow$ confinement channel $\rightarrow$ observable: confinement does not derive the strange eigenvalue, it controls how already-defined coloured spectral content appears through colour-singlet observables. The correct forward target is therefore **B_col from the pre-confinement closure/record response**, with confinement and RG matching as downstream filters. The missing microscopic statement is exactly

$$\mathbf{B_col} = \mathbf{L} [\mathbf{P_col} \mathbf{K_cl}^{(-1/2)} (2\mathbf{J_C}^\dagger \mathbf{W_C} \mathbf{J_C} + 2\mathbf{J_R}^\dagger \mathbf{W_R} \mathbf{J_R}) \mathbf{K_cl}^{(-1/2)} \mathbf{P_col}],$$

with $\mathbf{J_C}$, $\mathbf{J_R}$, $\mathbf{W_C}$, $\mathbf{W_R}$ and $\mathbf{L}$ all derived *before* $\mathbf{B_col}$ is opened.

Step 2 — $\mathbf{V_q}$: one conditional discharge, one non-identifiability

Proposition 10 — Conditional $\mathbf{C_3}$ Democracy Discharge. If the pre-readout functional is $\mathbf{C_3}$ -invariant, its background $\mathbf{C_3}$ -symmetric and $\mathbf{P_C}$ equivariant, then $[\mathbf{P_C} \mathbf{H_cl} \mathbf{P_C}, \mathbf{R_3}] = 0$, forcing equal Killing weights on the three branch supports, $w_{12} = w_{23} = w_{31} = b^2/3$, hence $|\zeta| = b/\sqrt{3}$. The minimal closure cell independently returns a democratic quadratic closure response $\propto 2\mathbf{I_3}$ on the Killing-orthogonal branch directions. The $\mathbf{C_3}$ support democracy of the conditional frame is thereby *discharged*, conditional on explicit functional-level $\mathbf{C_3}$ invariance and a symmetric vacuum. ■

But a block proportional to $\mathbf{I_3}$ generates none of the rest: it does not select the visible $2 \leftrightarrow 3$ branch, derive the generator entries a , b , c , produce a complex phase, select the CP orientation, derive $\mathbf{P_Y}$, or split the up and down left frames.

Proposition 11 — Relative-Frame Non-Identifiability. Every unitary $\mathbf{V}$ connected to the identity is $e^{(i\mathbf{H_V})}$ for some Hermitian $\mathbf{H_V}$ ($= -i \log \mathbf{V}$ on a declared branch). Therefore, until the projected block $\mathbf{P_W} \mathbf{H_cl} \mathbf{P_W}$ is independently evaluated, the statement " $\mathbf{V_q}$ is the unitary generated by $\mathbf{H_W}$ " selects nothing: any desired frame can be generated by a suitably chosen Hermitian block. ■

The exact forward requirements are: $\Omega_{\text{odd}} = \Omega_{\text{q}}$ with a , b , c returned from the projection; the $\mathbf{C_3}$ commutation on the role-even support; the democratic weights; the entry $(\mathbf{H_even}^{\wedge \mathbf{q}})_{23} = \lambda \varepsilon_{\text{W}} (b/\sqrt{3}) e^{(i\pi/3)}$; and $\|\Omega_{\text{mix}}\| \ll \|\Omega_{\text{even}}\| + \|\Omega_{\text{odd}}\|$. Only then does $\mathbf{V_q} = e^{(-\Omega_0 + \Omega_{\text{q}}/2)} e^{(+\Omega_0 + \Omega_{\text{q}}/2)}$ become a prediction rather than a reconstruction. The missing item is not further BCH algebra; it is the actual complex projected block $\mathbf{P_W} \mathbf{H_cl} \mathbf{P_W}$.

Step 3 — $\mathbf{CY3'}$: the arbitrary-residue theorem and the completeness fork

Proposition 12 — $\mathbf{CY3'}$ Non-Identifiability. Choose any matrix $\mathbf{R_f}$ and define a candidate bilinear $\mathcal{B_f}(\eta, \eta) = \mathcal{B_f0} + \eta \mathbf{K_fL}^{(1/2)} (\mathbf{Y_f}^\dagger \mathbf{H} + \mathbf{R_f}) \mathbf{K_fR}^{(1/2)}$ with no η -dependence. Then $\mathcal{D}_{\text{q}} \mathcal{B_f} = \partial_{\text{q}} \mathcal{B_f}$ and $\mathbf{Y_f}^\dagger \mathbf{M} = \mathbf{Y_f}^\dagger \mathbf{H} + \mathbf{R_f}$, so $\Delta_{\text{f}} \mathbf{CY3'} = \mathbf{R_f}$. Since $\mathbf{R_f}$ was arbitrary, every residue — including zero — is attainable by choice of the undetermined dressed bilinear. Before $\Gamma_{\text{sub}}^{\wedge \text{dressed}}$ is independently fixed, $\mathbf{CY3'}$ is neither proved nor falsified; defining the interaction so that $\mathbf{R_f} = 0$ makes the equality true by construction and destroys it as a test. ■

Corollary — the completeness fork. If the presently explicit bosonic functional plus the reconstructed free fermion action were regarded as *complete*, there is no radial left–right interaction, $[\mathcal{D}_Q \mathcal{B}_f]_\star = 0$, hence $Y_f^M = 0$ and $\Delta_f^{CY3'} = -Y_f^H \neq 0$: the Yukawa route would fail outright. If it is not complete — as this paper holds — the missing chiral fermion bilinear must be derived before $CY3'$ can be evaluated. The structural Fold-and-Record action currently supplies interface, closure and record dynamics, but not that bilinear. Either branch of the fork is decisive; neither is presently discharged.

The certificate

Theorem — Threefold Non-Identifiability Certificate. The current VERSF corpus uniquely fixes the *targets* of the three remaining charged-Yukawa calculations but does not determine their *values*. The positive residual-Hessian form admits infinitely many (J_C, J_R) realisations of any proposed colour operator (Proposition 9); the projected-Hamiltonian architecture admits every unitary relative frame until the projected blocks are evaluated (Proposition 11); and the unspecified dressed fermion bilinear admits every $CY3'$ residue (Proposition 12). Consequently none of the three targets may be promoted from conditional reconstruction to microscopic derivation without new constitutive information. ■

Step	Forward result
B_{col} target operator	Exact
B_{col} from J_C, J_R	Underdetermined (Proposition 9)
Confinement as source of B_{col}	Type-invalid — confinement is downstream
C_3 branch orbit	Proved
C_3 democratic support	Conditionally discharged (Proposition 10)
Complete blind V_q	Underdetermined until $P_W H_{cl} P_W$ is evaluated (Proposition 11)
$CY3'$ formula	Exact
$CY3'$ numerical residue	Arbitrary until $\Gamma_{sub}^{dressed}$ is fixed (Proposition 12)
Physical Yukawa calculation	Not closed

The irreducible new inputs are now exact:

$U_{ij}(\rho), \Phi_{ij}(\rho), W_C, W_R, P_Y^{micro}$ — for B_{col} and V_q ; $\Gamma_{sub}^{dressed} \supset \bar{\psi}_f L \mathcal{K}_f[\rho, \eta] \psi_f R + h.c.$, with $\mathcal{K}_f$ derived independently — for $CY3'$.

The honest summary: the keyhole is now completely specified, but the substrate metal from which the key must be forged is still absent from the action. Claiming the three discharges now would be precisely the backward insertion this gate was designed to prevent.

39. The staged executable programme

The calculation proceeds in frozen stages. No stage consumes flavour data, and no stage may be revisited after a later stage has been opened.

Stage A — Symmetric-diamond Hessian

Calculate the 64×64 real physical Hessian on $\ell^2(\{a,b,c,d\}) \otimes \mathbb{C}^8$ at the frozen diamond background, using the complex 32×32 representation only on the complex-linear sector per the Notation convention. Required explicit inputs: $U_{ij}(\rho)$, L_J , $\Phi_{ij}(\rho)$. Outputs: J_C , $D_\star$, J_R , and H_{bare} via Theorems 2–3.

Stage B — Physical reduction

Remove gauge and exact null directions; perform the auxiliary Schur complement of Theorem 4:

$$\mathcal{H}_{\text{red}} = H_{cc} - H_{c\eta} H_{\eta\eta}^{-1} H_{\eta c}, \hat{H}_{cl} = K_{cl}^{-1/2} \mathcal{H}_{\text{red}} K_{cl}^{-1/2}.$$

Stage C — Scalar-compatible profile space

Find the scalar-interface-compatible eigenprofiles of $\hat{H}_{cl}$ and **calculate** P_Y — its rank included — rather than selecting the rank in advance. Theorem 6 supplies the acceptance test: a returned rank of three or less cannot support four distinct full-rank spectra under the canonical discipline.

Stage D — Chiral address maps

Construct the eight raw address maps

$$A_{uL/R}, A_{dL/R}, A_{eL/R}, A_{\nu L/R}$$

from the frozen carrier census and exact representation projectors, then canonicalise by the polar construction of Theorem 9. Paper 2 supplies three terminal neutral carriers, but the neutral support and transversality still require direct evaluation.

Stage E — Blind matrices

The baseline branch of §36 constitutes the first pass through Stages A–E with frozen stand-in premises. The physical pass then calculates

$$Y_f = \mathcal{E}_{fL}^\dagger P_Y \hat{H}_{cl} P_Y \mathcal{E}_{fR},$$

together with the mixed-vertex side of Theorem 7 where the dressed functional permits, the CY3' residues of Theorem 8, and the neutral completion of §28. The output is accepted whatever it is.

40. Frozen input manifest

Before the target flavour dataset is accessible, freeze:

1. substrate event set or regulator sequence;
2. substrate functional;
3. coefficients of A_C , A_T , A_R ;
4. current operator J_ρ ;
5. record morphisms;
6. admissibility selector;
7. background-selection rule;
8. gauge and null-mode treatment;
9. physical kinetic metrics;
10. trace/radial extension;
11. auxiliary carrier;
12. scalar-compatible readout definition;
13. occupancy address maps;
14. chiral intertwiner selection rule (polar construction of Theorem 9);
15. neutral carrier ledger;
16. matching scale;
17. numerical truncation;
18. solver tolerances;
19. singular-value rank threshold;
20. degeneracy threshold;
21. phase convention;
22. second-variation convention (real, per Notation);
23. 1PI, momentum-configuration, scheme, residue and form-factor declarations for the mixed vertex (per §15);
24. publication repository and hash procedure.

Every object receives a content hash before the first flavour comparison. Any transport, closure or record coefficient not yet supplied by the corpus and fixed at this step is a **new VERSF premise**: it must be named, frozen and hashed before any output is seen, and its falsifier is inherited from F6 below.

41. Evaluation algorithm

Stage 0 — Data quarantine. Remove observed fermion masses, CKM and PMNS values, neutrino splittings and CP phases from the computational environment.

Stage 1 — Background. Solve $\delta\Gamma_{\text{sub}}^{\text{dressed}} = 0$ for the frozen admissible background.

Stage 2 — Physical tangent carrier. Construct gauge, null, radial, profile and auxiliary projectors.

Stage 3 — Parent Hessian. Evaluate the complete bosonic Hessian via Theorems 2–3 and the mixed fermion blocks.

Stage 4 — Canonicalisation. Calculate the physical kinetic metrics and their positive square roots.

Stage 5 — Auxiliary relaxation. Evaluate the exact Schur complement.

Stage 6 — Scalar-compatible profile support. Calculate P_Y , its rank and its relation to the radial field direction.

Stage 7 — Mixed vertex. Calculate $Y_f^{\text{mixed}} = K_{fL}^{(-1/2)} \mathcal{D}_Q \mathcal{B}_f K_{fR}^{(-1/2)}$.

Stage 8 — Hessian compression. Calculate $Y_f^{\text{Hess}} = \mathcal{E}_{fL}^\dagger P_Y \hat{H}_{\text{red}} P_Y \mathcal{E}_{fR}$.

Stage 9 — CY3' audit. Evaluate $\Delta_f^{\text{CY3'}}$.

Stage 10 — Neutral completion. Evaluate M_R and M_L^{irr} .

Stage 11 — Spectral extraction. Calculate singular and Takagi spectra, frames and invariants.

Stage 12 — Freeze publication. Publish all raw matrices and invariant ledgers.

Stage 13 — Only then open observations. No pre-output comparison is permitted.

42. Conditioning and numerical thresholds

Every inverse is accompanied by a conditioning report. For the auxiliary Hessian,

$$\kappa_\eta = \sigma_{\max}(H_{\eta\eta}) / \sigma_{\min}(H_{\eta\eta}).$$

For each sector,

$$\kappa_f = \sigma_{\max}(Y_f) / \sigma_{\min}(Y_f)$$

when full rank. A small singular value is declared zero only if the threshold was frozen before evaluation and is justified by rigorous numerical error bounds. Publish both raw floating-point outputs and certified intervals or precision-stability results. A matrix whose smallest singular value changes sign or rank under precision refinement is unresolved, not rank deficient.

43. Publication package

The blind package must contain: complete substrate input manifest; code and solver version; random seeds, if any; raw parent Hessian; kinetic metrics; gauge/null projectors; Schur-complement blocks; P_Y ; all raw address maps and canonical embeddings; Y_u, Y_d, Y_e, Y_v ; $M_L^{\text{irr}}, M_D, M_R$; singular and Takagi decompositions; left and right physical projectors; CKM and PMNS-compatible matrices; determinant and cycle phases; CP invariants;

transversality values; numerical uncertainty reports; CY3' matching residues; hashes and timestamps.

Publishing only masses or angles is insufficient. The operators themselves are the result. Any numerically converged exclusion bound quoted in the paper (in particular the trine $\lambda_{\min}^{\max}$) must ship with a formal certificate: the exact affine matrix family, the optimisation or SDP formulation, primal and dual solutions with a duality-gap bound, precision sufficient to certify the sign, and code and input hashes.

44. Non-insertion firewall

The calculation fails if any of the following is changed because the output looks wrong: substrate coefficient; background branch; readout rank; embedding representative outside unitary-commutant gauge; phase orientation; solver truncation; singular-value threshold; matching scale; treatment of an auxiliary mode; neutral branch; higher-order term; normalisation coefficient; second-variation convention.

A pre-registered alternative branch may be evaluated and published. An unregistered rescue branch is a fit. Existing VERSF mass-ratio and frame papers are comparison ledgers, not optimisation targets.

45. Current corpus evaluation

45.1 Available

- a candidate bosonic substrate functional;
- the $K = 7$, $\mathbb{C}^8$ closure-normalised carrier;
- admissibility decomposition;
- gradient-flow selector theorem;
- candidate refinement dynamics;
- one explicit nonuniform background (symmetric diamond);
- the exact radial potential Hessian (Theorem 2);
- the universal positive residual-Hessian formula (Theorem 3);
- the exact minimal flat-transport spectrum and its degeneracy verdict (Theorem 12);
- electroweak completion interface;
- scalar radial embedding and conditional effective normalisation;
- exact Schur-complement formalism;
- matter occupancy addresses;
- canonical polar embedding construction (Theorem 9);
- free relativistic fermion carrier;
- neutral carrier count and branch target;
- non-insertion and frame-invariant rulebook;
- the complete blind baseline branch MCHY-1(0): $H_{cl}^{(0)}$, $P_Y^{(0)}$, $\mathcal{E}_{fL/R}^{(0)}$ and the four matrices $Y_f^{(0)}$ of §36, with their common spectrum, relative frame and null CP invariants;

- exact address selectors S_{fX} with $S_{fX} \dagger S_{fX} = I_3$ from the occupancy map, with the sterile S_{NR} inherited separately from Paper 2 (§18);
- the six-dimensional pre-readout address carrier $P_{\Omega}^{(3 \times 2)} = P_{\text{gen}} \otimes P_{\text{comp}}$ (§13);
- the isotype-resolved closure-Hessian normal form, conditional on multiplicity-freeness (§34);
- one explicit conditional charged-lepton operator candidate Y_e^{CLAS} and pre-RG quark coefficient ledgers $D_u^{\text{grid}}, D_d^{\text{grid}}$, with declared Hermitian representatives $Y_u^{\text{C}}, Y_d^{\text{C}}$ (§37, under inherited premises);
- conditional relative-frame candidates $V_{\text{CKM}}^{\text{C}}$ and $U_{\text{PMNS}}^{\text{C}}$, quarantined with their underived inputs named (§37.4);
- the typed frozen neutral block B_N^{frozen} with its structural branch facts — rank $Y_v = 3$, symmetric invertible R_N , $M_L^{(0)} = 0$, type-I dominance, normal ordering, $\epsilon\tau$ attachment, $\theta_{23} > 45^\circ$, $J_{\text{lep}} < 0$ (§37.5);
- the matching-point audit branch $\text{MCHY-1C}(\alpha_0)$: full conditional charged matrices, an unmatched tree-level mass table, rank and hierarchy certificates, and the neutral texture as explicit linear forms in (m_1, m_2, m_3) (§37.7);
- the first execution of the inverse universal-closure consistency test: exact universal magnitude kernel with CLAS baseline and half-lazy fold, the analytic binary-carrier no-go, and the certified trine bound (§38.1–38.4);
- the second execution: the analytic rank-nine three-address realisation with doublet-locked selectors — universal for any positive targets under the declared convention (vacuity note, §38.5) — with exact recovery, an exactly constructed rank-six positive witness, and hierarchy-bound conditioning (Propositions 4–5);
- the third execution: the exact operator decomposition of the magnitude kernel, the forward assembly of S_{cl} and T_{cl} from the operator algebra and conditional frame, colour-superselected global rank-six minimality, and the complete positive-completion classification (§38.6, Premise CSE, Propositions 6–8) — closing the charged inverse calculation conditionally;
- the forward discharge attempt and Threefold Non-Identifiability Certificate: the conditional C_3 democracy discharge, and exact proofs that B_{col} , the blind V_q and the CY3' residue are non-identifiable from present inputs (§38.7, Propositions 9–12).

45.2 Not yet available

- the address-realisation bridge J_{Ω} , and a proof or refutation of its universality (Premise UAR);
- the isotype coefficients $\beta_A, \beta_L, \beta_{\perp}$ (equivalently the nine $\alpha_{X^{\text{C}}(R)}$), pending the constitutive ρ -dependence of U_{ij} and Φ_{ij} ;
- the full trace-plus-angular positive Hessian beyond the isotype normal form;
- the fermion-dressed substrate functional and the numerical mixed blocks $\mathcal{B}_f(q, \eta)$ and $\mathcal{D}_q \mathcal{B}_f$;
- the independently derived physical scalar-profile projector P_Y^{micro} , and hence P_Y^{phys} and its rank (bounded by nine via the charged address carrier; a rank-six witness exists, but neither minimality nor the physical rank is proved);
- the blind relative frames $U_{uL}, U_{dL}, U_{eL}, U_v$ from a substrate projection (the corpus CKM candidate's complex component z_{C3} remains underived);

- the RG matching operators $R_u(\mu_\star)$, $R_d(\mu_\star)$ that convert the grid ledgers into common-scale spectra;
- the CKM curvature phase and amplitude z_{C3} from the projected closure Hamiltonian, and the frozen Jarlskog orientation convention;
- the assumed lepton-frame premises (charged-lepton frame, canonical weak bridge, Frame–Takagi compatibility, active completeness) as derived results;
- the numerical neutral matrices Y_v , R_N , K_N and the heavy scale $M_\star$, and hence any numerical 6×6 block;
- both sides of every $CY3'$ residue $\Delta_f^{CY3'}$;
- the light masses m_1 , m_2 , m_3 and Majorana phases that would numericalise the returned neutral texture;
- the irreducible constitutive inputs of §38.7: $U_{ij}(\rho)$, $\Phi_{ij}(\rho)$, the response metrics W_C , W_R , P_Y^{micro} , and the derived chiral bilinear $\mathcal{K}_f[\rho, \eta]$ in $\Gamma_{\text{sub}}^{\text{dressed}}$ — proved necessary, not merely convenient;
- the projected block $P_W H_{cl} P_W$, without which every relative frame remains generable (Proposition 11);
- the extension from the three-address quotient to the full five-slot occupancy realisation — the registered $\mathcal{H}_{15} = \mathbb{C}^5 \otimes \mathbb{C}^3$ fixed-selector test with distinct L_L and e_R addresses — the derivation of the quotient π_{charged} from J_Ω , and the right-handed physical frames.

The unknown ledger is therefore no longer "everything numerical". It is concentrated in the bridge and its universality premise, three isotype coefficients (conditional on multiplicity-freeness), the physical readout rank, the dressed mixed vertex, the blind frames, the RG matching and the neutral completion.

45.3 Verdict

The four physical Yukawa matrices have not yet been calculated blind from one substrate. This paper does calculate and publish four explicit baseline compression matrices $Y_f^{(0)}$ from frozen stand-in premises (§36), and it assembles the conditional corpus branch MCHY-1C — an explicit Y_e^{CLAS} and conditional quark spectra — under inherited premises (§37). Neither stratum is identified with the physical Yukawa operators: the baseline lies in the wrong structural class ($Y_d^{(0)} = Y_e^{(0)}$, $Y_u^{(0)} = Y_v^{(0)}$, one common spectrum, no CP), and the corpus branch is conditional on its source gates' open premises, with RG matching, relative frames and the neutral sector absent. More precisely: the prior corpus fixes the complete address-level embedding selectors, the six-dimensional pre-readout address carrier, a conditionally exact three-isotype closure-Hessian normal form, one explicit conditional charged-lepton operator candidate and pre-RG up- and down-quark coefficient ledgers; the remaining microscopic work is concentrated in the address-realisation bridge and its universality premise, the isotype coefficients, the physical readout, the dressed mixed vertex, the blind relative frames, the RG matching and the neutral completion. Promoting either stratum to a physical claim, or repairing the baseline by importing corpus-branch structure retrospectively, is prohibited by the firewall. The current numerical closure test is therefore:

Not yet passed — baseline returned, structural class refuted.

This is not a failure of the construction theorem. It is the exact result of applying its closure standard. What has changed since the construction gate is that the first Hessian components are evaluated exactly, the first transport implementation is excluded, the first end-to-end blind branch is published, and the remaining unknowns are reduced to a named, finite list with sharply assigned jobs.

46. Falsification conditions

F1 — No dressed substrate vertex. If the substrate action cannot generate a mixed left–right radial vertex, the VERSF Yukawa route fails. The completeness fork of §38.7 makes this concrete: if the presently explicit bosonic-plus-free-fermion action is complete, then $Y_f^{\text{mixed}} = 0$, $\Delta_f^{\text{CY3}'} = -Y_f^{\text{H}} \neq 0$, and F1 fires.

F2 — Hessian–generator conflation required. If the construction requires one nonzero operator to be simultaneously positive and traceless, its typing is inconsistent.

F3 — Scalar-profile rank failure. If the derived scalar-compatible readout support has rank below the required Yukawa rank, the predicted full-rank sector fails.

F4 — Common-spectrum failure. If the derived readout support has rank exactly three under the canonical partial-isometry discipline, and the returned sectors nevertheless require distinct singular spectra, the leading construction fails by Theorem 6 and only the typed higher-order channels of §18 remain admissible.

F5 — CY3' mismatch. If $\Delta_f^{\text{CY3}'} \neq 0$ at declared leading order, the closure-Hessian identity fails for that sector.

F6 — Target-data dependence. If an embedding, Hessian term, projector, coefficient premise or branch is selected using a measured flavour quantity, the derivation fails.

F7 — Embedding non-uniqueness. If multiple inequivalent admissible unitary-commutant embedding orbits produce different matrices and no substrate selector distinguishes them, flavour is underdetermined.

F8 — Instability. If the physical reduced Hessian has a negative eigenvalue at the chosen background, the background fails.

F9 — Charged-sector rank failure. If a charged sector predicted to contain three nonzero physical masses returns $\sigma_{\min}(Y_f) = 0$ in the physical evaluation, the predicted branch fails.

F10 — Neutral-branch failure. If M_D , M_R , ρ_L or the hierarchy returns contradict Paper 2's selected branch, Paper 2's stand-in selection is overturned.

F11 — CP source absent. If the complete returned operator system admits a simultaneous real representative or generalised CP symmetry, any claimed nonzero physical CP phase fails.

Theorem 12 already shows the flat-transport sector alone lies inside this real class; a nontrivial CP return therefore requires holonomy, closure or record structure beyond it.

F12 — Residual degeneracy. If the fully evaluated H_{bare} inherits the diamond degeneracy — one common spectrum across all internal directions — after J_C , J_R and refinement weights are included, then the bosonic route to flavour structure fails and only the fermion-dressed mixed blocks remain. The baseline branch of §36 instantiates this degeneracy at the flat level and fixes the ledger against which the inclusion of each missing term is scored.

F13 — Cross-sector vacuum inconsistency. If different sectors require different scalar vacua, common electroweak completion fails.

F14 — Numerical instability. If matrix rank, phase or hierarchy changes under certified precision refinement, the corresponding claim remains unresolved.

47. The Microscopic Closure-Hessian and Full Yukawa Evaluation Theorem

Theorem 13 — MCHY-1

Let:

1. $\Gamma_{\text{sub}}^{\text{dressed}}$ be a frozen, gauge- and anomaly-admissible fermion-dressed VERSF substrate functional;
2. $\Xi_{\star}$ be its stable admissible vacuum;
3. $\mathcal{H}_{\text{red}}$ be the positive physical Schur-reduced closure Hessian, whose bare parent is $H_{\text{bare}} = 2\lambda_{\text{sub}} \rho_{\text{c}^2} P_{\text{rad}} + 2\Sigma_X J_X \dagger W_X J_X$ in the declared real second-variation convention;
4. K_{cl} , K_{fL} , K_{fR} be positive frozen kinetic metrics;
5. P_Y be the independently calculated scalar-compatible profile projector, with rank $P_Y > 3$ or a typed non-isometric or higher-order resolution of Theorem 6;
6. $\iota_{\text{fL/R}}$ be independently calculated admissible equivariant partial isometries, canonicalised by the polar construction of Theorem 9;
7. the completion–Hessian matching identity of Theorem 8 hold;
8. the neutral carrier and branch be those frozen in Paper 2;
9. all inputs and numerical choices be frozen before flavour observations are consulted.

Then the complete leading local, single-radial-insertion Yukawa operators are uniquely determined, up to physically trivial unitary-commutant and canonical-frame gauge, by

$$Y_{\text{f}_e} = \mathcal{E}_{\text{fL}} \dagger P_Y K_{\text{cl}}^{(-1/2)} \mathcal{H}_{\text{red}} K_{\text{cl}}^{(-1/2)} P_Y \mathcal{E}_{\text{fR}}, \text{ f = u, d, e, } \nu.$$

Equivalently, they are the exact relaxed mixed variations

$$\mathbf{Y}_{\mathbf{f}_e} = \mathbf{K}_{\mathbf{f}_L}^{(-1/2)} [\mathcal{D}_{\mathbf{q}} \delta^2 \Gamma_{\text{sub}}^{\text{dressed}} / \delta \bar{\psi}_{\mathbf{f}_L} \delta \psi_{\mathbf{f}_R}]_{\star} \mathbf{K}_{\mathbf{f}_R}^{(-1/2)}.$$

The equality of these two expressions is the microscopic CY3' certificate.

The charged physical magnitudes are the singular values of $\mathbf{Y}_u, \mathbf{Y}_d, \mathbf{Y}_e$. The quark mixing matrix is $\mathbf{V}_{\text{CKM}} = \mathbf{U}_u \mathbf{L}^\dagger \mathbf{U}_d \mathbf{L}$. The neutral completion is

$$\mathbf{B}_N = \left(\begin{array}{c} \mathbf{M}_L^{\text{irr}} (v_c / \sqrt{2}) \mathbf{Y}_\nu \\ (v_c / \sqrt{2}) \mathbf{Y}_\nu^T \mathbf{M}_R \end{array} \right),$$

with physical masses given by its Takagi singular values and PMNS-compatible mixing obtained from the relative charged-lepton and active-neutrino frames.

Every rank, hierarchy, texture and phase is thereby an output of the frozen substrate objects. No independent Yukawa entry, texture coefficient, hierarchy factor, mixing angle or physical phase remains admissible.

Current-corpus clause

The theorem's hypotheses are not yet numerically discharged because the present corpus does not provide the complete dressed mixed vertex, the transport-normalised Jacobians $\mathbf{J}_C, \mathbf{D}_{\star}, \mathbf{J}_R$ beyond the excluded flat case, the independently derived scalar-profile projector or the numerical chiral address maps. The baseline branch of §36 discharges a stand-in version of hypotheses 3–6 with frozen premises and returns four matrices in the wrong structural class, which is the expected outcome given Theorem 12 and the baseline isospectrality of Proposition 1. Therefore MCHY-1 is a closed conditional evaluation theorem, with its first Hessian components evaluated exactly and its first blind branch executed, but not yet a passed physical four-matrix evaluation.

■

48. MCHY-1 Closure Certificate

Closed

- correct parent object identified;
- Hessian and transport generator separated (Theorem 1);
- second-variation convention declared; factor-of-two ambiguity removed;
- exact radial potential Hessian derived (Theorem 2);
- universal positive residual-Hessian formula derived (Theorem 3);
- full physical carrier typed; trace/radial sector restored;
- exact auxiliary relaxation derived (Theorem 4);
- scalar-field/readout-rank distinction proved (Theorem 5);

- common-spectrum no-go proved; rank $P_Y > 3$ requirement established (Theorem 6);
- exact mixed-variation Yukawa formula derived (Theorem 7);
- CY3' converted into a direct equality test (Theorem 8);
- canonical polar embedding construction derived (Theorem 9);
- embedding uniqueness criterion derived (Theorem 10);
- transversality certificate established (Theorem 11);
- minimal flat transport evaluated exactly and excluded as sole flavour carrier (Theorem 12);
- first blind baseline branch executed end to end; four explicit matrices, spectrum, relative frame and CP invariants published (§36);
- baseline signature-conjugate isospectrality proved: on the evaluated branch, interface conjugation alone cannot split sector spectra (Proposition 1);
- raw embeddings coordinatised through exact address selectors, $A_{fX} = J_{\Omega} S_{fX}$, with universality of the bridge registered as Premise UAR (§18);
- six-dimensional pre-readout address carrier established as completion subcarrier, and superseded by the nine-dimensional charged carrier $P_{\Omega}^{(3 \times 3)}$ with rank $P_Y^{\text{phys}} \leq 9$ (§13, §38);
- closure Hessian reduced to the three-isotype normal form, conditional on multiplicity-freeness (§34);
- conditional corpus branch MCHY-1C assembled, numerically verified and quarantined, now including declared charged-sector representatives, relative-frame candidates and the typed neutral block (§37);
- CY3' residues certified unevaluated — not zero — with the true-by-construction failure mode named (§37.6);
- matching-point audit branch MCHY-1C(α_0) assembled, numerically verified to displayed precision, and firewall-typed (§37.7);
- neutral texture returned as explicit linear forms (§37.7);
- inverse universal-closure consistency test executed in constrained form: magnitude-rule pass exact, binary rank-six no-go proved analytically and upgraded to frame invariance from singular values alone (Propositions 3, 3'), trine bound numerically converged at -0.533 across three methods and all six assignments for the declared representatives, with the pole-pinned mixed variant separately bounded at -0.401 (§38.1–38.4);
- exact half-lazy identity proved: fold increments $\frac{1}{2} \ln(77/3)$ and $\frac{1}{4} \ln(77/3)$ (Proposition 2);
- rank-nine three-address realisation proved: positive universal operator, exact recovery, doublet lock respected (Proposition 4);
- a positive rank-six universal witness constructed exactly inside the nine-dimensional carrier (Proposition 5), upgraded to the global minimum under colour superselection (Proposition 7);
- the diagonal magnitude kernel identified exactly with named corpus coefficients (Proposition 6);
- S_{cl} and T_{cl} assembled from named conditional corpus operators and the conditional frame, without the target matrices — forward reassembly within the corpus algebra, not substrate derivation (§38.6);
- the complete positive-completion classification closed: $R = T^\dagger S^{-1} T + Z$, $Z \geq 0$ (Proposition 8);
- the charged inverse universal-closure calculation closed conditionally (§38.6);

- the Threefold Non-Identifiability Certificate proved: targets fixed, values undetermined (§38.7, Propositions 9, 11, 12);
- C_3 support democracy conditionally discharged from functional C_3 invariance and a symmetric vacuum (Proposition 10);
- the completeness fork stated: a complete present action forces $Y_f^{\text{mixed}} = 0$ and fails the route (§38.7);
- conditioning certified at $1.39\times$ the unavoidable hierarchy bound;
- charged and neutral outputs typed;
- CKM and PMNS relative-frame extraction fixed;
- phase and determinant audit fixed;
- staged executable programme fixed (Stages A–E);
- blind numerical protocol fixed;
- all principal smuggling channels closed.

Conditional

- uniqueness of the dressed substrate functional;
- completeness of $A_C + A_T + A_R$;
- basin uniqueness;
- continuum and kinetic matching;
- unique admissible embedding orbit;
- scalar-compatible profile factorisation and the $\mathbb{C}^3 \otimes \mathbb{C}^2$ candidate carrier (frozen as ansatz in the baseline branch; not yet independently derived);
- the baseline coefficient premises $\kappa_T = 1$, $r = (32/63)^2$ (HR-1 inheritance) and $\theta = \pi/6$ (illustrative);
- Premise UAR: one sector-independent equivariant address-realisation bridge exists;
- the bridge-factorisation premise: J_Ω preserves the generation-by-completion structure of $P_\Omega^{(3\times 2)}$, and $\text{rank } P_Y^{\text{phys}} = 6$;
- multiplicity-freeness of the diamond representation at the retained order;
- the MCHY-1C branch: the candidate Y_e^{CLAS} , the grid ledgers D_d^{grid} , D_u^{grid} , the Hermitian representatives, the frame candidates V_{CKM}^C and U_{PMNS}^C , and the typed B_N^{frozen} , inherited under the source gates' open localisation, one-twelfth, tau-channel, interface-projection, geometric-coupling, scheme/RG, curvature (z_{C3}) and lepton-frame premises;
- the Route-A scalar-branch scale $v_{cl} = 243.722807638$ GeV used only in typing B_N^{frozen} ;
- the MCHY-1C(α_0) anchors $\alpha(0) = 1/137.035999084$ and $v_\star = 246.21965$ GeV: measured non-flavour inputs, admissible for the audit branch, inadmissible in the blind physical evaluation;
- the $77/3$ charm hinge and the $e^{(163)/4}$ bottom step, the branch's most exposed inherited factors (the hinge now also carries the exact half-lazy identity of Proposition 2);
- the $\mathbb{C}^3 \otimes \mathbb{C}^3$ three-address carrier: realised and positive on the corpus targets, but reconstructed rather than substrate-derived, and a quotient of the full five-slot occupancy space via the explicit map π_{charged} , itself not yet derived from J_Ω ;
- the physical readout rank: bounded by nine, globally minimised at six under Premise CSE, not yet calculated blind;

- Premise CSE as a named premise about the microscopic address realisation (the channel level does not force it: $\mathbf{3} \otimes \mathbf{3} \supset \mathbf{1}$); falsifier: a colour-singlet component of the quark readout support overlapping the lepton channel reverts Proposition 7 to the rank-greater-than-three bound;
- functional-level C_3 invariance and the symmetric vacuum, on which the democracy discharge of Proposition 10 rests;
- leading-order closure-side isotropy;
- canonical unit inheritance;
- Paper 2 neutral branch.

Open numerical objects

J_Ω (with Premise UAR), $\beta_A, \beta_L, \beta_\perp$ (equivalently the nine $\alpha_X(R)$, conditional on multiplicity-freeness),

and through them

$\mathcal{H}_{\text{red}}, P_{Y^{\text{phys}}}$ and its rank, $\mathcal{E}_{uL/R}, \mathcal{E}_{dL/R}, \mathcal{E}_{eL/R}, \mathcal{E}_{vL/N_R}$,

together with the dressed mixed blocks $\mathcal{D}_Q \mathcal{B}_f$, the blind relative frames $U_{uL}, U_{dL}, U_{eL}, U_{vL}$, the RG matching operators R_u, R_d , the curvature input z_{C3} , the neutral matrices Y_v, R_N, K_N and scale $M_\star$, every CY3' residue $\Delta_f^{\text{CY3'}}$, and therefore the physical

Y_u, Y_d, Y_e, Y_v .

The conditional $Y_e^{\text{CLAS}}, D_u^{\text{grid}}$ and D_d^{grid} of §37 stand outside this list as quarantined comparison ledgers, not as discharged outputs.

Numerical closure verdict

Charged inverse universal-closure calculation: conditionally closed (§38.6). Blind microscopic Yukawa calculation: open — the blind baseline branch is published and structurally refuted; the conditional corpus branches are assembled and quarantined; no blind physical four-matrix pass is claimed.

49. Conclusion

The VERSF Yukawa programme began by proposing that fermion masses are readings of closure stiffness. The construction theorem then proved that the leading matrices must have the form

$$Y_f = \mathcal{E}_{fL}^\dagger P_Y \hat{H}_{cl} P_Y \mathcal{E}_{fR}.$$

The present paper moves the programme from the architecture of that statement to its microscopic edge — and across the first stretch of the calculation itself.

It identifies the dressed substrate functional that must be differentiated; separates the positive stiffness operator from the traceless transport generator; fixes the second-variation convention; derives the exact radial Hessian and the universal positive residual-Hessian formula, so that the entire bosonic Hessian is now reduced to three named Jacobians; evaluates the minimal flat-transport implementation exactly on the symmetric diamond and proves it too degenerate to carry flavour; restores the radial and scalar-profile carrier; relaxes all orthogonal modes exactly; derives the mixed third variation defining the Yukawa vertex; proves that the readout support must exceed rank three; supplies the canonical polar construction that removes every normalisation ambiguity from the embeddings; states every invariant output required for a blind four-sector evaluation; executes the first end-to-end blind branch, publishing four explicit matrices whose failure is as informative as any success could have been at this stage; and assembles the conditional corpus branch, so that the physical evaluation, when it returns, meets a pre-registered comparison ledger rather than an empty page.

The central microscopic identity is now

$$\mathbf{K}_{fL}^{(-1/2)} [\mathcal{D}_q \delta^2 \Gamma_{\text{sub}}^{\text{dressed}} / \delta \bar{\psi}_{fL} \delta \psi_{fR}] \star \mathbf{K}_{fR}^{(-1/2)} = \mathcal{E}_{fL}^\dagger \mathbf{P}_Y \hat{\mathbf{H}}_{\text{red}} \mathbf{P}_Y \mathcal{E}_{fR}.$$

That equality is the point at which the closure Hessian becomes the Yukawa operator rather than merely resembling one.

The paper also records the honest boundary of the present corpus. The bosonic substrate architecture is now concrete enough not merely to define the calculation but to run it once: the potential block is exact, the flat transport block is exact and excluded, and the baseline branch returns $Y_d^{(0)} = Y_e^{(0)}$ and $Y_u^{(0)} = Y_v^{(0)}$ with one common spectrum $\simeq (5.818, 3.011, 2.964)$, a single 2.24° real rotation and no CP phase. The flavour information is thereby localised, with proofs rather than hopes, to J_C, J_R , refinement-dependent transport weights, holonomy or complex record structure, and the fermion-dressed mixed blocks — each with a named job the baseline cannot do. The embedding problem has contracted: the address selectors are exact and corpus-fixed, so the eight raw embeddings are coordinatised by one bridge J_Ω — whose genuine universality is the named Premise UAR, an open theorem target — after which the metric polar construction fixes them without further choice. The Hessian problem has contracted likewise, to three isotype coefficients atop the exactly known radial term, conditional on multiplicity-freeness. And the corpus already holds a full conditional comparison ledger — the charged-lepton candidate, the quark grid ledgers with declared Hermitian representatives, the CKM and PMNS frame candidates with their underived inputs named, the typed neutral block, and now the matching-point audit branch with its mass table and neutral texture — against which the physical returns will be scored after RG matching, with every CY3' residue standing open until both of its sides exist. The inverse universal-closure consistency test has now delivered decisive returns before the dressed functional exists: the single-universal-rule architecture survives at the level of magnitudes — with its lepton baseline and quark fold matching the independently proposed CLAS and half-lazy mechanisms identically — the simplest binary

readout carrier is analytically excluded, and the three-address carrier is shown sufficient to host the conditional charged operators — a construction available for any positive targets under the declared convention (§38.5), its evidential weight residing in the exclusions it survives rather than the recovery it guarantees — with an exactly constructed rank-six witness sitting within $1.39\times$ the unavoidable hierarchy conditioning. The three-channel carrier is sufficient to host the conditional charged operators and admits a completely classified positive completion — assembled from named conditional corpus operators, block-separated under the named colour-superselection premise (CSE), and globally rank-minimal conditional on exactly that premise. This establishes a viable target architecture for the microscopic calculation, not evidence that the microscopic substrate selects that architecture or operator. The charged inverse calculation is closed conditionally, and the forward attempt then proves exactly why nothing more can be closed today: the remaining targets are uniquely fixed but provably non-identifiable — infinitely many microscopic realisations of the colour operator, every possible relative frame, every possible CY3' residue — until the constitutive maps, response metrics, microscopic readout and chiral fermion bilinear enter the action. One genuine forward result did land: the C_3 democracy of the quark frame now descends from functional symmetry rather than assumption. What separates this gate from the physical result is therefore no longer a list of calculations to run but a list of objects the action must first contain. The keyhole is completely specified; the metal for the key is not yet in the theory — and CY3' waits at the end to judge whether the closure Hessian and the completed architecture are one object or two. No legitimate route presently converts the remaining missing objects — the bridge and its premise, the isotype coefficients, the physical readout, the dressed vertex, the blind frames, the RG matching, the neutral completion — into four blind *physical* matrices.

The next action is therefore one direct calculation:

fix the constitutive ρ -dependence of U_{ij} and Φ_{ij} as frozen premises; evaluate J_C , $D_\star$, J_R — hence the isotype coefficients β_A , β_L , $\beta_\perp$ and H_{bare} — on the symmetric diamond; reduce, canonicalise and extract P_Y^{micro} , then P_Y^{phys} , testing the pre-readout factorisation $P_{\Omega^{(3\times 2)}}$ and calculating its rank; derive the universal bridge J_Ω and canonicalise the embeddings $\mathcal{E}_{fX} = B_{fX}(B_{fX}^\dagger B_{fX})^{(-1/2)}$, $B_{fX} = K_{cl}^{(1/2)} J_\Omega S_{fX} K_{fX}^{(-1/2)}$; first, supply the irreducible constitutive inputs of §38.7 — $U_{ij}(\rho)$, $\Phi_{ij}(\rho)$, W_C , W_R , P_Y^{micro} , and the derived chiral bilinear $\mathcal{K}_f[\rho, \eta]$ — since without them the three remaining steps are proved non-identifiable; alongside, run the $\mathcal{H}_{15}$ five-slot occupancy test with distinct L_L and e_R addresses; then calculate Y_u , Y_d , Y_e , Y_v ; score each against the frozen baseline of §36, the quarantined corpus ledger of §37 and the inverse-test constraints of §38; and publish before opening the flavour data.

If the returned operators reproduce the observed spectra and relative frames, VERSF will have crossed from a structural theory of flavour into an actual microscopic derivation. If they do not, the framework will fail at a precisely identified Jacobian, projector, embedding or matching surface.

Either outcome is progress, and the symmetric-diamond positive Hessian — not a guessed Yukawa matrix — is the correct first target.

The correct status of this gate is therefore Paper 3, Part I: a complete microscopic reduction, a first blind baseline evaluation, and a numerical non-closure theorem. Part II is the physical evaluation itself.