The Projected Weak-Doublet Closure Hamiltonian in VERSF

Deriving CKM Curvature and the PMNS Weak-Commitment Kernel from the H_cl Sector Projection

Keith Taylor *VERSF Theoretical Physics Programme — Standard Model Flavour Series Working manuscript*

Successor to *The CKM Curvature Residue in VERSF*, *The Weak-Commitment Closure Kernel in VERSF*, *The Weak-Commitment Neutrino Operator in VERSF*, *The Electroweak Flavour-Frame Operator in VERSF*, and *The Yukawa Operator from Completion-Channel Misalignment in VERSF*. Anchored on the substrate-dynamics result identifying H_cl as the su(8)-restricted admissibility Hessian.

Summary for the General Reader

Every kind of matter particle comes in three "generations" — three progressively heavier copies of the same basic thing. Quarks have three; so do the neutrinos. The generations are not perfectly separate: they blend into one another a little, and physicists call this blending *mixing*. It is not a technicality. Mixing controls real effects in the world, including the faint imbalance between matter and antimatter that let any matter at all survive the early universe.

Here is the puzzle this paper is about. **Quarks barely blend** — their mixing is tiny. **Neutrinos blend enormously** — their mixing is huge. Two kinds of particle built on the very same three-generation pattern end up at opposite extremes. The usual response is to tell two separate stories, one for each. This paper argues they are really one story told in two settings.

The central idea is that the rule governing this blending is not something invented just to explain it. The VERSF programme already has a single "master rulebook" describing how its underlying fabric updates itself. The blending rule for particle generations turns out to be one **slice** of that rulebook — a shadow it casts on the part of the theory that deals with families. Nothing new has to be bolted on; you look at a corner of what is already there.

That slice explains the two extremes through a single difference: **how firmly each particle is anchored.**

- A quark comes with two partners — an "up-type" and a "down-type" member of a pair — and both are held firmly in place. Because both are anchored, they can differ only by a tiny twist relative to each other, and a tiny twist means tiny blending. That is why quark mixing is small. A faint leftover *warp* in the twist they share then accounts for the fine details, including the matter–antimatter asymmetry.

- A neutrino, by contrast, is barely anchored at all. The stiffness that normally stops generations from blending goes slack, the resistance collapses, and the blending becomes large. And — this is the subtle part — the *pattern* of that blending is set by the shape of what is left over, not by the neutrino's tiny mass. So the famous smallness of neutrino mass does not shrink the mixing.

The same rulebook therefore gives small quark mixing in the "firmly anchored" setting and large neutrino mixing in the "barely anchored" setting. That is the claim in one line: **one rule, two regimes.**

What this paper does *not* do is finish the job — it sharpens the problem rather than closing it. A handful of specific numbers (a particular angle on the quark side; two ratios and a small "leakage" size on the neutrino side) still have to be *derived* from the master rulebook rather than chosen because they happen to fit. The paper is also open about one honest gap: for neutrinos it can predict *how much* the blending tips away from the perfectly balanced case, but not yet *which way* it tips — that direction hinges on a detail the theory has not pinned down. The value of the work is that these are now sharp, testable questions about a single object, instead of loose guesses about two.

Abstract

The remaining CKM problem has been reduced to a weak-doublet curvature residue $\kappa = [\Omega_0, \Omega_q]$, and the remaining PMNS problem to a weak-commitment neutrino kernel M_ν whose eigenframe stays finite as the commitment scale $\varepsilon \to 0$. This paper anchors both in the substrate-dynamics result: H_{cl} is the su(8)-restricted Hessian of the substrate free-energy functional at the admissible manifold and enters the minimal admissible refinement update through $\exp(i\cdot\varepsilon\cdot H_{cl})$. The weak-doublet closure Hamiltonian is therefore defined as the electroweak / flavour / readout projection

$$H_W = P_W \cdot H_{cl} \cdot P_W, \text{ with frame generator } \Omega_W = i\cdot\varepsilon_W\cdot H_W,$$

and role decomposition

$$\Omega_W = \Omega_{even} \otimes I + \Omega_{odd} \otimes \tau_3 + \Omega_{mix}.$$

In the committed quark regime, Ω_{odd} generates the leading CKM split and Ω_{even} generates the common-mode curvature, with BCH residue

$$\Delta_{CKM} = -\tfrac{1}{2}[\Omega_0, \Omega_q] + (1/6)[\Omega_0, [\Omega_0, \Omega_q]] + \cdots.$$

The minimal C_3 target is the role-even 2↔3 curvature $z_{C3} = (b/\sqrt{3})\cdot e^{(5\pi i/6)}$, whose commutator with the Cabibbo 1↔2 split gives $[\Omega_0, \Omega_q]_{13} = -a\cdot z$ (exact) and hence the missing 1↔3 triangle residue. In the weak-commitment neutrino regime the projection gives $T_\nu = D_0\cdot I + \varepsilon\cdot M_\nu$, so the neutrino eigenframe is exactly the eigenframe of M_ν for every $\varepsilon > 0$. The same

root-normalised closure-sharing principle that motivates $|z| = b/\sqrt{3}$ gives the solar ratio $\rho_\odot = \sqrt{3/2}$; weak neutrality gives $\delta = 0 + O(\beta^2)$; and a shared half-transport phase grammar — $h^2 = C\cdot O$, the square root of complement reversal against a sector orientation — gives both $\theta_q = 5\pi/6$ and $\psi_{wc} = 3\pi/4$. The pair ratio $B/A = 6/5$ is tied to the five-dimensional residual support space $R = C \oplus G$ but remains conditional on a one-extra-continuation rule; the leakage $\beta = \sqrt{3}/20$ remains the weakest target. One subtlety made explicit here: the atmospheric **octant** and the leptonic **CP sign** are not yet predicted — they are fixed by which electron-attachment channel carries the leading μ–τ breaking (premise WF-2), while the *magnitude* $|\theta_{23} - 45°| \approx 3.4°$ is robust. The result is a projected-Hamiltonian formulation with clear derive-or-reject tests, not a completed derivation.

Contents

---

0. Predictive-Content Ledger

This paper addresses one target and isolates one premise.

WF-1 — derive or reject a projected weak-doublet closure Hamiltonian whose committed quark limit gives the CKM common-mode curvature residue and whose weak-commitment neutrino limit gives the PMNS closure kernel.

WF-2 — derive, from closure geometry, *which electron-attachment channel* carries the leading μ–τ breaking. This binary choice fixes the atmospheric octant and the sign of δ_CP, which are otherwise undetermined (§8.6, §10.2).

Object	Construction	Status
Substrate generator	$H_{cl} = G\|_{su(8)}$	inherited; foundational projection target
Weak-doublet projection	$H_W = P_W \cdot H_{cl} \cdot P_W$	central definition
Frame generator	$\Omega_W = i\cdot\varepsilon_W\cdot H_W$	anti-Hermitian frame readout
Role decomposition	$\Omega_{even} \otimes I + \Omega_{odd} \otimes \tau_3 + \Omega_{mix}$	matches EW frame paper
Charge-blindness	$[H_{readout}, \chi] = 0$	inherited selection rule (Yukawa hierarchy)
Quark committed limit	$\gamma_u \approx \gamma_d \approx 1$	small CKM split + small curvature
CKM role-odd split	Ω_q	inherited leading generator
CKM common curvature	$\Omega_0(z_{C3})$	target output of role-even projection
C_3 curvature target	$z_{C3} = (b/\sqrt{3})\cdot e^{(5\pi i/6)}$	candidate; amplitude stronger than phase
BCH residue	$-½[\Omega_0, \Omega_q] + \cdots$	exact algebra (verified)
Neutrino weak limit	$T_\nu = D_0\cdot I + \varepsilon\cdot M_\nu$	exact frame theorem
PMNS solar ratio	$\rho_\odot = \sqrt{(3/2)}$	conditionally derived from root-count norm
PMNS pair ratio	$B/A = 6/5$	strengthened via residual support dim 5; still conditional
PMNS diagonal breaking	$\delta = 0 + O(\beta^2)$	weak-neutrality selection rule
PMNS phase	$\psi = 3\pi/4$	half-transport phase grammar
PMNS leakage	$\beta = \sqrt{3}/20$	weakest target; needs support-trace proof
Octant / CP-sign channel	e–μ vs e–τ leading breaking	**WF-2, open**

Three levels stay separated throughout: (i) **exact algebra** — the BCH residue and the ε-independence of the PMNS eigenframe; (ii) the **projection proposal** — that P_W·H_cl·P_W produces both structures; (iii) the **physical derivation still owed** — that the projected blocks *force* the claimed ratios and phases rather than merely admitting them.

Discipline. A projected Hamiltonian that reproduces CKM and PMNS is not enough. It must explain *why* those exact ratios and phases are selected — and, for the octant, *which* breaking channel closure geometry forces.

---

1. The Projection Principle

The weak-doublet flavour Hamiltonian is not a new primitive. The substrate-dynamics result supplies a closure Hamiltonian H_{cl} as the su(8)-restricted Hessian of the substrate free-energy functional at the admissible manifold, and the minimal admissible refinement update uses $\exp(i \cdot \varepsilon_m \cdot H_{cl})$. The weak-doublet flavour operator is therefore defined as a **projected sector** of that generator:

$$H_W = P_W \cdot H_{cl} \cdot P_W, \quad \Omega_W = i \cdot \varepsilon_W \cdot H_W,$$

where P_W projects onto the electroweak, generation, weak-role, and mass-readout subspace relevant to left-handed doublets, and Ω_W is anti-Hermitian because the sector-frame matrices are unitary.

This fixes the shape of the whole programme. The task is not "write a Hamiltonian that gives CKM and PMNS." It is: **project the already-defined closure Hamiltonian onto the weak-doublet flavour subspace and test what it forces.**

Status. The projector P_W is not yet computed from substrate microphysics. This paper defines the required projection and states the output tests. The derivation succeeds only when $P_W \cdot H_{cl} \cdot P_W$ is calculated and shown to force the claimed entries.

---

2. Source Architecture Inherited from the Foundational Papers

2.1 The closure Hamiltonian H_cl

The substrate-dynamics work identifies H_{cl} as the su(8)-sector restriction of the Hessian G of the free-energy functional F at the admissible manifold. This gives the flavour programme an existing finite generator to project. The same work cautions that the *unitary* use of G is a refinement-analysis layer, while the physical substrate dynamics are gradient-flow dynamics; H_W therefore uses H_{cl} as a frame-transport generator, not as a replacement for the dissipative substrate flow.

2.2 The Yukawa operator as closure-depth overlap

The structural-origin Yukawa work identifies Yukawa matrix elements as closure-depth overlap amplitudes, so the projected Hamiltonian acts not on arbitrary masses but on closure-depth readout operators whose eigenvectors become sector frames. The left Yukawa-square operator remains the charged-current object,

$$Y_S \cdot Y_S^\dagger = U_S \cdot \Lambda_S^2 \cdot U_S^\dagger,$$

and CKM and PMNS are relative left frames.

2.3 Charge-blindness of mass readout

The Yukawa hierarchy theorem proves the mass-amplification operator is a function of the completion-density operator and commutes with charge operators. The projected Hamiltonian may distinguish generation/readout depth, but it must not change electric charge or gauge representation — a constraint, not a decoration.

2.4 Closure-norm condensation and weak commitment

The closure-norm condensation work reads electroweak symmetry breaking as closure-norm condensation, with the Higgs-like mode as the radial amplitude mode of the closure-norm field, and states that neutrinos couple weakly because their PFD structure has minimal overlap with the closure-norm direction. This is the physical basis for the lepton asymmetry: the charged lepton is condensate-anchored; the neutrino is weakly committed.

2.5 Residual support dimensions

The flavour-mixing work derives a five-dimensional residual space $R = C \oplus G$ with $\dim C = 3$, $\dim G = 2$, and uses it to obtain $\eta = 3/5$, $\chi = 2/5$. The same five-support unit becomes the natural denominator for the PMNS pair-ratio count in §8.2.

---

3. Definition of the Projected Weak-Doublet Closure Hamiltonian

The weak-doublet flavour subspace is the sector of the substrate space surviving three projections:

- generation-completion projection P_C onto the three generation-depth directions;
- weak-role projection P_role onto the two components of a left-handed weak doublet;
- mass/readout projection P_Y onto the closure-norm / Yukawa readout channel.

$P\_W = P\_Y \cdot P\_role \cdot P\_C$, $H\_W = P\_W \cdot H\_cl \cdot P\_W$.

On the weak-role space introduce τ_3, with eigenvalues ±1 distinguishing the two doublet readings. The most general role-decomposed leading generator is

$\Omega\_W = \Omega\_even \otimes I + \Omega\_odd \otimes \tau_3 + \Omega\_mix$,

where Ω_even is common-mode transport seen by both members, Ω_odd is the role-odd split that survives into CKM (or the lepton weak-commitment frame), and Ω_mix changes weak role and must be gapped from the mass-readout basis for a clean sector-frame construction.

4. Selection Rules: Charge-Blindness, Weak Role, Commitment Class

4.1 Charge-blindness

The mass/readout part of H_W may depend on refinement depth and closure support, but it must commute with gauge-charge operators χ:

$$[H_W^{readout}, \chi] = 0.$$

If this failed, the same operation that changes mass across generations could shift charge eigenvalues, producing generation-dependent electric charge. The Yukawa hierarchy theorem explicitly rejects that failure mode.

4.2 Commitment class

A commitment parameter γ_S records how strongly a sector overlaps the closure-norm condensate:

$\gamma_u \approx \gamma_d \approx 1$ (committed quark regime), $\gamma_e \approx 1$, $\gamma_\nu = \varepsilon \ll 1$ (weak-commitment lepton regime).

This is bookkeeping, not an extra fit: the closure-norm condensation result gives the structural reason neutrinos live near the weak-commitment end.

4.3 Weak-role selection

A clean flavour-frame readout requires the role-changing piece Ω_{mix} to be subleading or gapped. The first projected-Hamiltonian test is therefore not numerical: check that $P_W \cdot H_{cl} \cdot P_W$ is approximately role-diagonal in the committed quark regime and asymmetrically role-diagonal in the lepton regime.

5. The Closure-Committed Quark Regime

For quarks both doublet partners are closure-committed. In the product reading,

$$U_u = U_0 \cdot \exp(-\Omega_q/2),\ U_d = U_0 \cdot \exp(+\Omega_q/2),\ V_{CKM} = U_u^\dagger \cdot U_d = \exp(\Omega_q),$$

the clean exponential is exact. The single-projected-generator (additive) reading is the physically relevant one:

$U\_u = \exp(\Omega_0 - \Omega\_q/2)$, $U\_d = \exp(\Omega_0 + \Omega\_q/2)$,

so common-mode transport does not cancel unless it commutes with the role-odd split. BCH gives

$$\log(U\_u^\dagger \cdot U\_d) = \Omega\_q - \tfrac{1}{2}[\Omega_0, \Omega\_q] + (1/6)[\Omega_0, [\Omega_0, \Omega\_q]] + \cdots.$$

The projection formulation makes Ω_0 and Ω_q two faces of the same H_cl block:

$$\Omega_0 = P\_even \cdot \Omega\_W \cdot P\_even, \quad \Omega\_q/2 = P\_odd \cdot \Omega\_W \cdot P\_odd.$$

The CKM curvature residue is therefore **the failure of two projections of the same substrate Hamiltonian to commute**, $\kappa = [\Omega_0, \Omega\_q]$.

---

6. CKM Curvature from the Role-Even Projection

6.1 Exact algebra

Use the inherited quark relative generator

$$\Omega\_q = \begin{bmatrix} 0 & a & c \cdot e^{(i\varphi)} \\ -a & 0 & b \\ -c \cdot e^{(-i\varphi)} & -b & 0 \end{bmatrix}$$

with a = 9/40, b = 81/2000, c = 243/100000, φ = 2π/3. The minimal role-even curvature lives in the 2↔3 common-mode channel, $\Omega_0(z) = z \cdot E_{23} - z^* \cdot E_{32}$. Then **exactly** (verified)

$$[\Omega_0, \Omega\_q]_{13} = -a \cdot z, \quad \Delta_{13} = -\tfrac{1}{2}[\Omega_0, \Omega\_q]_{13} = \tfrac{1}{2} \cdot a \cdot z.$$

A 2↔3 common curvature dragged through the wide 1↔2 Cabibbo doorway produces a 1↔3 triangle residue — exactly the sector the clean exponential underproduces.

6.2 Selection of the 2↔3 common-mode channel

Among role-even common-mode corrections, 2↔3 is the minimal admissible candidate. A 1↔2 correction would attack the already-inherited Cabibbo entry; a direct 1↔3 correction would act like a new direct c-entry; a 2↔3 correction is invisible to the leading 1↔2 entry at first order but becomes visible in the triangle through the Lie bracket [2↔3, 1↔2] → 1↔3. This is an exclusion-style argument, not yet a theorem of H_cl: the first admissible common-mode channel should be the one that repairs the missing triangle while protecting the inherited Cabibbo scale.

6.3 Amplitude as root-normalised C_3 sharing

If the projected role-even curvature inherits the $2\leftrightarrow3$ transport norm b but distributes it democratically over a threefold C_3 closure orbit, norm conservation gives the component amplitude

$$|z| = b/\sqrt{3}.$$

This follows from ordinary Hilbert-space norm sharing if the C_3 branches are Killing-orthonormal at the projected H_cl level. The open derivation is now precise: prove that the relevant role-even $2\leftrightarrow3$ curvature is a C_3-democratic component of the H_cl projection.

6.4 Phase as half-transport against complement reversal

The phase target is $\arg z = (\pi + 2\pi/3)/2 = 5\pi/6$. The strengthened reading is a **half-transport (square-root) rule**. Let $C = e^{i\pi}$ be complement reversal and $O_q = e^{2\pi i/3}$ the C_3 generation-loop orientation. The common-mode half-transport h satisfies

$$h^2 = C\cdot O_q, \quad h = e^{i(\pi + 2\pi/3)/2} = e^{5\pi i/6}.$$

(Verified: $(e^{5\pi i/6})^2 = e^{5\pi i/3} = e^{i\pi}\cdot e^{2\pi i/3}$.) This is elegant but remains the weaker half of the CKM derivation: the closure Hamiltonian must show the common-mode curvature is the square-root transport between role complement and C_3 loop orientation, rather than a direct holonomy or another branch.

6.5 CKM output audit (verified)

Quantity	Clean exp(Ω_q)	C_3 curvature	Interpretation
$\lvert V_{us}\rvert$	0.22307	0.22307	stable; inherited Cabibbo scale
$\lvert V_{cb}\rvert$	0.04028	0.04054	stable
$\lvert V_{ub}\rvert$	0.00393	0.00356	cost of curvature
$\lvert V_{td}\rvert$	0.00611	0.00869	repaired
$\lvert J\rvert$	1.84×10^{-5}	3.00×10^{-5}	near observed scale
$\lvert V_{ub}\rvert/\lvert V_{cb}\rvert$	0.0976	0.0877	improved but still high

PDG-2024 ballpark for orientation: $|V_{us}| \approx 0.2243$, $|V_{cb}| \approx 0.0408$, $|V_{ub}| \approx 0.00382$, $|V_{td}| \approx 0.00857$, $J \approx 3.08\times10^{-5}$. The curvature improves the triangle and CP outputs while spending $|V_{ub}|$; the $|V_{ub}|$-ratio wall remains an important near-falsifying constraint.

---

7. The Weak-Commitment Neutrino Regime

The lepton doublet is asymmetric. The charged lepton is closure-norm anchored and mass-resolved; the neutrino is weakly anchored and nearly neutral in the closure-norm readout. The projection should therefore give an anchored charged-lepton frame and a residual neutrino frame:

$U\_e \approx U_0\hat{}\ell$, $U\_\nu = U_0\hat{}\ell \cdot U\_wc$, $U\_PMNS = U\_e^\dagger \cdot U\_\nu \approx U\_wc$.

Weak commitment means the neutrino stiffness gaps collapse together with the off-diagonal transport scale:

$D\_\nu = D_0 \cdot I + \varepsilon \cdot \Delta$, $\Pi\_\nu = \varepsilon \cdot K$, $T\_\nu = D_0 \cdot I + \varepsilon \cdot (\Delta + K) = D_0 \cdot I + \varepsilon \cdot M\_\nu$.

For every $\varepsilon > 0$, T_ν and M_ν share eigenvectors. The small neutrino scale sets the eigenvalues, not the PMNS frame. This is the exact weak-commitment frame theorem — the cleanest part of the PMNS architecture.

---

8. PMNS Closure-Kernel Rules from the Projected Hamiltonian

The target Hermitian weak-commitment operator is

$$M\_\nu = \begin{bmatrix} r\_e & A & A(1 + \beta \cdot e\hat{}(i\psi)) \\ A & r\_s + \delta & B \\ A(1 + \beta \cdot e\hat{}(-i\psi)) & B & r\_s - \delta \end{bmatrix}$$

8.1 Symmetric core and solar ratio

In the μ–τ-symmetric core, rotate to $\nu_+ = (\nu\_\mu + \nu\_\tau)/\sqrt{2}$ and $\nu_- = (\nu\_\mu - \nu\_\tau)/\sqrt{2}$. The solar block gives

$\tan 2\theta_{12} = 2\sqrt{2} / \rho\_\odot$, $\rho\_\odot = (r\_s + B - r\_e)/A$.

The projection strengthening is that $\rho\_\odot$ is not a free shape parameter. If the weak-commitment residual cloud has threefold support norm and the coherent non-electron pair has twofold support norm, then

$\rho\_\odot = \sqrt{3}/\sqrt{2} = \sqrt{(3/2)}$,

giving $\theta_{12} \approx 33.4°$. This is the PMNS analogue of the CKM amplitude rule $|z| = b/\sqrt{3}$ — both are root-normalised support comparisons.

8.2 Pair ratio B/A = 6/5

The flavour-mixing work supplies a five-dimensional residual space $R = C \oplus G$ (dim $C = 3$, dim $G = 2$, dim $R = 5$), the natural denominator for weak-commitment continuation counts. If the internal μ/τ pair carries one additional continuation channel relative to the electron-to-pair link, then

$B/A = (5 + 1)/5 = 6/5.$

The denominator 5 is now inherited from the residual support dimension rather than chosen. The remaining proof debt is the numerator: show from the H_cl weak-commitment projection that the internal pair receives exactly one extra continuation, not zero, two, or a root-normalised alternative.

8.3 Diagonal μ–τ breaking

Weak neutrality suppresses first-order diagonal self-stiffness splitting between μ and τ, so the leading breaking sits in the attachment channel:

$\delta = 0 + O(\beta^2).$

This is a selection rule, not a fit. If the H_cl projection produces a first-order diagonal μ–τ split, the kernel loses much of its predictive content.

8.4 Weak-commitment phase

The half-transport grammar of §6.4 carries over, replacing C_3 orientation with weak-neutral quadrature $O_\nu = e^{i\pi/2}$:

$h_\nu^2 = C \cdot O_\nu, \psi = (\pi + \pi/2)/2 = 3\pi/4.$

(Verified: $(e^{3\pi i/4})^2 = e^{3\pi i/2} = e^{i\pi} \cdot e^{i\pi/2}$.) Thus the CKM and PMNS phases are two sectoral square roots against complement reversal:

$\theta_q = (\pi + 2\pi/3)/2 = 5\pi/6, \psi_\nu = (\pi + \pi/2)/2 = 3\pi/4.$

8.5 Leakage amplitude β

The benchmark leakage amplitude remains the weakest target:

$\beta^\star = \sqrt{3}/20.$

The numerator $\sqrt{3}$ fits the root-normalised threefold leakage idea; the denominator 20 plausibly connects to the twentyfold support trace of the mass-trace programme, but that bridge is not discharged here. β is therefore a named support-trace target, not a derived rule.

8.6 Octant and CP sign: the breaking-channel premise (WF-2)

Section 8.3 places the leading breaking in "the attachment channel" — but *which* electron-attachment channel is not innocuous. Direct diagonalisation shows the **magnitude** of the atmospheric departure from maximal mixing is robust,

$|\theta_{23} - 45°| \approx 3.4°,$

while the **octant** and the **sign of δ_CP** are jointly fixed by a binary choice:

Leading breaking channel	θ_{23} (standard PDG)	δ_CP
electron–μ, the (1,2) entry	**41.6° (lower octant)**	≈ 47° (J > 0)
electron–τ, the (1,3) entry	**48.4° (upper octant)**	≈ 227° (J < 0)

These are exact μ↔τ reflections (41.6° = 90° − 48.4°) with opposite CP sign; θ_{12} ≈ 33.4°, θ_{13} ≈ 8.6°, and the magnitude |J_lep| ≈ 0.024 are identical in both, but the *sign* of J_lep flips between branches. The δ_CP quadrant is then forced, not free: in the PDG convention sgn(J) = sgn(sin δ_CP), so J < 0 ⟺ δ_CP ∈ (180°, 360°). The electron–τ branch has J = −0.024, fixing δ_CP = 227° (sin δ < 0, cos δ < 0); the electron–μ branch has J = +0.024, fixing δ_CP = 47°. The M_ν written above places the (1 + β·e^(±iψ)) factor on the (1,3)/(3,1) **electron–τ** entries, which under standard PDG conventions yields the **upper** octant (48.4°) with δ_CP = 227°. Any benchmark quoting the lower octant (41.6°) is instead the **electron–μ** branch with δ_CP = 47°. The two must not be conflated.

WF-2 (octant channel premise). Closure geometry must determine which electron-attachment channel carries the leading μ–τ breaking. Until WF-2 is derived, the framework predicts the *magnitude* $|\theta_{23} - 45°|$ ≈ 3.4°, together with θ_{12}, θ_{13}, and |J_lep| — but the octant and CP sign are predicted only up to a twofold reflection.

This converts a hidden ambiguity into a precise named target: resolving WF-2 would upgrade four robust outputs to six.

---

9. The Shared Phase Grammar

The strongest unifying pattern is the phase grammar:

phase = ½·(sector orientation + complement reversal), equivalently $h^2 = C \cdot O$.

Complement reversal C = e^(iπ) is shared because both sectors are read through weak-doublet complementarity. The sector orientation differs by regime: C_3 generation holonomy O_q = e^(2πi/3) for quarks, weak-neutral quadrature O_ν = e^(iπ/2) for neutrinos. Both identities hold exactly:

CKM: $e^{(5\pi i/6)^2}$ = e^(iπ)·e^(2πi/3), PMNS: $e^{(3\pi i/4)^2}$ = e^(iπ)·e^(iπ/2).

If H_cl derives this square-root transport rule, the CKM common-mode phase and the PMNS weak-commitment phase become one theorem rather than two choices.

---

10. Unified Theorem Package

Theorem A — Projected Hamiltonian. If H_cl exists as the su(8)-restricted admissibility Hessian and P_W is the electroweak/generation/readout projector, then $H_W = P_W \cdot H_{cl} \cdot P_W$ is the unique candidate weak-doublet flavour Hamiltonian at this level of the programme.

Theorem B — Role decomposition. If role-changing components are gapped, H_W induces role-even and role-odd frame generators; the role-odd piece gives the leading relative sector split, the role-even piece gives common-mode curvature.

Theorem C — CKM residue. In the committed quark regime, additive sector frames give $\Delta_{CKM} = -\frac{1}{2}[\Omega_0, \Omega_q] + \cdots$. For a pure 2↔3 role-even curvature, the dominant triangle correction is $\Delta_{13} = \frac{1}{2} \cdot a \cdot z$.

Theorem D — Weak-commitment frame. If $T_\nu = D_0 \cdot I + \varepsilon \cdot M_\nu$ with $\varepsilon > 0$, the PMNS neutrino frame is the eigenframe of M_ν, independent of ε.

Theorem E — Shared root-count and phase grammar. If the H_cl projection is C_3/Killing-normalised and complement-half-transport selected, it yields $|z| = b/\sqrt{3}$, $\rho_\odot = \sqrt{3/2}$, $\theta_q = 5\pi/6$, and $\psi_\nu = 3\pi/4$.

The theorems are conditional on the projection $P_W \cdot H_{cl} \cdot P_W$ being computed; they state what that computation must deliver.

---

11. Derive-or-Reject Audit

Target	Strength	Reason
$H_W = P_W \cdot H_{cl} \cdot P_W$	strong	Anchors the weak-doublet Hamiltonian in an inherited substrate generator.
Charge-blindness	strong	Inherited commutator rule prevents generation-dependent charge.
CKM channel 2↔3	moderate	Minimal channel that repairs the triangle while protecting Cabibbo.
CKM amplitude $b/\sqrt{3}$	conditional, plausible	Follows from root-normalised C_3 sharing if branches are orthonormal/democratic.
CKM phase $5\pi/6$	conditional, soft	Needs the half-transport theorem from H_cl, not just bisector arithmetic.
PMNS $\rho_\odot = \sqrt{3/2}$	conditional, strong	Same root-count logic as the CKM amplitude.
PMNS $B/A = 6/5$	improved, not proven	Denominator 5 inherited from residual support; numerator one-extra-continuation still owed.

Target	Strength	Reason
PMNS $\delta = 0$	good selection rule	Follows if weak neutrality forbids first-order self-stiffness split.
PMNS $\psi = 3\pi/4$	conditional, unified	Same half-transport grammar as CKM.
PMNS $\beta = \sqrt{3}/20$	weakest	Needs an explicit twenty-support leakage theorem.
Octant / CP sign (WF-2)	**open**	Fixed by the breaking-channel choice; only $\|\theta_{23} - 45°\|$ is currently predicted.

---

12. Falsification Conditions

1. If P_W·H_cl·P_W cannot be defined without fitting sector-frame matrices, the programme collapses back to a phenomenological ansatz.
2. If the projected Hamiltonian fails charge-blindness, it predicts forbidden generation-dependent charge shifts and must be rejected.
3. If the role-even projection does not produce a small 2↔3 common-mode curvature, the CKM C_3 residue is not derived.
4. If the projected curvature phase is far from $5\pi/6 \approx 150°$ (or from the empirically preferred phase window), the CKM phase grammar fails.
5. If the weak-commitment projection does not give $T_\nu = D_0 \cdot I + \varepsilon \cdot M_\nu$, the large-PMNS explanation fails.
6. If the weak-commitment support catalogue gives B/A far from 6/5, the kernel must be revised.
7. If the support trace does not produce or bound β near $\sqrt{3}/20$, the reactor-angle magnitude remains fitted rather than derived.
8. **Octant / channel (WF-2), sharpened.** The kernel predicts a definite octant *only once WF-2 fixes the breaking channel*. Its channel-independent claim is the magnitude $|\theta_{23} - 45°| \approx 3.4°$. If data place θ_{23} consistent with maximal mixing to well within this deviation, β is too large as written. If WF-2 is derived and selects the channel opposite to the measured octant, the kernel is falsified.
9. If |V_ub| and |V_ub|/|V_cb| remain incompatible with the 2↔3 curvature wall after experimental tensions settle, a second curvature component or a different common-mode channel is required.
10. If CKM requires one closure rule and PMNS a genuinely unrelated rule, the unified projection fails and the flavour programme loses its strongest unification claim.

---

13. Conclusion

The central object is no longer a four-term ansatz but H_W = P_W·H_cl·P_W, where H_cl is the inherited su(8) closure Hamiltonian. This moves the programme from model-building toward a concrete projection calculation.

In the quark committed regime, the role-odd projection supplies the leading CKM generator and the role-even projection supplies common-mode curvature. The CKM C_3 residue is a precise target: derive a 2↔3 common-mode curvature with amplitude b/√3 and phase 5π/6 from the H_cl projection. Its commutator with the Cabibbo link repairs J and |V_td| (verified: |V_td| 0.0061 → 0.0087, |J| 1.84 → 3.00 $\times 10^{-5}$) while exposing the |V_ub|-ratio wall.

In the neutrino weak-commitment regime, closure-norm condensation explains why the charged lepton is anchored while the neutrino is weakly committed. The projected neutrino operator has the scale–shape form $T_\nu = D_0 \cdot I + \varepsilon \cdot M_\nu$, so the PMNS frame is set by M_ν rather than by the tiny neutrino scale. The solar ratio and weak-commitment phase join the same root-count and half-transport grammar as the CKM curvature; B/A and β remain support-count targets, and the atmospheric octant and CP sign await the breaking-channel premise WF-2.

CKM and PMNS are the committed and weak-commitment projections of one closure Hamiltonian.

The bottom line is stronger but still conditional. The architecture is now anchored in an inherited generator rather than an invented one. But the programme must still compute the projection of H_cl. Until that calculation derives z_C3, B/A = 6/5, and β = √3/20 — and fixes the octant channel — the result is a sharpened derivation target, not a completed Standard Model flavour derivation. If it succeeds, VERSF gains a unified flavour mechanism: small CKM from committed doublet splitting, large PMNS from weak-commitment stiffness collapse, and both CP structures from one complement-reversal phase grammar. If it fails, it fails cleanly.

---

14. References and Uploaded Source Map

- *Substrate Dynamics and the Higgs Ratio* — H_cl as the su(8)-restricted closure Hamiltonian; the minimal admissible refinement update exp(i·ε·H_cl).
- *Closure-Norm Condensation and Electroweak Symmetry Breaking in VERSF* — closure-norm condensation, electroweak mass generation, weak neutrino overlap with the closure-norm direction.
- *The Yukawa Hierarchy Theorem* — charge-blindness [Ŷ, χ] = 0 and the completion-density mass-readout architecture.
- *The Structural Origin of Yukawa Operators in VERSF* — closure-depth overlap reading of Yukawa operators.
- *Realised Readout and Charge-Sector Maintenance* — the realised-readout and signed-sector-contrast proof pattern.
- *The Yukawa Operator from Completion-Channel Misalignment in VERSF* — left Yukawa-square operator; sector-frame mismatch architecture.
- *The Electroweak Flavour-Frame Operator in VERSF* — role decomposition; QF-1/QF-2 distinction; weak-commitment frame theorem.
- *The CKM Curvature Residue in VERSF* — z_C3; the exact CKM audit; the |V_ub|-ratio wall.

- *The Weak-Commitment Closure Kernel in VERSF* — PMNS solar ratio; pair ratio; μ–τ breaking relation; weak-commitment phase.
- *The Weak-Commitment Neutrino Operator in VERSF* — PMNS mixing from stiffness-gap collapse; μ–τ symmetric core.
- *Deriving Flavour Mixing from Closure Geometry* — $D = \mathrm{diag}(1, 2, 4)$; $\eta = 3/5$, $\chi = 2/5$; the five-dimensional residual support space $R = C \oplus G$.

---

15. Numerical Note

CKM. The audit uses Ω_q with $a = 9/40$, $b = 81/2000$, $c = 243/100000$, $\varphi = 2\pi/3$, and common-mode curvature $\Omega_0(z) = z\cdot E_{23} - z^*\cdot E_{32}$, $z = z_{C3} = (b/\sqrt{3})\cdot e^{(5\pi i/6)}$. The CKM matrix is $V = \exp(-\Omega_0 + \tfrac{1}{2}\Omega_q)\cdot\exp(\Omega_0 + \tfrac{1}{2}\Omega_q)$, unitary to machine precision, giving $|V_{us}| = 0.22307$, $|V_{cb}| = 0.04054$, $|V_{ub}| = 0.00356$, $|V_{td}| = 0.00869$, $|J| = 3.00\times10^{-5}$, $|V_{ub}|/|V_{cb}| = 0.0877$. With $z = 0$ the additive reading reduces to $\exp(\Omega_q)$, giving the "clean" column ($|V_{td}| = 0.00611$, $|J| = 1.84\times10^{-5}$, $|V_{ub}| = 0.00393$). The identity $[\Omega_0, \Omega_q]_{13} = -a\cdot z$, hence $\Delta_{13} = \tfrac{1}{2}\cdot a\cdot z$, is confirmed symbolically and numerically.

PMNS. Scale/shift convention $A = 1$, $r_e = 0$, with $B = 6/5$, $r_s = \sqrt{(3/2)} - 6/5$, $\delta = 0$, $\psi = 3\pi/4$, $\beta = \sqrt{3/20}$, so $\rho_\odot = (r_s + B - r_e)/A = \sqrt{(3/2)}$. The Hermitian operator M_ν has diagonal entries $(0, \sqrt{(3/2)} - 6/5, \sqrt{(3/2)} - 6/5)$; the electron–μ link is the real entry $A = 1$; the μ–τ link is $B = 6/5$; and the leading breaking sits on the **electron–τ** link, which carries the complex factor $1 + (\sqrt{3/20})\cdot e^{(+3\pi i/4)}$ above the diagonal and its conjugate below.

Exact Hermitian diagonalisation (ν_3 = column of smallest electron content) gives $\theta_{12} \approx 33.4°$, $\theta_{13} \approx 8.6°$, $J_{lep} \approx -2.4\times10^{-2}$, and $\theta_{23} \approx$ **48.4° (upper octant)**, with $\delta_{CP} \approx 227°$ (forced by $J < 0$ via $\mathrm{sgn}\, J = \mathrm{sgn} \sin \delta_{CP}$; full quadrant: $\sin \delta = -0.73$, $\cos \delta = -0.68$).

Octant branch (WF-2). Moving the same breaking factor to the (1,2)/(2,1) **electron–μ** entries leaves θ_{12}, θ_{13}, and $|J_{lep}|$ unchanged but flips J_{lep} to $+2.4\times10^{-2}$, mapping $\theta_{23} \to 41.6°$ (lower octant) and $\delta_{CP} \to 47°$. The two branches are exact μ↔τ reflections with opposite CP sign. The channel-independent invariant is $|\theta_{23} - 45°| \approx 3.4°$. Until WF-2 is derived, only this magnitude (with θ_{12}, θ_{13}, $|J_{lep}|$) is a prediction; the octant and CP sign are not.

The half-transport phase identities are exact: $e^{(5\pi i/6)^2} = e^{(i\pi)}\cdot e^{(2\pi i/3)}$ for CKM, $e^{(3\pi i/4)^2} = e^{(i\pi)}\cdot e^{(i\pi/2)}$ for PMNS.

These numbers audit the projected-Hamiltonian target; they are not a substitute for computing the projection of H_{cl}.