The Weak-Doublet Projection Theorem in VERSF

Computing the Admissible Block of H_cl that Selects CKM Curvature, Weak-Commitment Neutrino Mixing, and the Octant Branch

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

Successor to *The Projected Weak-Doublet Closure Hamiltonian in VERSF*, *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*, *The Yukawa Operator from Completion-Channel Misalignment in VERSF*, and the substrate-dynamics result that identifies H_cl as the su(8)-restricted admissibility Hessian.

Summary for the General Reader

Matter comes in three "generations" — three copies of the same handful of particles, identical except that each copy is heavier than the last. Particles don't stay perfectly inside their own generation, though. They *mix*: a particle from one generation can behave a little like its cousins in the others. Quarks mix one way; their lighter cousins, the neutrinos, mix in a strikingly different way. For decades these two patterns of mixing have been treated as separate puzzles, each described by its own table of numbers measured in experiments but explained by neither.

The previous paper in this series argued that the two puzzles are really one. Both kinds of mixing, it claimed, are two views of a *single underlying object* — like a sculpture that looks different depending on which side you walk around. This paper asks the obvious next question: **if there really is one object, what shape is it allowed to have?** Crucially, it tries to answer that *before* fitting anything to the measured numbers, so the answer can't just be reverse-engineered from the data.

It turns out three simple, almost commonsense requirements pin the shape down hard:

- You can't change a particle's electric charge just by re-reading its mass. (Mixing is allowed to shuffle generations, not to turn one kind of particle into another.)
- There's a clean separation between the two partners in a pair moving *together* and the two moving *oppositely*.
- Quarks and neutrinos are "anchored" to the underlying structure very differently — and that single difference is what makes their mixing patterns look so unlike each other.

The quark story is the cleanest new result, and it has a nice intuition. Picture the three generations as three rungs of a ladder. The mixing between rungs 1 and 2 is large and very well measured — you don't want to disturb it. The mixing between rungs 1 and 3 is small and *under*-produced by the simple version of the theory: there's a missing piece. Now, there are three ways you might try to supply that missing piece:

- Push directly on the 1–3 rungs. But that's just pasting in the answer by hand — it explains nothing.
- Push on the 1–2 rungs. But that wrecks the one big number the theory already gets right.
- Push gently on the 2–3 rungs. By itself this seems to do almost nothing. **But combined with the large 1–2 motion that's already there, it produces exactly the missing 1–3 piece.**

Under the admissibility rules used in this paper, the third route is the only honest one left. It preserves the successful 1–2 mixing, avoids simply inserting the missing 1–3 answer by hand, and still generates the missing 1–3 effect through the fact that two allowed motions fail to commute. The missing effect isn't inserted — it *emerges* from two motions happening in an order that matters (do A then B, and you don't get the same thing as B then A; the difference is precisely the missing piece). That's a genuine explanation rather than a fudge factor.

The neutrino story answers a long-standing oddity: neutrinos weigh almost nothing, yet they mix enormously — far more than quarks do. Naively you'd expect tiny masses to mean tiny mixing. The structure here resolves it neatly. The tiny mass scale acts like an overall volume knob: it makes everything faint, but it doesn't touch the *directions* in which the neutrinos mix. Mass and mixing are set by two different things, so you can have feather-light masses sitting alongside wide-open mixing with no contradiction.

The paper also cleans up one stubborn loose end. Experiments know that one neutrino mixing angle is close to 45°, but not whether it sits a touch *above* or a touch *below*. That used to be a genuinely open question in the framework. Here it's reduced to the **sign of a single quantity** — essentially one tie-breaker you can compute: positive means the angle tips one way, negative the other. A fuzzy ambiguity has become a sharp yes/no test.

The honest bottom line: this paper proves what *shape* the unified object must have, and writes down the short, specific list of numbers left to calculate to confirm or kill the picture. It does **not** claim that hard calculation is finished. Its real achievement is turning a vague conceptual question — "are these two puzzles secretly the same?" — into a concrete, finite to-do list. And if any item on that list comes out wrong, it fails *cleanly*: you'll know exactly which assumption to revisit.

Abstract

The projected weak-doublet closure programme reduces the residual Standard Model flavour problem to the sector block

$H_W = P_W H_{cl} P_W$,

where H_cl is the inherited su(8)-restricted admissibility Hessian of the substrate free-energy functional and P_W projects onto generation completion, weak role, and mass readout. This paper strengthens that reduction into a **projection theorem**. Under charge-blind mass readout, weak-role gap admissibility, generation-depth covariance, complement reversal, and commitment-class separation, the weak-doublet block decomposes as

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

with Ω_mix gapped at leading order. In the committed quark regime, Ω_odd supplies the leading CKM split, while the first Cabibbo-protected role-even curvature is the **2↔3** common-mode channel. Its Baker–Campbell–Hausdorff residue obeys the exact identity

$$\boldsymbol{[\Omega_0(z) , \Omega_q]_{13} = -a\cdot z , \Delta_{13} = \tfrac{1}{4}\cdot a\cdot z} ,$$

so the missing triangle correction is generated by *dragging the 2↔3 curvature through the 1↔2 Cabibbo doorway*. If the projected role-even block is a Killing-orthonormal democratic C_3 component, its amplitude is $z = b/\sqrt{3}$; if common-mode transport is the square root of complement reversal against the C_3 generation orientation, its phase is $\arg z = 5\pi/6$.

In the weak-commitment neutrino regime, closure-norm decoupling gives

$$\boldsymbol{T_\nu = D_0\cdot I + \varepsilon\cdot M_\nu} ,$$

so the PMNS frame is the eigenframe of M_ν, independent of ε. Root-normalised residual-support comparison gives the solar target $\rho_\odot = \sqrt{3/2}$. The five-dimensional residual support space gives the denominator of the pair ratio $B/A = 6/5$, the numerator following from a single-return continuation rule. The weak-neutrality selection rule suppresses first-order diagonal μ–τ splitting, placing the leading atmospheric breaking in an electron-attachment channel. The remaining octant ambiguity collapses to one substrate sign,

$$\boldsymbol{\sigma_W = \mathrm{sgn}(\|P_e H_W P_\tau\|^2 - \|P_e H_W P_\mu\|^2)} ,$$

which selects the electron–tau or electron–muon branch and hence fixes the atmospheric octant and leptonic CP sign. The leakage amplitude $\beta = \sqrt{3}/20$ remains the weakest target and is now isolated as a **K = 7-descended support-trace invariant**: the numerator $\sqrt{3}$ is the root-normalised threefold leakage source, while the denominator is no longer a free number but the projected support trace $\mathrm{Tr\_support}(\Pi_{leak}) = (K - 2) \times 4 = 5 \times 4 = 20$, where $K = 7$ is the closure-admissibility carrier, two directions are fixed by boundary/readout anchoring, and the remaining five residual directions tensor with a fourfold weak-attachment orientation fibre whose second factor of two is the load-bearing open premise. The paper therefore moves the programme from a *unified projection proposal* to a *finite set of explicit projected-Hamiltonian tests*. The result is not yet the completed microscopic computation of P_W H_cl P_W, but a strengthened admissibility theorem that identifies the exact projected blocks, signs, ratios, and support traces the next calculation must derive or reject.

Notation

Symbol	Meaning
$H_{cl} = G \mid su(8)$	closure Hamiltonian: su(8)-restricted Hessian G of the substrate free-energy functional
P_C , P_R , P_Y	generation-completion, weak-role, and mass/readout projectors
$P_W = P_Y P_R P_C$	weak-doublet projector; $H_W = P_W H_{cl} P_W$
$\Omega_W = i\cdot\varepsilon_W\cdot H_W$	anti-Hermitian frame generator
τ_1, τ_2, τ_3	Pauli matrices on the 2-D weak-role space
Ω_{even} , Ω_{odd} , Ω_{mix}	role-diagonal common-mode, role-odd, and role-changing blocks
$\gamma_u, \gamma_d, \gamma_e, \gamma_\nu$	closure-commitment weights of the doublet partners
$a \gg b \gg c, \varphi$	inherited quark generator entries and CP phase
$\Omega_0(z) = z\cdot E_{23} - z^*\cdot E_{32}$	role-even 2↔3 common-mode curvature
$T_\nu = D_0\cdot I + \varepsilon\cdot M_\nu$	weak-commitment neutrino readout
$\rho_\odot$, B/A, δ, ψ, β	PMNS kernel invariants (solar ratio, pair ratio, diagonal split, weak phase, leakage)
σ_W	weak-attachment orientation (octant sign)
$\lVert\cdot\rVert_K$	Killing norm; $\langle\cdot,\cdot\rangle_K$ Killing inner product
K = 7	closure-admissibility carrier count, inherited from the K = 7 closure programme
R	residual weak-commitment support after boundary/readout anchoring; dim R = K − 2 = 5
Q	fourfold weak-attachment orientation fibre; dim Q = 4 = 2 × 2 (weak-neutral quadrature × closure two-cycle binary)
Π_{leak}	projector onto the weak-commitment leakage support; $\mathrm{Tr}_{support}(\Pi_{leak}) = 20$

Contents

---

0. Predictive-Content Ledger

This paper addresses the next target after the projected-Hamiltonian formulation.

- **WP-1 — Weak-doublet projection theorem.** Derive the admissible block structure of P_W H_cl P_W from charge-blindness, weak-role decomposition, and commitment class.
- **WP-2 — CKM curvature selection.** Show that the first Cabibbo-protected common-mode curvature is the role-even 2↔3 channel.
- **WP-3 — Residual-support kernel.** Show that the weak-commitment neutrino block has the scale–shape form $D_0 \cdot I + \varepsilon \cdot M\_\nu$, with support ratios matching the PMNS kernel.
- **WP-4 — Octant branch reduction.** Replace the previous open octant premise by a single computable substrate sign σ_W.

Object	Result in this paper	Status
H_W = P_W H_cl P_W	unique weak-doublet object	inherited

Object	Result in this paper	Status
Role decomposition	Pauli decomposition on weak-role space	exact once P_W is defined
Role-changing block	gapped/suppressed for clean flavour frames	admissibility test
CKM common curvature channel	2↔3 selected by Cabibbo protection	selection theorem under stated axioms
BCH triangle residue	$[\Omega_0, \Omega\_q]_{13} = -a\cdot z$; $\Delta_{13} = \frac{1}{4}\cdot a\cdot z$	exact algebra (coefficient ¼ fixed by Axiom C split)
CKM amplitude	$z = b/\sqrt{3}$	derived if projected C_3 branches are orthonormal/democratic
CKM phase	$\arg z = 5\pi/6$	conditional on complement-half-transport **and** branch invariant σ_φ^q
Neutrino weak limit	$T\_\nu = D_0\cdot I + \varepsilon\cdot M\_\nu$	exact eigenframe theorem
Solar ratio	$\rho\_\odot = \sqrt{(3/2)}$	root-support theorem under residual-comparison axiom
Pair ratio	$B/A = 6/5$	denominator from dim R = 5; numerator from single-return continuation
Diagonal μ–τ split	$\delta = 0 + O(\beta^2)$	weak-neutrality selection rule
Weak phase	$\psi = 3\pi/4$	conditional on asserted O_ν **and** branch invariant σ_φ^ν
Leakage β	$\sqrt{3}/20$	K = 7-descended support-trace target: numerator √3 from root-normalised threefold leakage; denominator 20 from Tr_support(Π_leak) = (K − 2) × 4
Leakage support trace	Tr_support(Π_leak) = 20	conditional theorem if leakage support factors as R ⊗ Q with dim R = K − 2 = 5 and dim Q = 4
Octant & CP sign	fixed by σ_W	precise substrate sign test

Three classes of statement are kept rigorously separate throughout.

1. **Exact algebra** — Pauli role decomposition, the BCH commutator identity, and the ε-independence of the weak-commitment eigenframe. These hold with no further assumption.
2. **Projection theorems under axioms** — channel selection, root-normalised amplitudes, half-transport phases, residual-support ratios. These hold *given* explicitly stated projection axioms.
3. **Substrate computation still owed** — the actual evaluation of the relevant P_W H_cl P_W matrix elements.

The discipline of the paper is one rule: **every surviving number must be either an exact algebraic consequence, a theorem from explicitly stated projection axioms, or a named substrate invariant still awaiting computation.**

---

1. The Problem After the Projected-Hamiltonian Paper

The previous stage established the correct object:

$$H_W = P_W H_{cl} P_W .$$

This is stronger than writing down a flavour Hamiltonian by hand. H_{cl} is *inherited* from the substrate-dynamics layer as the su(8)-restricted admissibility Hessian; the weak-doublet operator is its electroweak, generation, role, and readout projection.

That move solved one problem and exposed a harder one.

- It **solved the architectural problem**: CKM and PMNS can be read as two commitment regimes of the same object.
- It **exposed the projection problem**: what does $P_W H_{cl} P_W$ actually *force*?

The present paper introduces **no new flavour ansatz**. Its task is narrower and sharper:

Determine the admissible block structure of $P_W H_{cl} P_W$, then identify which entries must be computed to derive — or reject — the CKM and PMNS targets.

The intended progress is therefore not another numerical fit but a **reduction** of the remaining flavour programme to a finite projected-Hamiltonian calculation. The central claim is that the projection is *already* constrained enough to select the first CKM curvature channel and to reduce the PMNS octant ambiguity to a single substrate sign — real progress before the full microscopic Hessian is evaluated.

---

2. Projection Data and Notation

Let S be the finite substrate readout space on which the admissibility Hessian acts after su(8) restriction. The closure Hamiltonian is

$$H_{cl} = G \mid su(8) ,$$

where G is the Hessian of the substrate free-energy functional at the admissible manifold. The weak-doublet projection is built from three projectors — P_C (generation completion), P_R (weak role), P_Y (mass/readout) — combined as

$$P\_W = P\_Y P\_R P\_C , H\_W = P\_W H\_{cl} P\_W .$$

The role space is two-dimensional; we use Pauli matrices τ_i on it, with τ_3 distinguishing the two role readings of a left-handed doublet. The generation/readout part is three-dimensional at leading flavour order.

The frame generator is anti-Hermitian,

$$\Omega\_W = i\cdot\varepsilon\_W\cdot H\_W ,$$

and every weak-role block admits the exact Pauli decomposition

$$\Omega\_W = \Omega_0 \otimes I + \Omega_1 \otimes \tau_1 + \Omega_2 \otimes \tau_2 + \Omega_3 \otimes \tau_3 .$$

We identify

$$\Omega\_even = \Omega_0 , \Omega\_odd = \Omega_3 , \Omega\_mix = \Omega_1 \otimes \tau_1 + \Omega_2 \otimes \tau_2 .$$

The **clean flavour-frame regime** is the regime in which Ω_mix is gapped or subleading relative to the role-diagonal blocks.

The commitment structure distinguishes the two doublets:

quarks: $\gamma\_u \simeq \gamma\_d \simeq 1$ (both closure-committed) leptons: $\gamma\_e \simeq 1$, $\gamma\_\nu = \varepsilon \ll 1$ (neutrino weakly committed)

This commitment split is the origin of the CKM/PMNS contrast.

2.1 Projector Compatibility Conditions

The weak-doublet projection is not merely a formal sandwich $P\_W H\_{cl} P\_W$. The projectors must satisfy compatibility conditions that prevent the construction from smuggling in flavour structure by definition.

First, the mass/readout projector must be charge-blind:

$$[P\_Y , \chi] = 0 , \text{ and therefore } [P\_Y H\_{cl} P\_Y , \chi] = 0$$

on the admissible readout subspace. This ensures that generation-depth readout can change mass scale and frame orientation, but cannot alter electric charge or gauge representation.

Second, the weak-role projector must preserve the role grading at leading order:

$$P\_R \tau_3 = \tau_3 P\_R ,$$

with role-changing components confined to the τ_1 and τ_2 blocks. The role-gap condition used later is therefore not a claim that role-changing terms are impossible, but that they are subleading or gapped in the clean flavour-frame regime.

Third, the generation-completion projector must close on the three-generation readout module:

$$P\_C\ H\_cl\ P\_C \subset \mathrm{End}(G_3)\ ,$$

where G_3 is the three-dimensional generation-completion space seen by the leading flavour operator. This is the condition that lets the projected block be represented by 3×3 generation-frame matrices.

Finally, the charged-lepton attachment projectors P_e, P_μ, P_τ used in the octant invariant are to be read as projectors onto charged-lepton-labelled weak-commitment attachment channels *after* the charged-lepton frame has been fixed by closure-norm anchoring. They are not free phenomenological labels: they are the three residual attachment directions inherited by the neutrino block once the charged-lepton readout has selected its stiff frame.

The octant sign

$$\sigma\_W = \mathrm{sgn}(\ \|P\_e\ H\_W\ P\_\tau\|^2 - \|P\_e\ H\_W\ P\_\mu\|^2\ )$$

is therefore a genuine projected-Hamiltonian question: does the weak-commitment block attach more strongly to the electron–tau residual channel or to the electron–muon residual channel? The answer fixes the atmospheric octant branch.

2.2 K = 7 Descent of the Weak-Commitment Support Count

The weak-commitment leakage trace that later fixes β should not be introduced as an isolated denominator. It is a descendant of the K = 7 closure carrier.

The K = 7 closure programme supplies a sevenfold admissibility carrier. In the weak-commitment projection, the leakage channel does not have free access to all seven closure directions: two are already fixed by boundary/readout anchoring — one by the committed charged-lepton frame, one by the weak-doublet/readout complement that defines the attachment basis. The live residual support is therefore

$$\dim R = K - 2 = 7 - 2 = 5\ ,$$

in apparent agreement with the residual-support logic used elsewhere in the flavour programme, where the weak-commitment cloud is organised as a five-dimensional residual module. One reconciliation is still owed here: §8.2 obtains dim R = 5 as the direct sum $C \oplus G = 3 + 2$, whereas §2.2 obtains it as $K - 2 = 7 - 2$. The numerical agreement is suggestive, but the identification $C \oplus G \cong$ (K = 7 carrier $\ominus$ 2 anchored directions) is asserted, not shown; if the two constructions pick out the *same* five-dimensional R it is a genuine consistency, and if not, the symbol R is doing two jobs. The substrate projection must confirm they coincide.

The leakage channel then carries an additional fourfold orientation fibre,

dim Q = 4 = 2 × 2 ,

whose first factor is **weak-neutral quadrature** (the complex attachment occupies two real quadrature directions in the weak-neutral plane). The second factor of two is the genuine open premise of the whole leakage trace, and it carries a non-trivial admissibility burden. To supply a full leading-order factor of two it must be a binary that is *simultaneously*: (i) genuinely two-dimensional; (ii) orthogonal to the Re/Im quadrature plane — otherwise it is the quadrature factor relocated, i.e. a double-count; and (iii) **not** the Axiom-B-gapped role-changing channel.

Requirements (ii) and (iii) pull against each other, and together they exclude the obvious candidates. The naïve reading of the second factor as the **weak-role complement** — both members of the (ν, e_L) doublet partner carrying the attachment — fails (iii): that contribution lives in exactly the role-changing (τ_1, τ_2) sector that Axiom B gaps in the clean flavour-frame regime, and a gapped channel cannot supply a leading-order factor of two to the support trace. Two further readings fail in the other direction: the Hermitian return entry M_τe = M_eτ* is the Im → −Im reflection already counted in the quadrature factor and so violates (ii); and any binary internal to the single neutrino role risks being weak-neutral again, i.e. the quadrature plane once more.

The candidate that may thread all three constraints is the **closure two-cycle binary** (forward/return on the closure cycle) inherited from the half-integer-spin closure structure. The a priori reason to expect it to be independent is categorical: the closure two-cycle is a *topological* Z_2 — the spinor double-cover sign acquired under a 2π rotation — and a topological Z_2 is a different kind of object from a U(1) quadrature phase or a weak-isospin τ flip. That categorical difference is exactly what the rejected "Hermitian forward/return" reading lacked: forward/return there was merely a phase reflection (Im → −Im), living in the same U(1) plane as the quadrature, whereas the double-cover sign is not a phase at all.

Naming it does not discharge the premise; it relocates the burden. Two checks remain owed, and the harder one is *not* orthogonality to the quadrature plane. (i)–(ii): the projection must still exhibit the two-cycle as a genuine two-dimensional subspace not reducible to the quadrature plane — plausible, given the topological-vs-U(1) distinction, but still to be shown. (iii) is the real hazard: because spinor structure and weak-doublet structure are *both* two-valued, the projection must specifically confirm that this binary lives in the closure/holonomy sector and **not** in the weak-role sector — otherwise it fails (iii) in exactly the way the role complement did, and Axiom B gaps it. The load-bearing open claim behind the denominator 20 is therefore the closure-sector residence of the two-cycle, not its mere two-dimensionality. At the support-counting level the leakage projector is then expected to factor as

Π_leak ≃ Π_R ⊗ Π_Q , Tr_support(Π_leak) = dim R · dim Q = (K − 2) × 4 = 5 × 4 = 20 .

This is the support-trace origin of the denominator in $\beta = \sqrt{3}/20$ (§8.5). The result is conditional — the factorisation $\Pi_{leak} \simeq \Pi_R \otimes \Pi_Q$ must ultimately be verified inside $P_W H_{cl} P_W$ — but the denominator 20 is no longer an arbitrary support number: it is the $K = 7$ residual-support trace of the weak-commitment attachment channel.

---

3. Admissibility Axioms for the Weak-Doublet Block

The projection theorem rests on five admissibility conditions. They are stated as **axioms** because the full substrate Hessian calculation has not yet been performed; stating them explicitly is precisely what makes them checkable.

Axiom A — Charge-blind mass readout. The mass/readout projection commutes with the gauge-charge operators χ:

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

This forbids flavour generation from changing electric charge or gauge representation. It permits mass hierarchy and frame transport; it forbids charge drift.

Axiom B — Role-gap admissibility. The role-changing block is gapped at leading flavour-frame order:

$$\|\Omega_{mix}\| \ll \|\Omega_{even}\| + \|\Omega_{odd}\| .$$

This permits a clean role decomposition. Without it, the weak-doublet components would not define separate sector frames in the required approximation.

Axiom C — Committed quark doublet symmetry. In the quark sector both doublet partners are closure-committed, so the leading relative split is role-odd and the leading common curvature is role-even:

$$U_u = \exp((\Omega_0 - \Omega_q)/2) , U_d = \exp((\Omega_0 + \Omega_q)/2) , V_{CKM} = U_u^\dagger U_d .$$

Axiom D — Weak-neutral commitment collapse. In the neutrino sector, weak commitment collapses the nontrivial stiffness scale together with the off-diagonal transport scale:

$$T_\nu = D_0 \cdot I + \varepsilon \cdot M_\nu , \varepsilon > 0 .$$

This does not make neutrinos unstructured; it makes their mass scale small while their residual eigenframe stays finite.

Axiom E — Complement-half-transport. Common-mode curvature is a half-transport between weak-role complement reversal and the sector orientation:

$$\mathbf{h^2 = C \cdot O\ ,\ C = e^{(i\pi)}\ ,\ O_q = e^{(2\pi i/3)}\ ,\ O_\nu = e^{(i\pi/2)}\ .}$$

This is the most delicate axiom. It is **not** forced by ordinary matrix algebra and must ultimately be derived from the geometry of H_cl. The present paper shows what follows *if* it holds and isolates exactly where it must be checked.

Two further inputs inside this axiom carry their own burden and must be tagged **asserted, not derived**, with the same status as Axiom E itself:

- The quark orientation $\mathbf{O_q = e^{(2\pi i/3)}}$ reads naturally as the C_3 generation-loop orientation, and is the more defensible of the two.
- The neutrino orientation $\mathbf{O_\nu = e^{(i\pi/2)}}$ ("weak-neutral quadrature") is the weaker input. It must be derived from the weak-commitment geometry *independently of the phase* $\psi = 3\pi/4$ *it produces*; otherwise a skeptical reader is entitled to see $\pi/2$ as reverse-engineered from the target. The entire neutrino-phase prediction rests on this orientation, so deriving O_ν from H_cl — not from the answer — is a named owed step.

---

4. The Weak-Doublet Projection Theorem

Theorem 1 — Admissible weak-doublet block

Assume Axioms A and B. Then the leading weak-doublet frame generator has the admissible form

$$\mathbf{\Omega_W = \Omega_{even} \otimes I + \Omega_{odd} \otimes \tau_3 + O(\Omega_{mix})\ ,}$$

where Ω_{even} is common-mode generation/readout transport and Ω_{odd} is role-odd generation/readout splitting. Charge-changing entries are forbidden, and role-changing entries are subleading or gapped.

Proof. The Pauli decomposition on role space is exact:

$$\Omega_W = \Omega_0 \otimes I + \Omega_1 \otimes \tau_1 + \Omega_2 \otimes \tau_2 + \Omega_3 \otimes \tau_3\ .$$

Axiom A removes every entry that would change gauge charge or representation through mass readout. Axiom B suppresses the τ_1 and τ_2 components at leading flavour-frame order. The surviving role-diagonal components are $\Omega_0 \otimes I$ and $\Omega_3 \otimes \tau_3$, identified respectively as common-mode and role-odd transport. ■

Corollary 1 — The CKM residue is a noncommutation residue

Under Axiom C,

$V\_CKM = \exp(-(\Omega_0 - \Omega\_q)/2) \cdot \exp((\Omega_0 + \Omega\_q)/2)$.

Baker–Campbell–Hausdorff gives

$$\log V\_CKM = \Omega\_q - \tfrac{1}{4}\cdot[\Omega_0 , \Omega\_q] + \text{(higher nested commutators)} .$$

The curvature residue is therefore not an extra interaction. It is the **failure of the role-even and role-odd projections of the same weak-doublet block to commute**.

Corollary 2 — The PMNS frame is a weak-commitment residual eigenframe

Under Axiom D, $T\_\nu = D_0\cdot I + \varepsilon\cdot M\_\nu$, so for every $\varepsilon > 0$, T_ν and M_ν share eigenvectors. The neutrino mixing frame is controlled by M_ν, not by the small scale ε.

This is the mathematical core of the CKM/PMNS contrast: **committed doublets produce small relative splitting; weak commitment permits a large residual eigenframe.**

---

5. Quark Committed Sector: Why the First Common Curvature is 2↔3

Use the inherited leading quark generator

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

with the hierarchy **a ≫ b ≫ c**.

The largest doorway is the Cabibbo channel 1↔2. The underproduced triangle information lives mainly in the 1↔3 sector. We seek the *smallest* common-mode curvature Ω_0 that can affect 1↔3 without spoiling 1↔2.

Let the elementary anti-Hermitian generation rotations be X_{12}, X_{23}, X_{13}. Their Lie brackets close schematically:

$$[X_{12} , X_{23}] \sim X_{13} , [X_{12} , X_{13}] \sim X_{23} , [X_{23} , X_{13}] \sim X_{12} .$$

The common-mode candidate is therefore *not* free to choose.

§5.1 Why not 1↔2? A common 1↔2 curvature lies in the already dominant Cabibbo plane. It is not Cabibbo-protected: it modifies the leading entry directly and risks converting the most successful inherited piece of the CKM generator into a fitted common-mode correction. It cannot be the first admissible curvature if the inherited Cabibbo scale is to be preserved.

§5.2 Why not direct 1↔3? A common 1↔3 curvature inserts the very triangle sector the clean exponential underproduces. That may repair a number, but it does not explain *why* the missing triangle should be a curvature residue — it collapses the mechanism into a hand-inserted correction to the target entry.

§5.3 Why 2↔3? A common 2↔3 curvature is Cabibbo-protected at first order: it does not directly alter the 1↔2 doorway. But through the bracket $[X_{23}, X_{12}] \sim X_{13}$ it produces the desired triangle residue using the largest available lever arm, a. This is exactly the required geometry — the missing 1↔3 effect is generated by **transporting a smaller common curvature through the Cabibbo doorway.**

Selection Lemma 2 — Cabibbo-protected curvature selection

Conditional on the weak-doublet admissibility rules, assume the leading role-odd quark split is Cabibbo-dominated; assume ***single-channel minimality****, i.e. that the first admissible common-mode curvature is carried by one elementary generation channel rather than a combination; and assume the common-mode curvature must (i) preserve the leading 1↔2 scale at first order, (ii) avoid direct insertion of a new 1↔3 term, and (iii) generate a 1↔3 residue through noncommutation with the role-odd split. Then the first admissible role-even common-mode curvature is the 2↔3 channel.*

Proof. Single-channel minimality restricts the candidate set to the three elementary generation channels {1↔2, 1↔3, 2↔3}. A 1↔2 common curvature violates (i). A 1↔3 common curvature violates (ii). The remaining elementary generation channel is 2↔3, and it satisfies (iii) because its commutator with the Cabibbo 1↔2 component lies in the 1↔3 plane. ■

The lemma selects within the minimal (single-channel) candidate set; it does not exclude a generic admissible Ω_0 that combines channels. Single-channel minimality is therefore a genuine extra premise, not a consequence of the other rules — a generic combination is admissible, and the claim is only that the *first* and simplest admissible curvature is the 2↔3 channel.

This is a *selection lemma*, not yet the full microscopic derivation of the projected H_cl block. It shows that once Cabibbo protection, no direct triangle insertion, and noncommutator generation are imposed, the 2↔3 common-mode channel is the first admissible candidate. The substrate calculation must still show that P_W H_cl P_W actually realises this channel.

§5.4 Exact commutator identity

Take the role-even 2↔3 curvature

$\Omega_0(z) = z \cdot E_{23} - z \cdot E_{32}$.*

Direct multiplication against Ω_q gives the exact identity

$[\Omega_0(z), \Omega\_q]_{13} = -a \cdot z$, so the leading BCH correction to the triangle entry is

$$\Delta_{13} = -\tfrac{1}{4}\cdot[\Omega_0 , \Omega\_q]_{13} = \tfrac{1}{4}\cdot a\cdot z .$$

This identity is the **algebraic spine** of the CKM curvature mechanism. The commutator value $-a\cdot z$ is convention-independent. The prefactor is *not* free: once Axiom C fixes the doublet split with the explicit factor of ½ in each exponential, Baker–Campbell–Hausdorff determines the second-order coefficient to be exactly −¼ (the halving of the split is what consumes the naïve factor of two). Direct numerical evaluation confirms the coefficient as −0.250000. The leading triangle residue is therefore $\tfrac{1}{4}\cdot a\cdot z$ under the stated convention; the alternative −½ would require dropping the ½ from the Axiom C exponentials, which would also shift the leading term away from Ω_q and is not the convention used here.

---

6. CKM Amplitude and Phase Targets from Projected Orbit Structure

Section 5 selects the *channel*. It does not yet fix the complex number z. The projected-Hamiltonian calculation must determine its amplitude and phase; this section isolates the two geometric inputs that calculation needs.

§6.1 Amplitude from root-normalised C_3 sharing

Suppose the role-even 2↔3 curvature inherits the committed 2↔3 transport norm b, but the common-mode component is distributed *democratically* over three Killing-orthonormal branches of a C_3 generation orbit:

$$Z = Z_1 + Z_2 + Z_3 , \langle Z\_i , Z\_j\rangle\_K = 0 \ (i \neq j) , \|Z\_i\|\_K = z .$$

Norm conservation gives

$$\|Z\|\_K^2 = 3\cdot z^2 = b^2 \Longrightarrow z = b/\sqrt{3} .$$

Theorem 3 — C_3-normalised curvature amplitude

If the role-even 2↔3 projected curvature is a democratic component of a three-branch Killing-orthonormal C_3 orbit carrying total norm b, then $z = b/\sqrt{3}$.

Status. Elementary once the orbit assumptions hold. The nontrivial substrate task is to prove that the relevant H_cl projection genuinely carries those orthonormal C_3 branches and that the inherited norm is the committed b norm.

§6.2 Phase from complement-half-transport

With complement reversal $C = e^{i\pi}$ and quark generation-loop orientation $O_q = e^{2\pi i/3}$, Axiom E states that the common-mode curvature phase is the half-transport

$$h_q^2 = C \cdot O_q = e^{i\pi} \cdot e^{2\pi i/3} = e^{5\pi i/3} .$$

The square root has **two branches**, $e^{5\pi i/6}$ and $e^{11\pi i/6}$, and these are *not* interchangeable: they are **antipodal** (z and $-z$; $e^{11\pi i/6} = -e^{5\pi i/6}$), and because they carry opposite-sign imaginary parts ($\mathrm{Im} = +\tfrac{1}{2}$ and $-\tfrac{1}{2}$) they select opposite CP signs. Note these two roots are *not* complex conjugates of each other — the conjugate $e^{7\pi i/6}$ is a third, distinct point that does not solve the square-root equation; branch selection (z vs $-z$) and CP conjugation (z vs $\bar{z}$) are different operations, and it is the branch that is at stake here. Choosing $e^{5\pi i/6}$ by calling it "the principal complement-continuous branch" is a *naming convention, not a derivation*. Structurally this is the same kind of discrete sign that §9 (rightly) converts into the computable invariant σ_W; it deserves the same treatment. Define the **phase-branch invariant** $\sigma_\varphi^q = \mathrm{sgn}(\mathrm{Im}\, z)$ as the sign that selects between the two antipodal roots from the projected H_{cl} holonomy. Because the antipodal roots have opposite-sign imaginary parts, this single sign both labels the branch and fixes the CP sign. Then

$$h_q = e^{5\pi i/6} \text{ if } \sigma_\varphi^q > 0, \text{ so } z_{C3} = (b/\sqrt{3})\cdot e^{5\pi i/6} .$$

Until σ_φ^q is computed, $\arg z_{C3} = 5\pi/6$ is one assertion away from its antipode $11\pi/6$.

Theorem 4 — CKM half-transport phase

If role-even common curvature is complement-half-transport against the C_3 generation orientation ***and the phase-branch invariant σ_φ^q selects the principal root****, then $\arg z_{C3} = 5\pi/6$.*

Status. Structurally powerful but conditional on two things, not one. Matrix algebra verifies the square-root identity; the substrate projection must confirm both (a) that common curvature is the half-transport rather than a direct holonomy, and (b) which of the two roots σ_φ^q selects. The branch choice is a named owed invariant, on the same footing as σ_W, not a convention.

---

7. Weak-Commitment Sector: The Neutrino Frame Theorem

The lepton doublet is asymmetric in commitment: the charged lepton is closure-norm anchored, while the neutrino has only weak overlap with the closure-norm direction. Hence the charged-lepton frame is **stiff** and the neutrino frame is **residual**. Write

$$T_\nu = D_0\cdot I + \varepsilon\cdot M_\nu , \quad \varepsilon > 0 .$$

Theorem 5 — Weak-commitment eigenframe invariance

For every $\varepsilon > 0$, T_ν and M_ν have identical eigenvectors. The PMNS neutrino frame is therefore independent of the weak-commitment scale ε.

Proof. Let $M_\nu v = m \cdot v$. Then

$$T_\nu v = (D_0 \cdot I + \varepsilon \cdot M_\nu) v = (D_0 + \varepsilon \cdot m) \cdot v ,$$

so every eigenvector of M_ν is an eigenvector of T_ν. Conversely, subtracting $D_0 \cdot I$ and dividing by ε maps the T_ν eigenproblem onto the M_ν eigenproblem. ■

This theorem is **exact**. It is the cleanest account of how a tiny neutrino mass scale and large neutrino mixing coexist inside the closure programme.

---

8. The PMNS Kernel from Residual Support and Weak Neutrality

The target weak-commitment Hermitian kernel is

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

up to the choice of whether the leading attachment breaking sits in the electron–tau or electron–muon entry. Its symmetric core is

$$M_0 = \begin{pmatrix} r_e & A & A \\ A & r_s & B \\ A & B & r_s \end{pmatrix}$$

Rotate to the μ–τ symmetric/antisymmetric basis,

$$\boldsymbol{\nu_+ = (\nu_\mu + \nu_\tau)/\sqrt{2} \, , \; \nu_- = (\nu_\mu - \nu_\tau)/\sqrt{2} \, .}$$

The ν_e–ν_+ solar block is

$$\begin{pmatrix} r_e & \sqrt{2} \cdot A \\ \sqrt{2} \cdot A & r_s + B \end{pmatrix}$$

so that

$$\boldsymbol{\tan 2\theta_{12} = (2\sqrt{2} \cdot A)/(r_s + B - r_e) = 2\sqrt{2}/\rho_\odot \, , \; \rho_\odot = (r_s + B - r_e)/A \, .}$$

§8.1 Solar ratio from root support

The projected residual cloud carries two relevant support counts: a **threefold** completion support for the electron-to-residual contrast, and a **twofold** coherent non-electron pair support. Comparing them by root-normalised norms gives

$$\rho_{\odot} = \sqrt{3} / \sqrt{2} = \sqrt{(3/2)} .$$

Theorem 6 — Solar root-support ratio

If the weak-commitment solar diagonal contrast is the root-norm comparison between threefold completion support and twofold coherent pair support, then $\rho_{\odot} = \sqrt{(3/2)}$.

This is the PMNS counterpart of the CKM amplitude theorem $z = b/\sqrt{3}$ — both are root-normalised support statements.

§8.2 Pair ratio from five-dimensional residual support

Let the residual weak-commitment support space be

$$R = C \oplus G \text{ , dim } C = 3 \text{ , dim } G = 2 \text{ , dim } R = 5 .$$

The electron-to-pair coupling A samples the base residual support; the internal μ–τ pair coupling B samples the same support *plus one return-continuation channel*. Hence

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

Theorem 7 — Pair ratio under single-return continuation

If the weak-commitment residual support has dimension five and the internal non-electron pair has exactly one additional admissible return continuation relative to the electron-to-pair link, then $B/A = 6/5$.

Status. The denominator is structurally inherited from dim R = 5. The numerator is now isolated as a precise continuation claim: the substrate projection must show *exactly one* additional return channel. Zero, two, or a root-normalised alternative would force a revision.

§8.3 Weak neutrality and diagonal μ–τ splitting

Weak neutrality suppresses first-order self-stiffness distinction between μ and τ in the neutrino kernel, so the leading breaking appears in an *attachment channel* rather than a diagonal split:

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

Theorem 8 — Weak-neutral diagonal suppression

If the weak-neutral residual projection is invariant under first-order μ–τ self-stiffness exchange, then $\delta = 0 + O(\beta^2)$.

This matters because a first-order diagonal split would make the atmospheric angle easy to fit but would weaken the architecture. The VERSF claim is sharper: the dominant breaking is **off-diagonal attachment leakage**.

§8.4 Weak-commitment phase

For neutrinos the sector orientation is weak-neutral quadrature, $O_\nu = e^{i\pi/2}$ — itself an asserted input that must be derived independently (see Axiom E). Complement-half-transport gives

$$h_\nu^2 = C \cdot O_\nu = e^{i\pi} \cdot e^{i\pi/2} = e^{3\pi i/2}\ ,\ h_\nu = e^{3\pi i/4}\ ,\ \text{so}\ \psi = 3\pi/4\ .$$

As in the quark case, the square root carries a genuine branch choice between **antipodal** roots: $e^{3\pi i/4}$ and $e^{7\pi i/4} = -e^{3\pi i/4}$, with opposite-sign imaginary parts ($\mathrm{Im} = +\sqrt{2}/2$ and $-\sqrt{2}/2$) selecting opposite CP signs. (The conjugate $e^{5\pi i/4}$ is again a distinct third point, not the second root.) This is fixed by a neutrino phase-branch invariant $\sigma_\varphi^\nu = \mathrm{sgn}(\mathrm{Im}\ h_\nu)$ on the same footing as σ_φ^q and σ_W, not by the word "principal"; as before, the single sign both labels the branch and fixes the CP sign. Until σ_φ^ν is computed, $\psi = 3\pi/4$ is one assertion away from its antipode $7\pi/4$.

Theorem 9 — Shared phase grammar

If both the quark common curvature and the neutrino weak-commitment leakage are complement-half-transports, ***and the phase-branch invariants*** $\sigma_\varphi^q, \sigma_\varphi^\nu$ ***select the principal roots****, then*

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

The two phases are not independent choices: they are the two sectoral **square roots of complement reversal** — each carrying one discrete branch sign that the projection must select.

§8.5 Leakage amplitude from the K = 7-descended support trace

The benchmark leakage amplitude is

$\beta = \sqrt{3}/20$.*

This quantity is deliberately kept separate from the better-established solar and pair ratios; its numerator and denominator have different origins, and — importantly — they enter under *different* counting conventions. The numerator is rooted; the denominator is linear. This is not a convenience: the governing principle is that **a source amplitude enters as an RMS norm over the branches it is distributed across (hence a root), while a support trace counts the independent channels the amplitude spreads across (hence a linear count).** Numerator and denominator are different kinds of quantity, so it would be wrong to apply one convention to both.

The numerator √3 is the root-normalised threefold leakage *amplitude*: a source spread over three branches contributes as the RMS norm $\sqrt{1^2+1^2+1^2} = \sqrt{3}$, the same amplitude convention as the CKM curvature ($b/\sqrt{3}$) and the solar ratio ($\sqrt{3/2}$). The denominator 20 is a *support trace* — a count of independent admissible channels — and counts therefore enter linearly, not under a root. The asymmetry is forced by the distinction between an amplitude and a channel count, and it is precisely what lands the reactor angle: the fully-linear form $3/20 \approx 0.15$ gives $\theta_{13} \approx 13°$, and the fully-rooted form $\sqrt{3}/\sqrt{20} \approx 0.39$ gives $\theta_{13} \approx 22°$; only the mixed rooted-numerator / linear-denominator form $\sqrt{3}/20 \approx 0.087$ gives $\theta_{13} \approx 8.6°$.

The denominator 20 is the support trace of the weak-commitment leakage projector, and it descends from $K = 7$ (§2.2). The leakage projection removes two boundary/readout-anchored directions from the sevenfold closure carrier, leaving $\dim R = K - 2 = 5$, which then tensors with the fourfold weak-attachment orientation fibre $\dim Q = 4 = 2 \times 2$ (weak-neutral quadrature × closure two-cycle binary):

$$\Pi\_leak \simeq \Pi\_R \otimes \Pi\_Q \,, \quad Tr\_support(\Pi\_leak) = (K - 2) \times 4 = 20 \,,$$

so that

$$\beta = \sqrt{3} \,/\, Tr\_support(\Pi\_leak) = \sqrt{3}/20 \,.$$

The honest status is therefore sharper than before: the numerator is structurally motivated by root-normalised threefold leakage, and the denominator is a **K = 7-descended support-trace target** rather than a free number. Within $\dim Q = 4$, the two factors are not equally secure: the quadrature factor of two is secure (a complex entry has two real parts), but the **second factor of two is the genuine open premise** — and it is doubly constrained (§2.2): the substrate projection must exhibit a binary that is two-dimensional, orthogonal to the quadrature plane, and *not* the Axiom-B-gapped role-changing channel, then confirm it has exactly multiplicity two (the candidate is the closure two-cycle binary). Multiplicity one or three would fix the denominator at 10 or 30, and hence the reactor angle at ≈ 14° or ≈ 6° rather than ≈ 8.6° — so this single factor of two is what pins θ_{13} to its measured value. The remaining task is to compute the weak-commitment attachment block of P_W H_cl P_W and verify that the leakage support really factors as $R \otimes Q$ with dimensions 5 and 4. If the projection fails to produce this support trace, the reactor-angle magnitude remains benchmarked rather than derived. (Replacing $\sqrt{3}/20$ by 3/20 would define a different, much larger leakage benchmark — raising θ_{13} from ≈ 8.6° to ≈ 13° — not the one inherited from the weak-commitment closure kernel.)

---

9. The Octant-Channel Invariant σ_W

The previous formulation left one ambiguity: the same magnitude of attachment breaking can sit in either the electron–tau or electron–muon entry. The two cases are $\mu \leftrightarrow \tau$ reflections — they preserve θ_{12}, θ_{13}, and J_lep, but flip the atmospheric octant and the CP sign. This paper replaces that ambiguity with **one substrate invariant**.

Define the projected electron-attachment strengths

$A_{e\mu} = \|P_e H_W P_\mu\|^2$, $A_{e\tau} = \|P_e H_W P_\tau\|^2$,

and the weak-attachment orientation

$\sigma_W = \mathrm{sgn}(A_{e\tau} - A_{e\mu})$.

Then:

- **$\sigma_W > 0$** → leading breaking is electron–tau → **upper-octant** branch.
- **$\sigma_W < 0$** → leading breaking is electron–muon → **lower-octant** branch.
- **$\sigma_W = 0$** → no first-order octant selection; the benchmark leakage cannot predict an octant.

The leakage trace and the octant sign are distinct but connected. The K = 7-descended support trace fixes the *magnitude* of the leading attachment breaking through $\beta = \sqrt{3/20}$; the sign invariant σ_W fixes *which* electron-attachment channel receives that breaking. So β controls the size of the atmospheric departure from maximality while σ_W controls the branch, and the projection calculation must determine two separate pieces of information: the support trace $\mathrm{Tr\_support}(\Pi_{leak}) = 20$ and the attachment orientation sign σ_W.

Theorem 10 — Octant branch from weak-attachment orientation

Assume weak neutrality suppresses first-order diagonal μ–τ splitting and the leading atmospheric breaking is a single electron-attachment leakage. Then the atmospheric octant and leptonic CP sign are determined by σ_W: the magnitude of the octant departure is fixed by β, the sign of the departure by σ_W.

Proof sketch. The μ↔τ reflection maps the electron–tau leakage branch to the electron–muon branch, sending

$$\theta_{23} \to 90° - \theta_{23}$$

and conjugating the CP orientation, while preserving the solar angle, reactor angle, and J_{lep}. The only missing information is which attachment channel the projection selects — exactly the sign of $A_{e\tau} - A_{e\mu}$. ■

Status. A real sharpening. The programme no longer says vaguely that the octant is open; it says the octant is the **sign of one computable projected-Hamiltonian quantity**. The next substrate calculation must compute σ_W.

---

10. The Finite Computation Now Required

The weak-doublet flavour problem is now reduced to a finite list of projected-Hamiltonian matrix elements and support traces.

§10.1 Required quark-sector computations

Compute the role-odd block $\Omega_q = P_{odd}\, \Omega_W\, P_{odd}$ and check whether its leading entries reproduce or constrain **a, b, $c \cdot e^{i\varphi}$**. Compute the role-even block $\Omega_0 = P_{even}\, \Omega_W\, P_{even}$ and check whether the first nontrivial common curvature is

$$\Omega_0(z) = z \cdot E_{23} - z^* \cdot E_{32} ,$$

then test **$z = b/\sqrt{3}$** and **$\arg z = 5\pi/6$**. The phase test is two-part: it requires both the half-transport structure and the sign of the phase-branch invariant **$\sigma_\varphi^q = \mathrm{sgn}(\mathrm{Im}\, z)$** that selects which of the two antipodal roots ($5\pi/6$ or $11\pi/6$) the projection realises. If the projected block gives a different leading common-mode channel, the CKM curvature mechanism is not derived.

§10.2 Required neutrino-sector computations

Compute the weak-commitment readout $T_\nu = P_\nu\, H_W\, P_\nu$ and test the form $T_\nu = D_0 \cdot I + \varepsilon \cdot M_\nu$. The weak phase ψ likewise requires the orientation O_ν *and* the neutrino phase-branch sign **σ_φ^ν** to be computed from H_{cl} rather than asserted. Then compute the kernel ratios

$$\rho_\odot = (r_s + B - r_e)/A\ ,\ B/A\ ,\ \delta\ ,\ \psi\ ,\ \beta\ ,$$

against the targets

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

§10.3 Required octant computation

Compute

$$\sigma_W = \mathrm{sgn}\left(\|P_e\, H_W\, P_\tau\|^2 - \|P_e\, H_W\, P_\mu\|^2 \right)$$

to determine whether the projected Hamiltonian selects the upper or lower atmospheric branch.

§10.4 Required support-trace computation

The unresolved leakage amplitude requires a projected support-trace theorem of factored form

$$\Pi_{leak} \simeq \Pi_R \otimes \Pi_Q\ ,\ \dim R = K - 2 = 5\ ,\ \dim Q = 4\ ,\ \mathrm{Tr}_{support}(\Pi_{leak}) = 5 \times 4 = 20\ .$$

The $K - 2$ factor records the residual support left after two boundary/readout-anchored directions are removed from the $K = 7$ closure carrier; the fourfold factor records the weak-attachment orientation fibre (weak-neutral quadrature × closure two-cycle binary, the latter required to be orthogonal to the quadrature plane and distinct from the Axiom-B-gapped role-changing channel

— see §2.2). If this factorisation is verified, the leakage amplitude becomes $\beta = \sqrt{3}$ / Tr_support(Π_leak) = $\sqrt{3}/20$. If no such factorisation is found, β remains a successful benchmark but not a derived output of the weak-doublet projection.

---

11. Unified Theorem Package

- **Theorem A — Weak-doublet projection.** If H_cl is the inherited su(8)-restricted admissibility Hessian and P_W is the generation/role/readout projector, then H_W = P_W H_cl P_W is the unique weak-doublet closure Hamiltonian at this level of the programme.
- **Theorem B — Role decomposition.** If role-changing components are gapped, then Ω_W = Ω_even ⊗ I + Ω_odd ⊗ τ_3 + O(Ω_mix).
- **Theorem C — CKM noncommutation residue.** In the committed quark regime, log V_CKM = Ω_q − ¼·[Ω_0, Ω_q] + ⋯ ; for a pure 2↔3 common curvature, $[\Omega_0, \Omega_q]_{13}$ = −a·z and Δ_{13} = ¼·a·z.
- **Theorem D — Cabibbo-protected channel selection.** Under the stated admissibility rules, Cabibbo protection, and no direct triangle insertion, the first admissible common curvature is the role-even 2↔3 channel.
- **Theorem E — Root-normalised C_3 amplitude.** If the role-even curvature is a democratic Killing-orthonormal C_3 component with total inherited norm b, then $z = b/\sqrt{3}$.
- **Theorem F — Shared half-transport phases.** If common curvature and weak leakage are complement-half-transports, *and the phase-branch invariants σ_φ^q, σ_φ^ν select the principal roots*, then arg z = 5π/6 and ψ = 3π/4. Each phase carries one discrete branch sign the projection must select.
- **Theorem G — Weak-commitment eigenframe.** If T_ν = D_0·I + ε·M_ν, then the neutrino frame is the eigenframe of M_ν for every ε > 0.
- **Theorem H — PMNS residual-support ratios.** If the solar contrast is a root-normalised 3:2 support comparison, then ρ_⊙ = $\sqrt{3/2}$. If the residual support has dimension five and the internal pair has one return continuation, then B/A = 6/5.
- **Theorem H2 — K = 7 leakage support trace.** If the weak-commitment leakage projector factors as Π_leak ≃ Π_R ⊗ Π_Q, with residual support dim R = K − 2 = 5 and fourfold orientation fibre dim Q = 4 = (quadrature 2) × (closure two-cycle binary 2), where the second binary is two-dimensional, orthogonal to the quadrature plane, and distinct from the Axiom-B-gapped role-changing channel, then Tr_support(Π_leak) = 20. If the leakage numerator is the root-normalised threefold source $\sqrt{3}$, then $\beta = \sqrt{3}/20$. The numerator enters as an RMS amplitude over three branches (hence rooted, as for $b/\sqrt{3}$ and $\sqrt{3/2}$); the denominator is a support trace counting independent channels (hence linear) — the two conventions differ because an amplitude and a channel count are different kinds of quantity. The second factor of two in that trace is the load-bearing open premise.
- **Theorem I — Octant branch.** If first-order diagonal μ–τ breaking is suppressed and the leading atmospheric breaking is electron-attachment leakage, then the octant and CP sign are fixed by σ_W = sgn(A_eτ − A_eμ).

---

12. Falsification Conditions

- If P_W H_cl P_W cannot be defined without fitting the observed flavour matrices, the programme collapses back to phenomenology.
- If the projected readout violates charge-blindness, the construction fails immediately.
- If Ω_mix is not gapped or subleading, clean weak-role flavour frames are not obtained.
- If the first role-even common-mode curvature is not 2↔3, the CKM C_3 curvature story is not derived.
- If the role-even 2↔3 amplitude is not close to b/√3, the root-normalised C_3 amplitude theorem is not realised.
- If the projected common-mode phase is not on the complement-half-transport branch, or the phase-branch invariants σ_φ^q, σ_φ^ν select the antipodal (opposite-CP) roots, the shared CKM/PMNS phase grammar (and its CP-sign predictions) fails.
- If the neutrino block is not of the form D_0·I + ε·M_ν, the weak-commitment explanation of large PMNS mixing fails.
- If the weak-commitment support calculation does not give ρ_⊙ = √(3/2), the solar-ratio claim fails.
- If the residual continuation count does not give B/A = 6/5, the pair-ratio rule must be revised.
- If first-order diagonal μ–τ breaking appears, the attachment-breaking architecture loses predictive power.
- If no K = 7-descended support-trace invariant yields β near √3/20, the reactor-angle magnitude remains fitted rather than derived.
- If the weak-commitment projection does not factor the leakage support as R ⊗ Q with dim R = K − 2 and dim Q = 4, the proposed denominator 20 fails; the framework may retain the weak-commitment kernel, but the reactor-angle amplitude would require a different leakage theorem.
- If the second factor of dim Q = 4 cannot be exhibited as a binary that is simultaneously two-dimensional, orthogonal to the Re/Im quadrature plane, and distinct from the Axiom-B-gapped role-changing channel, the support trace is not 20 and the denominator is unexplained. In particular, identifying that factor with the (ν, e_L) role complement is inconsistent with Axiom B and is not available.
- The (K − 2) × (2 × 2) support-trace rule has not yet reproduced a *second*, independently-known denominator out of sample. Until it predicts a different suppression count without retuning, a post-hoc reading of the value 20 remains available; the charged-lepton threshold-compression precedent establishes the genre (suppression-as-trace-count) but does not by itself close this concern.
- If σ_W selects the atmospheric branch opposite to stable experimental data, the weak-attachment channel rule is falsified.
- If CKM and PMNS require unrelated projection rules, the unified closure-Hamiltonian claim fails even if separate fits survive.

---

13. Conclusion

The previous projected-Hamiltonian paper identified the right object, $H_W = P_W H_{cl} P_W$. This paper turns that object into a **theorem-level programme**. The weak-doublet projection is not arbitrary: charge-blindness, role-gap admissibility, commitment class, generation support, and complement reversal constrain the projected block strongly enough to *select the first CKM curvature channel* and to *reduce the PMNS octant ambiguity to a single substrate sign*.

The quark result is the cleanest new selection. In the committed regime the role-odd block gives the leading CKM generator. The first common-mode correction that protects Cabibbo, avoids direct triangle insertion, and still repairs the triangle through noncommutation is the 2↔3 role-even curvature, whose commutator with the Cabibbo doorway is

$$[\Omega_0 , \Omega_q]_{13} = -a\cdot z .$$

If the projected orbit is C_3-democratic and Killing-orthonormal, $z = b/\sqrt{3}$; if common curvature is complement-half-transport, $\arg z = 5\pi/6$.

The neutrino result is the cleanest structural explanation. In the weak-commitment regime $T_\nu = D_0\cdot I + \varepsilon\cdot M_\nu$, so the small neutrino scale does not suppress the mixing frame — the PMNS frame is the eigenframe of the residual kernel. Root-normalised support gives the solar ratio; five-dimensional residual support gives the denominator of the pair ratio; weak neutrality suppresses diagonal μ–τ breaking; complement-half-transport gives the weak phase.

The atmospheric octant is no longer a loose ambiguity. It is the sign of

$$\sigma_W = \mathrm{sgn}(A_{e\tau} - A_{e\mu}) .$$

The status should be stated carefully. This paper does not yet compute the microscopic entries of H_{cl}. It derives the admissible projected shape that those entries must take if the unified flavour programme is correct. That is why the result is powerful but still conditional: it turns CKM curvature, PMNS weak commitment, and the atmospheric octant into concrete projected-Hamiltonian tests rather than completed substrate calculations.

The leakage amplitude now has a sharper address. The denominator in $\beta = \sqrt{3}/20$ is no longer an isolated numerical target: it is proposed to be the $K = 7$-descended support trace $\mathrm{Tr\_support}(\Pi_{leak}) = (K - 2) \times 4 = 20$. This moves the reactor-angle part of the PMNS kernel from a floating benchmark toward a specific projected-Hamiltonian test — the remaining task is not to motivate the number again, but to compute the weak-commitment attachment block and check whether its support really factors into the five-dimensional residual support and the fourfold weak-attachment orientation fibre.

The paper therefore converts the next stage of the Standard Model flavour programme into a **finite computation**: evaluate the relevant blocks and support traces of $P_W H_{cl} P_W$. If the computation returns the target blocks, VERSF gains a genuine unified flavour mechanism — small CKM from committed weak-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 and locally**: the failed invariant names exactly which part of the closure programme must be revised.

---

14. Numerical and Algebraic Note

CKM algebra

With $\Omega_0(z) = z \cdot E_{23} - z^* \cdot E_{32}$ and

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

direct computation gives the exact identity $[\Omega_0(z), \Omega\_q]_{13} = -a \cdot z$, hence the leading BCH correction $\Delta_{13} = ¼ \cdot a \cdot z$ (the −¼ coefficient is fixed by the halved Axiom C split and confirmed numerically to −0.250000). This identity is independent of the numerical choices for a, b, c, φ. The benchmark curvature target is

$z\_C3 = (b/\sqrt{3}) \cdot e\hat{}(5\pi i/6)$.

PMNS kernel

In the scale convention A = 1, r_e = 0, the benchmark weak-commitment kernel uses

$B = 6/5$, $r\_s = \sqrt{(3/2)} - 6/5$, $\delta = 0$, $\psi = 3\pi/4$, $\beta = \sqrt{3}/20$.

With $\beta = \sqrt{3}/20$, the benchmark weak-commitment kernel preserves the intended PMNS scale: $\theta_{12} \approx 33.4°$, $\theta_{13} \approx 8.6°$, an atmospheric deviation $|\theta_{23} - 45°| \approx 3.4°$ ($\theta_{23} \approx 48.4°$ on the electron–tau branch, 41.6° on the electron–muon reflection), and $|J\_lep| \approx 2.45 \times 10^{-2}$, with the branch selected by σ_W. The denominator 20 is interpreted as the support trace Tr_support(Π_leak) = (K − 2) × 4, not as a fitted numerical denominator. Using β = 3/20 would define a different and much larger leakage benchmark — raising θ_{13} to ≈ 13° — not the one inherited from the weak-commitment closure-kernel paper.

The two attachment branches are related by μ↔τ reflection. They preserve θ_{12}, θ_{13}, and J_lep, while reflecting $\theta_{23} \rightarrow 90° - \theta_{23}$ and reversing the CP orientation. The branch is selected by σ_W.

---

15. References and Source Map

- *The Projected Weak-Doublet Closure Hamiltonian in VERSF* — establishes $H_W = P_W H_{cl} P_W$, the committed/weak-commitment architecture, the CKM curvature target, and the PMNS octant-channel premise.
- *Substrate Dynamics and the Higgs Ratio* — identifies H_{cl} as the su(8)-restricted admissibility Hessian and supplies the refinement update $\exp(i \cdot \varepsilon \cdot H_{cl})$.
- *Closure-Norm Condensation and Electroweak Symmetry Breaking in VERSF* — gives closure-norm condensation and the weak neutrino overlap with the closure-norm direction.
- *The Yukawa Hierarchy Theorem* — gives charge-blindness of the mass/readout operator.
- *The Structural Origin of Yukawa Operators in VERSF* — reads Yukawa entries as closure-depth overlap amplitudes.
- *The Yukawa Operator from Completion-Channel Misalignment in VERSF* — supplies the sector-frame mismatch architecture.
- *The Electroweak Flavour-Frame Operator in VERSF* — develops weak-role decomposition and the weak-commitment frame theorem.
- *The CKM Curvature Residue in VERSF* — gives the z_{C3} target and the exact CKM numerical audit.
- *The Weak-Commitment Closure Kernel in VERSF* — gives the PMNS kernel, solar ratio, pair ratio, leakage phase, and atmospheric branch structure.
- *The Weak-Commitment Neutrino Operator in VERSF* — develops stiffness-gap collapse and large PMNS mixing from weak commitment.
- *Deriving Flavour Mixing from Closure Geometry* — supplies the residual support space $R = C \oplus G$ with dim $C = 3$, dim $G = 2$, dim $R = 5$.
- *The Weak-Commitment Leakage Trace in VERSF* — derives $\beta = \sqrt{3}/20$ as the $K = 7$-descended support-trace amplitude and audits the PMNS kernel ($\theta_{12} \approx 33.4°$, $\theta_{13} \approx 8.6°$, $|\theta_{23} - 45°| \approx 3.4°$, $|J_{lep}| \approx 2.45 \times 10^{-2}$).
- *A No-Go Theorem for Non-Simplicial Relational Substrates* — derives $K = 7$ as the minimal admissible encoding of the triangular hinge state via the Hamming-bound / Fano-plane route.
- *The K = 7 Closure Symmetry* — supplies the spatial–temporal closure-counting framework, including $N_{loop} = 14 = 2K$ and the 11 + 2 + 1 channel decomposition.
- *The K = 7 Wilson Limit* — uses $N_{loop} = 14$ and $\beta_{K=7}$ as the $K = 7$ closure-counting scale for electromagnetic transport.
- *Half-Integer Spin as Closure Two-Cycle Structure in VERSF* — supplies the closure two-cycle (the topological Z_2 double-cover sign under 2π rotation) proposed as the second factor of the weak-attachment orientation fibre. As a topological Z_2 it is categorically distinct from a U(1) quadrature phase and from a weak-isospin τ flip; whether it resides in the closure/holonomy sector rather than the Axiom-B-gapped weak-role sector is the check the leakage support trace still owes (§2.2).
- *Threshold Compression in the Charged-Lepton Spectrum* — precedent for boundary/interior support reduction and for treating suppression factors as trace/count targets rather than fitted constants.

Appendix A — Minimal Calculation Checklist

The next technical note should evaluate or bound the following quantities from the substrate Hessian:

$P_{even}\ P_W\ H_{cl}\ P_W\ P_{even}$, $P_{odd}\ P_W\ H_{cl}\ P_W\ P_{odd}$, $P_e\ P_W\ H_{cl}\ P_W\ P_\mu$, $P_e\ P_W\ H_{cl}\ P_W\ P_\tau$,

and the leakage support trace **$\mathrm{Tr}_{support}(\Pi_{leak})$**.

The decisive outputs are

$\Omega_0^{(23)}$, z_{C3} , M_ν , σ_W , σ_φ^q , σ_φ^ν , Π_{leak} , $\mathrm{Tr}_{support}(\Pi_{leak})$, $\beta = \sqrt{3/20}$.

The block outputs are $\Omega_0^{(23)}$, z_{C3}, M_ν; the discrete signs are σ_W (octant) and σ_φ^q, σ_φ^ν (the CKM and PMNS phase-root branches). The leakage calculation now splits into two independent checks: the projected Hamiltonian must determine the attachment branch sign σ_W (which fixes the octant), and it must verify the K = 7-descended support trace $\mathrm{Tr}_{support}(\Pi_{leak}) = (K - 2) \times 4 = 20$ (which fixes the reactor-angle amplitude $\beta = \sqrt{3/20}$). These are separate outputs — the octant sign and the support trace answer different questions.

These objects are the current frontier of the VERSF Standard Model flavour programme.