The Weak-Commitment Leakage Trace in VERSF

Deriving β = √3 / 20, the Reactor Angle, and the Atmospheric Octant Sign from the Projected H_cl Attachment Block

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

Successor to *The Weak-Doublet Projection Theorem in VERSF*, *The Projected Weak-Doublet Closure Hamiltonian 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 substrate-dynamics result identifying **H_cl** as the su(8)-restricted admissibility Hessian.

Summary for the General Reader

Matter comes in three "generations" — three copies of the same few particles, each copy heavier than the last. Particles don't stay neatly inside their own generation; they *mix*, meaning a particle of one generation behaves a little like its cousins in the others. Neutrinos — the lightest, most ghostly of these particles — mix in a remarkably wide-open way, far more than their heavier relatives the quarks. Describing that mixing takes a handful of angles, measured in experiments.

Two of those neutrino angles are the subject here. One of them sits close to 45° — close to a perfectly balanced arrangement between two of the generations. If the balance were *exact*, two things would follow automatically: that angle would be exactly 45°, and a second angle (the one reactor experiments measure) would be exactly zero. Experiments see neither. The first angle is tilted slightly off balance, and the second is small but definitely not zero. So nature must contain a small imbalance — a tiny "leak" that breaks the perfect symmetry by just the right amount. The whole job of this paper is to pin down the size of that leak, a single number called β.

The previous paper could show the leak *had* to be there, but not how big — it put the size in by hand, matched to the data. This paper tries instead to *calculate* the size, by counting structural features of the underlying theory rather than fitting it.

The recipe is a fraction: the strength of the leak equals *how much leaks* divided by *how many channels it spreads across*.

The top of the fraction works out to √3 (about 1.73). Why √3 and not plain 3 is worth a sentence: three separate contributions feed the leak, but they combine the way the two short sides of a right-angled triangle make the long one — by squares, not by simple addition. Three equal

contributions of size one give $\sqrt{1 + 1 + 1} = \sqrt{3}$, not 3. This part of the answer is forced by the rule; it is not a choice.

The bottom of the fraction works out to 20. It starts from a number — seven — that turns up all over this framework (it is the same seven that, elsewhere in the programme, sets the strength of electromagnetism). Two of the seven are already spoken for, leaving five; and because the leak can also point in different orientations, that five is multiplied by four, giving twenty.

Put together:

$\beta = \sqrt{3} / 20 \approx 0.087$.

This is the number the neutrino measurements were asking for. Fed into the model, the two angles come out about right: the reactor angle near 8.6°, and the tilt of the near-45° angle about 3.4° — both close to what experiments find.

The paper also tidies up a long-standing loose end. Experiments can tell that the near-45° angle is tilted, but not whether it tips a little *above* 45° or a little *below*. That used to be simply open. Here it becomes a single yes/no — one quantity whose sign can be computed decides which way it leans. The leak size β sets *how far* off balance the angle is; this sign sets *which side*.

Now the honest part, which matters as much as the result. The top of the fraction ($\sqrt{3}$) is genuinely forced by the counting rule. The bottom (20) is not — at least not yet. It is built on a solid starting number, the seven, but every step after that — subtracting two, multiplying by four, and a further bookkeeping choice about what counts as a separate channel — is currently settled by *which answer matches the measured angle*, not by the rules themselves. Change any of those steps and the predicted reactor angle moves to a value experiments rule out, so for now the data is doing the choosing. And the whole "divide how-much by how-many" only means anything if every channel genuinely counts equally; if it turns out they don't, the bottom number could be adjusted to almost anything, which would quietly drain the result of its power to be proven wrong.

So the fair summary is this: a promising recipe that lands the right number, with one half on firm ground and the other half still leaning on the very measurement it is trying to explain. The paper is deliberately blunt about that, and it names the one specific calculation that would settle it — showing that the underlying theory really does weight every channel equally and really does produce the count of twenty. That calculation, not more argument, is what would turn this from a strong benchmark into a genuine prediction.

Contents

Abstract

The weak-doublet projection theorem reduces the remaining Standard Model flavour problem to the projected sector block

$$\mathbf{H_W = P_W\, H_{cl}\, P_W,}$$

with H_{cl} the inherited su(8)-restricted admissibility Hessian and P_W the generation/role/readout projector. The theorem leaves one especially exposed PMNS target: the weak-commitment leakage amplitude β, which controls the reactor angle, the magnitude of the atmospheric deviation from maximality, and the leading electron-attachment breaking. This paper derives the benchmark

$$\boldsymbol{\beta = \sqrt{3} / 20}$$

from a support-trace calculation inside the projected weak-commitment attachment block.

The weak-commitment neutrino operator has the scale–shape form $\mathbf{T_\nu = D_0\, I + \varepsilon\, M_\nu}$, so the PMNS frame is the eigenframe of the residual Hermitian kernel M_ν for every $\varepsilon > 0$. In the μ–τ-symmetric core, root-normalised residual support gives $\rho_\odot = \sqrt{3/2}$, and the five-dimensional residual support $R = C \oplus G$ supplies the denominator of the pair ratio $B / A = 6 / 5$. The present paper extends the same support calculus to the leading μ–τ-breaking attachment leakage under one explicit normalisation convention (premise **WL-5**): the Killing norm is scaled so that each

admissible support slot carries unit weight, making a coherent leakage *norm* and a support *count* commensurable in a single dimensionless ratio.

Under that convention, the leakage numerator is a coherent three-completion residue whose Killing norm is $\sqrt{3}$ (premise **WL-1**), and the leakage denominator is the support trace of the weak-attachment fibre

$$\boldsymbol{\Pi_{leak} \simeq \Pi_R \otimes \Pi_Q,\ \mathrm{Tr}\ \Pi_R = 5,\ \mathrm{Tr}\ \Pi_Q = 4,}$$

with Π_R a residual support that descends from the $K = 7$ closure carrier minus two boundary/readout-anchored directions, $\dim R = K - 2 = 5$ (premise **WL-2**), Π_Q a fourfold weak-role orientation fibre (premise **WL-3**), and the factorisation itself premise **WL-4**. Hence

$$\boldsymbol{\mathrm{Tr}_{support}(\Pi_{leak}) = 5 \times 4 = 20,\ \beta = \sqrt{3} / 20.}$$

With $B / A = 6 / 5$, $\rho_\odot = \sqrt{(3/2)}$, $\psi = 3\pi/4$, and $\beta = \sqrt{3}/20$, exact Hermitian diagonalisation gives the intended PMNS scale: $\theta_{12} \approx 33.4°$, $\theta_{13} \approx 8.6°$, $|\theta_{23} - 45°| \approx 3.4°$, $|J_{lep}| \approx 2.45 \times 10^{-2}$ (reproduced independently; see §13). The octant remains the sign of the projected weak-attachment orientation,

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

so the paper derives the *magnitude* of the leading atmospheric breaking while isolating the octant as a computable substrate sign. The result is conditional but sharp: if the projected H_cl block does not realise the 5 × 4 leakage trace under WL-1…WL-5, β remains fitted rather than derived.

Notation

Symbol	Meaning
H_cl	su(8)-restricted closure Hamiltonian; admissibility Hessian of the substrate free-energy functional
P_W	weak-doublet projector; P_W = P_Y P_R P_C
H_W	projected weak-doublet Hamiltonian, H_W = P_W H_cl P_W
P_C	generation-completion projector
P_R	weak-role projector
P_Y	mass / readout projector
P_e, P_μ, P_τ	charged-lepton-labelled weak-commitment attachment projectors after charged-lepton anchoring
T_ν	weak-commitment neutrino readout operator
M_ν	residual Hermitian weak-commitment kernel
A	electron-to-non-electron symmetric attachment strength

Symbol	Meaning
B	internal μ–τ pair coupling
ρ_⊙	solar diagonal contrast ratio, (r_s + B − r_e) / A
ψ	weak-commitment leakage phase, target 3π/4
β	leading weak-commitment leakage amplitude
L_coh	coherent leakage residue surviving weak-neutral projection
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, organised as R = C ⊕ G
Q	weak-attachment orientation fibre, dim Q = 4
Π_R, Π_Q	support projectors onto R and Q
Π_leak	leakage support projector
Tr_support	support trace (dimension count) under the unit-slot normalisation
σ_W	weak-attachment orientation sign selecting the atmospheric octant
‖·‖_K	Killing norm on projected closure directions, unit-slot normalised (WL-5)

---

0. Predictive-Content Ledger

This paper addresses one target left exposed by the weak-doublet projection theorem, and sharpens one linked target.

- **LT-1 — Derive the weak-commitment leakage amplitude β = √3/20 from a support trace rather than a fit.**
- **LT-2 — Keep the atmospheric octant as the sign of one projected weak-attachment invariant σ_W.**

The derivation rests on five named structural premises, each of which is either inherited or owed to the substrate projection:

- **WL-1 — Threefold orthonormal leakage residue.** The coherent leakage numerator is the sum of three Killing-orthonormal completion branches of unit norm.
- **WL-2 — K = 7-descended residual support.** The leakage denominator contains the five-dimensional residual support dim R = K − 2 = 7 − 2 = 5: the sevenfold closure carrier minus two boundary/readout-anchored directions (the committed charged-lepton frame and the weak-doublet/readout complement). This is the one denominator factor with a closure pedigree. Two sub-claims ride on it — that *exactly* two directions anchor (not one, not three), and that the five survivors organise as R = C ⊕ G with dim C = 3, dim G = 2.
- **WL-3 — Fourfold weak-attachment orientation fibre.** The off-diagonal leakage is distributed over a weak-role orientation fibre Q with dim Q = 4.

- **WL-4 — Support factorisation.** The leakage support projector factors as $\Pi\_leak \simeq \Pi\_R \otimes \Pi\_Q$.
- **WL-5 — Commensurable unit-slot normalisation.** The Killing norm $\|\cdot\|\_K$ is scaled so that each admissible support slot carries unit weight; this is what makes the ratio (norm)/(count) a dimensionless amplitude.

Object	Claim	Status
$T\_\nu = D_0 I + \varepsilon M\_\nu$	weak-commitment scale–shape form	inherited exact frame theorem
$\rho\_\odot = \sqrt{(3/2)}$	solar ratio from root 3 : 2 residual support	inherited conditional theorem
B / A = 6 / 5	pair ratio from dim R = 5 plus one return continuation	inherited conditional theorem
$\psi = 3\pi/4$	weak phase from complement-half-transport	inherited conditional theorem
WL-5 normalisation	unit-slot Killing norm makes norm/count commensurable	**stated premise (new)**
β numerator = $\sqrt{3}$	WL-1: root-normalised threefold leakage	derived here under WL-1
dim R = K − 2 = 5	residual support as K = 7 carrier minus two anchored directions	K = 7-descended; bridges to $R = C \oplus G$
β denominator = 20	(K − 2) × 4: K = 7-descended residual support × orientation fibre	the 5 is K = 7-descended (WL-2); the 4 is not (WL-3, open)
$\beta = \sqrt{3} / 20$	benchmark leakage amplitude	**main result, conditional on WL-1…WL-5**
θ_{13}	reactor angle from β	numerical audit (verified §13)
$\lvert\theta_{23} - 45°\rvert$	atmospheric-deviation magnitude from β	numerical audit (verified §13)
σ_W	sign of electron–tau vs electron–muon attachment strength	projected-Hamiltonian octant test (owed)

The discipline is the same as in the previous paper. Three levels stay distinct:

1. **Exact algebra.** If $T\_\nu = D_0 I + \varepsilon M\_\nu$, the eigenframe is independent of ε. If the Hermitian PMNS kernel is specified, its mixing outputs are fixed. This level is verified arithmetic.
2. **Support-trace theorem under stated premises.** $\beta = \sqrt{3}/20$ follows *if and only if* WL-1…WL-5 hold.
3. **Substrate computation still owed.** The projected H_cl block must verify that the residual leakage space realises WL-2…WL-4 under the WL-5 normalisation, and must compute σ_W.

The purpose of this paper is to move β from level 3 toward level 2: not yet a microscopic matrix-element computation, but no longer an unexplained numerical choice — and to make the remaining gap a short, named, falsifiable list rather than a single opaque benchmark.

---

1. The Exposed Target After the Weak-Doublet Projection Theorem

The weak-doublet projection theorem established the correct object and the correct block logic:

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

It reduced the PMNS sector to a weak-commitment residual kernel. The charged lepton is closure-norm anchored; the neutrino is weakly committed. Therefore

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

and the mixing frame is the eigenframe of M_ν — a leading-order structure, not a small effect suppressed by ε.

That theorem left several finite targets. The least secure was β, for a clear reason. The solar ratio $\rho_\odot = \sqrt{3/2}$ has a clean root-support reading. The pair ratio $B / A = 6/5$ at least inherits its denominator from the five-dimensional residual space $R = C \oplus G$. The phase $\psi = 3\pi/4$ ties to the same complement-half-transport grammar as the CKM curvature phase. By contrast, β required a denominator 20 that had never been structurally counted, and a numerator whose root-normalisation had never been stated as a rule.

So β is the right next target: small, exposed, and physically decisive. In the weak-commitment kernel it appears as the leading attachment asymmetry between the electron channel and one non-electron channel. In the electron–tau branch,

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

with the electron–muon branch obtained by $\mu \leftrightarrow \tau$ reflection. The sign convention on ψ fixes the convention for the quoted CP phase; the invariant outputs are θ_{12}, θ_{13}, $|\theta_{23} - 45°|$, $|J_{lep}|$, and the exchange of CP orientation between reflected branches.

The target is $\boldsymbol{\beta = \sqrt{3} / 20}$, and this paper derives it as

β = (root-normalised threefold leakage) / (twentyfold leakage support trace).

---

2. The Weak-Commitment Attachment Block

The relevant part of the projected Hamiltonian is not the whole weak-doublet block. It is the attachment block connecting the electron-labelled channel to the residual non-electron pair, after the charged-lepton frame has been fixed by closure-norm anchoring:

$$H_{att} = P_e H_W (P_\mu + P_\tau).$$

This block splits into a symmetric part and a leakage part,

$$H_{att} = H_{sym} + H_{leak}.$$

The symmetric part gives the common electron-to-pair coupling A:

$$P_e H_{sym} P_\mu \simeq P_e H_{sym} P_\tau \simeq A.$$

The leakage part is the first admissible μ–τ-breaking attachment asymmetry,

$$P_e H_{leak} P_\tau - P_e H_{leak} P_\mu \neq 0$$

on one branch, with the reflected branch exchanging μ and τ.

The weak-neutrality rule suppresses first-order *diagonal* μ–τ breaking,

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

so the leading atmospheric and reactor perturbation must live in this off-diagonal attachment block, not in $r_\mu - r_\tau$. This is the structural reason the two observables are tied together. A diagonal split could move θ_{23}, but it would not explain why the reactor angle and the atmospheric deviation are linked. The leakage block links them — opening θ_{13} and shifting θ_{23} off 45° with a single amplitude.

2.1 Branch basis

After charged-lepton anchoring, P_e, P_μ, P_τ are not arbitrary labels: they are the three charged-lepton-labelled attachment directions inherited by the neutrino weak-commitment block. Two leading branches are admissible:

- **electron–tau branch:** H_{leak} appears first in $P_e H_W P_\tau$;
- **electron–muon branch:** H_{leak} appears first in $P_e H_W P_\mu$.

The two are related by $\mu \leftrightarrow \tau$ reflection. So β — with the inherited phase ψ — fixes the *magnitude* of the deviation from μ–τ symmetry, while the branch sign fixes the atmospheric octant.

3. Normalisation of Leakage Amplitudes and the Support Calculus

The PMNS kernel uses A as the unit electron-to-pair attachment strength, so β is dimensionless:

$P_e H_W P_{branch} = A(1 + \beta e^{(\pm i\psi)})$.

β is therefore not the raw norm of H_{leak}. It is the residual leakage norm divided by the symmetric attachment scale *and* by the number of support slots over which the leakage is distributed:

$\beta = \|L_{coh}\|_K / \mathrm{Tr}_{support}(\Pi_{leak})$,

where L_{coh} is the coherent leakage residue surviving weak-neutral projection, Π_{leak} is the support projector for all admissible leakage placements, and $\|\cdot\|_K$ is the Killing norm inherited from the projected H_{cl} block.

3.1 The commensurability premise (WL-5)

The ratio in §3 deserves a sharp eye: the numerator is a *norm* and the denominator is a *count* (a trace of a projector — an integer dimension). These are not automatically commensurable, and dividing one by the other is only legitimate under a stated convention. That convention is premise **WL-5**:

The Killing norm $\|\cdot\|_K$ is normalised so that each admissible support slot carries unit weight.

Under WL-5 the denominator is literally "how many unit-weight slots the leakage is spread across," and the numerator is "the Killing norm of the coherent committed residue measured in those same units." A coherent residue of norm $\sqrt{3}$ distributed over 20 unit slots then *is* an amplitude $\sqrt{3}/20$, with the units cancelling cleanly. WL-5 is the keystone of the whole construction; without it, "norm over count" is a category error rather than a derivation. It is stated here as an explicit premise so the substrate calculation knows precisely what normalisation it must reproduce.

One caution must be stated plainly, because WL-5 is also where a fit could hide. "Unit weight per slot" is itself a choice of slot-weighting, and a different substrate weighting would reweight the denominator to almost any value. So WL-5 is not merely the premise that makes $\sqrt{3}/20$ well-typed; it is the joint at which, if the projected Killing norm came back unfriendly, a discrepancy could be *absorbed* rather than registered as a falsification — and absorbed by a free function, since an arbitrary reweighting has as many parameters as there are slots. The construction therefore requires more than that WL-5 be assumed: it requires that equal slot-weighting **emerge as a derived property** of the projected Killing norm on the leakage support. If the natural norm is unequal-weighted, the construction does not simply need rescaling; it loses falsifiability, because the rescaling is unconstrained. Deriving WL-5 — not stipulating it — is thus on the same owed-step footing as the orientation count, and arguably ahead of it.

3.2 The two-rule support calculus

The construction uses one calculus, applied consistently, with two distinct rules for two distinct kinds of object:

- **Coherent committed residue → norm (quadrature addition).** When several branches form a single coherent object, their contribution is the Killing *norm* of the sum. Orthogonal unit branches therefore add in quadrature: n branches give $\sqrt{n}$, not n.
- **Admissible support manifold → count (trace).** When the leakage is *distributed* over a manifold of admissible placements, the normalising weight is the *dimension* of that manifold: a plain slot count.

This is the rule that fixes the numerator as $\sqrt{3}$ rather than 3 and the denominator as 20 rather than $\sqrt{20}$. It is not a per-quantity choice tuned to the answer; it is the same distinction — coherent object vs. distribution manifold — used throughout the flavour programme (see §9). The paper therefore has two concrete jobs:

1. derive $\|L_{coh}\|_K = \sqrt{3}$ under WL-1 and the norm rule;
2. derive $\mathrm{Tr}_{support}(\Pi_{leak}) = 20$ under WL-2…WL-4 and the count rule.

Then β follows immediately.

---

4. The Root-Normalised Threefold Numerator

The numerator uses the root-count grammar already established twice in the programme.

- In the CKM curvature paper, a 2↔3 curvature norm b distributed democratically over three Killing-orthonormal C_3 branches gives a component amplitude $\mathbf{b/\sqrt{3}}$.
- In the PMNS solar-ratio argument, threefold completion support against twofold coherent pair support gives $\boldsymbol{\rho_\odot = \sqrt{3}/\sqrt{2} = \sqrt{(3/2)}}$.

The leakage numerator follows the same principle (premise **WL-1**). The leading leakage residue is carried by a threefold completion orbit whose branches are independent norm components, not one linearly summed amplitude. Let the three leakage branches be ℓ_1, ℓ_2, ℓ_3 with

$$\boldsymbol{\langle \ell_i, \ell_j \rangle_K = 0 \text{ for } i \neq j,\ \|\ell_i\|_K = 1.}$$

The coherent leakage residue is $\mathbf{L_{coh} = \ell_1 + \ell_2 + \ell_3}$, so

$$\boldsymbol{\|L_{coh}\|_K^2 = 3,\ \|L_{coh}\|_K = \sqrt{3}.}$$

Lemma 1 — Threefold leakage numerator

If the weak-commitment leakage residue is the coherent sum of three Killing-orthonormal completion branches, each of unit leakage norm (WL-1), then the leakage numerator is √3.

Proof. By orthonormality, $\|\ell_1 + \ell_2 + \ell_3\|_K^2 = \|\ell_1\|_K^2 + \|\ell_2\|_K^2 + \|\ell_3\|_K^2 = 3$. Taking the positive root gives $\sqrt{3}$. ■

Status. The lemma is elementary linear algebra. The substantive content is premise WL-1 itself: that the projected weak-commitment leakage channel really is a three-branch Killing-orthonormal completion residue. That is owed to the substrate projection (see §10.2).

Relation to the residual support (one closed point, one open one). The three numerator branches are the completion directions: they span C, the dim-3 summand of $R = C \oplus G$, so $\{\ell_1, \ell_2, \ell_3\}$ lie *within* R rather than orthogonal to it — as they must, since the leakage residue has to live in the support it is diluted across. This much is settled and does **not** collapse WL-4: WL-4 concerns R vs Q (support vs orientation, which remain disjoint), and is untouched by the numerator branches lying inside R; the branches, in $C \subset R$, are automatically disjoint from the orientation fibre Q whenever WL-4 holds.

But placing the coherent carrier inside R opens a *third* choice the count rule does not by itself settle, and it must be named rather than waved past. The same three C-directions now do double duty: they build the numerator (norm √3, the forward Pythagorean reading) and they are also among the five counted in the denominator's R-factor (the backward dilution reading). Counting all of R — the carrier included among the slots it is diluted across — gives dim R = 5 and the benchmark 5 × 4 = 20. Excluding the carrier and diluting only over its complement gives dim $R_{dilution} = 2$, denominator 2 × 4 = 8, $\beta = \sqrt{3}/8 \approx 0.217$, and $\theta_{13} \approx 16.3°$ (verified) — far too large. No principle stated here decides whether a direction the residue coherently *occupies* is also a slot it is diluted *across*; "count all of R" is, once again, what the reactor angle needs, not what the calculus dictates. So this carrier-inclusion convention joins the subtraction (−2) and the orientation multiplicity (×4) as a third data-selected joint (see §10.3a, §11 item 6a), and §4 should not be read as closing the matter.

---

5. The Twentyfold Leakage Support Trace

The denominator 20 is the new work of this paper, and its weakest link is the orientation multiplicity, which this section states with deliberate care.

The residual weak-commitment support space descends from the K = 7 closure carrier (premise **WL-2**). 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

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

and these five survivors organise as the residual module already used elsewhere in the flavour programme,

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

This gives the denominator's first factor a closure pedigree rather than an inherited stipulation: the 5 is the residual support of the same $K = 7$ carrier that fixes $\alpha^{-1} \approx 137$, that the no-go theorem derives, and that the closure-symmetry and Wilson-limit papers turn on. Two riders remain owed — that exactly two directions anchor (§11, item 4), and that the five survivors really are $C \oplus G$ (§10.1a) — but the number 5 is no longer free.

This fivefold residual space already appears in the pair ratio $B / A = 6/5$. But leakage is not merely a residual support count: it is an *off-diagonal* weak attachment, and so it also carries weak-role orientation information. The claim (premise **WL-4**) is that the leakage support factors,

$\Pi_{leak} \simeq \Pi_R \otimes \Pi_Q$,

with Π_R the five-dimensional residual support and Π_Q a four-dimensional weak-attachment orientation fibre (premise **WL-3**). Hence

$\mathrm{Tr}_{support}(\Pi_{leak}) = \mathrm{Tr}(\Pi_R) \cdot \mathrm{Tr}(\Pi_Q) = 5 \times 4 = 20$.

The central question is therefore: **why $\dim Q = 4$?**

5.1 The fourfold weak-attachment orientation fibre — stated carefully

A warning is in order, because the obvious justification is wrong. The leakage enters M_ν as a single complex off-diagonal entry, $M_{e\tau} = A\cdot\beta\cdot e^{-i\psi}$, with $M_{\tau e} = M_{e\tau}^*$ fixed by Hermiticity. Counted as independent *real degrees of freedom of the matrix*, that is **2**, not 4: the conjugate return entry is determined, not free. So one cannot reach 4 by saying "two quadratures × two Hermitian orientations," because the Hermitian return is not an independent direction. Any derivation that double-counts the conjugate is invalid, and a referee will say so.

The factor 4 must therefore be a count of **support placements**, not of independent matrix DOF. Under WL-5, Π_Q is the orientation fibre over which the single complex attachment is *distributed*, and its multiplicity is the product of two genuinely independent two-element structures:

1. **Weak-neutral quadrature (×2).** The *support manifold* of admissible leakage placements spans two real directions in the weak-neutral plane (magnitude and weak-neutral phase). This is a dimension of the placement manifold, not a splitting of the committed residue: the actual residue is a single coherent complex object sitting at the sharp transported phase $\psi = 3\pi/4$ (carrying $O_\nu = e^{i\pi/2}$) at one point of that manifold.

2. **Weak-role complement pair (×2).** The attachment lives *across* a weak-role complement, and weak-doublet closure requires both members of that complement pair to carry the attachment. This factor comes from the doublet structure, independently of Hermitian conjugation.

So

dim Q = 2 (quadrature) × 2 (role complement) = 4.

This is the important structural point, now stated without the conjugate double-count: the denominator is not 5 alone, because leakage is not a scalar residual weight. It is a complex weak attachment distributed across a weak-role complement, so the residual denominator 5 is lifted to **5 × 4 = 20**. The two factors are sourced from two different and independent structures — the weak-neutral plane and the weak-doublet complement — which is why they multiply rather than collapse.

Consistency note (coherent phase vs. counted quadrature). The two-rule calculus of §3.2 treats the committed residue as a coherent object, which is exactly what licenses extracting ψ as a single *sharp* complement-half-transport phase. The quadrature factor above then counts the same weak-neutral plane as a *placement dimension* in the denominator. These two roles are compatible only under one reading: the denominator counts the dimension of the manifold of admissible placements, while the numerator is the norm of the one coherent residue that occupies a sharp phase *within* it — the residue's phase is its coordinate in the quadrature plane, not evidence that the residue is incoherently smeared across it. The tension this reading must avoid is real and load-bearing. If the two quadratures are instead internal coordinates of the single coherent residue — used up by the numerator object rather than available as denominator slots — then the quadrature factor collapses, dim Q = 2, the denominator is 10, and $\theta_{13} \approx 14.2°$ (excluded). The construction therefore *needs* the quadrature to be a manifold dimension, and at present that requirement is enforced by the reactor angle, not by the support calculus itself. Settling it — manifold dimension vs. residue-internal coordinate — is part of the WL-3 substrate check; §3.2 alone does not decide it.

Lemma 2 — Fourfold orientation fibre

If the leading weak-commitment leakage is a complex weak-neutral attachment distributed across a weak-role complement pair, then its weak-attachment orientation fibre Q has support multiplicity four: two weak-neutral quadrature placements times two weak-role complement members.

Proof sketch. Weak neutrality suppresses first-order diagonal μ–τ stiffness splitting (§2), forcing the leading perturbation into an off-diagonal attachment. As a complex object the attachment occupies two real quadrature placements in the weak-neutral plane. Weak-doublet closure requires the attachment to be carried across both members of the weak-role complement, supplying a second, independent factor of two. The two factors arise from distinct structures, so the orientation fibre has 2 × 2 = 4 support slots. ■

Status. This is the load-bearing premise of the paper (WL-3), and *both* of its factors of two are open in the same way. The weak-neutral plane certainly exists, but whether its two quadratures count as denominator dimensions or are absorbed as internal coordinates of the coherent numerator residue is unsettled (see the consistency note above); and the role-complement multiplicity must be confirmed to be exactly two. If the quadrature factor collapses, dim Q = 2 and the denominator is 10 ($\theta_{13} \approx 14.2°$); if the complement multiplicity is 1 or 3, the denominator is 10 or 30. The substrate projection must fix both. At present the reactor angle (§13) is what selects 4 — which is to say the orientation count is, for now, data-selected rather than rule-selected.

Lemma 3 — Twentyfold leakage trace

If the leakage support factors as $\Pi\_leak \simeq \Pi\_R \otimes \Pi\_Q$ (WL-4) with Tr $\Pi\_R = 5$ (WL-2) and Tr $\Pi\_Q = 4$ (WL-3), then Tr_support(Π_leak) = 20.

Proof. By factorisation, Tr_support(Π_leak) = Tr(Π_R) · Tr(Π_Q) = dim R · dim Q = 5 · 4 = 20. ■

Status. The residual factor dim R = 5 is K = 7-descended (WL-2). The orientation factor dim Q = 4 follows from Lemma 2 (WL-3). The factorisation Π_leak ≃ Π_R ⊗ Π_Q (WL-4) is itself owed: the substrate calculation must verify that the residual support and the orientation fibre enter as an honest tensor product inside P_W H_cl P_W, rather than sharing directions (which would lower the trace below 20).

5.2 What K = 7 does and does not buy

The descent sharpens the denominator asymmetrically, and the paper is explicit about it. The factor **5 is K = 7-descended** (WL-2): it is the residual support of the same closure carrier that runs through the rest of the programme, so it is no longer a stipulated number. The factor **4 is not** (WL-3): it comes from the weak-attachment orientation fibre of §5.1, which is independent of the closure-carrier count and remains the genuine open premise.

So the K = 7 descent moves the denominator from *two stipulated factors* to a closure-anchored carrier with two operations still hanging off it. The fair description is **anchored carrier, two free joints**: K = 7 is pedigreed (the same K that fixes $\alpha^{-1} \approx 137$), but neither the subtraction that turns 7 into 5 (the "exactly two anchor" claim, §10.1a) nor the multiplication by the orientation fibre (dim Q = 4, WL-3) is yet rule-fixed. That is real progress — the carrier is no longer free, and the reactor angle is wired back to the closure programme — but it would overstate it to say "half the denominator now has a spine," because both operations on the carrier remain open and are currently selected by θ_{13} rather than by the support calculus. The priority substrate check (§10.3) is still dim Q = 4.

---

6. The Leakage Theorem: β = √3 / 20

Theorem 1 — Weak-commitment leakage trace

Assume the leading weak-commitment leakage channel satisfies:

1. *(WL-1) threefold completion residue — the coherent leakage numerator is carried by three Killing-orthonormal completion branches;*
2. *(WL-2) $K = 7$-descended residual support — the leakage denominator includes the five-dimensional residual support $\dim R = K - 2 = 5$, organised as $R = C \oplus G$;*
3. *(WL-3) weak-attachment orientation — the leakage is a complex weak-neutral attachment across a weak-role complement, with four orientation placements;*
4. *(WL-4) support factorisation — $\Pi\_leak \simeq \Pi\_R \otimes \Pi\_Q$;*
5. *(WL-5) commensurable unit-slot normalisation of $\|\cdot\|\_K$.*

Then the dimensionless leakage amplitude is

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

Proof. By Lemma 1 and WL-1, the coherent leakage numerator is $\|L\_coh\|\_K = \sqrt{3}$. By Lemmas 2–3 and WL-2…WL-4, the leakage support trace is Tr_support(Π_leak) = 20. By WL-5 the two are commensurable, so

$\beta = \|L\_coh\|\_K / \mathrm{Tr\_support}(\Pi\_leak) = \sqrt{3} / 20$. ■

This is the proposed derivation of the benchmark leakage amplitude — not a phenomenological fit, but a support-trace statement about the projected weak-commitment attachment block, explicitly conditional on five named premises.

6.1 Why β is small but not negligible

The size of β is now explained structurally. It is small because the leakage is diluted over twenty support slots. It is not tiny because the numerator is a coherent root-three residue rather than a single suppressed branch:

$\beta = \sqrt{3} / 20 \approx 0.0866$.

That is exactly the scale needed for a reactor angle of order 8.6° in the weak-commitment kernel.

6.2 Why the numerator is √3 — and the limit of "the rule decides"

A *linear* threefold numerator would give $\beta = 3/20 = 0.15$, raising the reactor angle to ≈ 13° (verified, §13) and spoiling the benchmark. On the *numerator axis* the rule genuinely decides: √3 is forced by the support calculus of §3.2, not chosen to hit 8.6°. A coherent residue built from orthogonal branches contributes its *norm*, and the norm of three orthonormal branches is √3 — the same rule that gives $\rho\_\odot = \sqrt{3/2}$ and the CKM amplitude $b/\sqrt{3}$. So √3 vs 3 is settled by the rule first and merely corroborated by the data.

The denominator axis must be stated more cautiously, because here the situation is reversed and the numerator's clean win must not lend it borrowed authority. Among the family √3/{10, 16, 20, 24, 30, …} there is at present *no rule* that excludes the alternatives — WL-2 and WL-3 are precisely the premises that would do the excluding, and they are owed, not derived. What the rule (plus the K = 7 descent) currently fixes on the denominator is the *carrier*: that the first factor is the closure-carrier residual support. The two operations applied to that carrier — the subtraction (K − **2**) and the orientation multiplicity (× **4**) — are not yet rule-fixed; they are selected by the reactor angle. So on the denominator axis it is the data doing the selecting, and the paper claims only that √3/20 is the value the owed premises must reproduce, not that a rule has already excluded √3/16 or √3/30.

---

7. PMNS Kernel Audit

Use the scale convention

$A = 1, r\_e = 0, B = 6/5, r\_s = \sqrt{3/2} - 6/5, \delta = 0, \psi = 3\pi/4, \beta = \sqrt{3}/20.$

For the electron–tau branch,

$$M\_\nu^{+} = \begin{pmatrix} 0 & 1 & 1 + \beta\, e^{(-i\psi)} \\ 1 & r\_s & 6/5 \\ 1 + \beta\, e^{(+i\psi)} & 6/5 & r\_s \end{pmatrix}$$

The electron–muon branch is the μ ↔ τ reflection of the same leakage magnitude. Exact Hermitian diagonalisation gives the branch-independent outputs

$\theta_{12} \approx 33.40°, \theta_{13} \approx 8.59°, |J\_lep| \approx 2.45 \times 10^{-2},$

and the reflected atmospheric angle

electron–tau branch: $\theta_{23} \approx 48.37°$; electron–muon branch: $\theta_{23} \approx 41.63°$,

so the channel-independent atmospheric prediction is

$|\theta_{23} - 45°| \approx 3.37°.$

This is the quantitative value controlled by β.

7.1 What the audit proves

The audit proves that the support-trace value β = √3/20 lands on the intended PMNS leakage scale. It does not, by itself, prove which branch nature selects; that is σ_W.

7.2 What the audit does not prove

The audit does not replace the substrate calculation. It shows only that *if* the projected weak-commitment attachment block supplies the support trace of §5, *then* the PMNS kernel has the desired reactor and atmospheric scale. The support trace still has to be found inside H_W, and σ_W still has to be computed.

---

8. The Octant Sign σ_W

The leakage trace fixes the *magnitude* of μ–τ breaking, but not its orientation. The orientation is the sign of the projected attachment imbalance,

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

equivalently

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

The reading is direct:

- **σ_W > 0:** electron–tau attachment dominates; the upper-octant branch is selected.
- **σ_W < 0:** electron–muon attachment dominates; the lower-octant branch is selected.
- **σ_W = 0:** the leading leakage trace is μ–τ symmetric; the benchmark predicts the magnitude of possible splitting but not an octant at first order.

Theorem 2 — Leakage magnitude and octant sign separate

If weak neutrality suppresses first-order diagonal μ–τ splitting and the leading atmospheric breaking is electron-attachment leakage, then β and the inherited phase ψ jointly fix $|\theta_{23} - 45°|$, while σ_W fixes the sign of $\theta_{23} - 45°$.

Proof sketch. The two leakage branches are related by μ ↔ τ reflection. Reflection leaves θ_{12}, θ_{13}, and |J_lep| invariant but sends $\theta_{23} \to 90° - \theta_{23}$ and reverses the CP orientation. The leakage amplitude β and the transported phase ψ together control the distance from the symmetric point — at fixed β the deviation runs from ≈ 0.25° at ψ = π/2 to ≈ 4.6° at ψ = π (§13) — while the projected attachment imbalance σ_W selects which side of the symmetric point is realised. ■

The octant is therefore not guessed from experiment. It is reduced to one sign in the projected H_cl block.

8.1 What would derive σ_W?

A full derivation compares the two projected attachment norms,

$\|P\_e\ P\_W\ H\_cl\ P\_W\ P\_\tau\|^2$ vs. $\|P\_e\ P\_W\ H\_cl\ P\_W\ P\_\mu\|^2$.

The sign could arise from one of three closure-geometric sources:

1. **electron-channel continuation bias** — the electron attachment may have one extra continuation toward τ rather than μ, or vice versa;
2. **orientation branch bias** — complement-half-transport may select a preferred weak-neutral quadrature orientation;
3. **residual support asymmetry** — the $C \oplus G$ residual decomposition may couple unevenly to μ and τ after charged-lepton anchoring.

The paper does not choose among these by hand. It states the calculation: compute σ_W.

---

9. Relation to CKM Curvature and the Shared Phase Grammar

The leakage trace is not an isolated numerological move. It is the PMNS counterpart of the CKM curvature support logic, and the §3.2 calculus is precisely what the two sectors share.

The CKM side uses

$$\mathbf{z_{C3} = (b / \sqrt{3}) \cdot e^{(5\pi i/6)}.}$$

The amplitude $b/\sqrt{3}$ distributes a committed $2\leftrightarrow3$ curvature norm over a threefold C_3 orbit (the norm rule). The phase $5\pi/6$ comes from complement-half-transport against the C_3 orientation,

$$\mathbf{h_q^2 = e^{(i\pi)} \cdot e^{(2\pi i/3)}.}$$

The PMNS leakage side uses

$$\boldsymbol{\beta = \sqrt{3} / 20, \psi = 3\pi/4.}$$

The numerator $\sqrt{3}$ again reflects threefold root-normalised support (norm rule). The denominator 20 reflects the support trace appropriate to a distributed weak-commitment leakage (count rule). The phase $3\pi/4$ comes from complement-half-transport against the weak-neutral quadrature,

$$\mathbf{h_\nu^2 = e^{(i\pi)} \cdot e^{(i\pi/2)}.}$$

The shared grammar:

Sector	Support rule	Phase rule
CKM curvature	C_3 norm sharing → $b/\sqrt{3}$	complement-half-transport → $5\pi/6$
PMNS solar	root 3 : 2 support → $\sqrt{(3/2)}$	—

Sector	Support rule	Phase rule
PMNS leakage	root-three numerator over 20 support slots → $\sqrt{3}/20$	complement-half-transport → $3\pi/4$

The unification claim is *not* that the same number appears everywhere. It is that the same two rules — root-normalised norm for coherent residues, support count for distribution manifolds, plus complement-half-transport for phases — govern the allowed numbers in every sector.

---

10. Finite Substrate Tests

The paper leaves a short, concrete checklist, each item tied to a named premise.

10.1 Leakage support factorisation (WL-4, WL-2, WL-3)

Does the projected leakage support factor as $\Pi\_leak \simeq \Pi\_R \otimes \Pi\_Q$ with $\mathrm{Tr}\ \Pi\_R = 5$ and $\mathrm{Tr}\ \Pi\_Q = 4$? If the support trace is not 20, the β derivation fails or must be revised.

10.1a Residual support descent (WL-2)

Do exactly two closure directions anchor under boundary/readout projection, giving $\dim R = K - 2 = 5$ rather than $K - 1 = 6$ or $K - 3 = 4$? And do the five survivors organise as $C \oplus G$ (dim 3 + 2)? The first sub-check fixes the reactor-angle scale — $K - 3 = 4$ would give denominator 16 and $\theta_{13} \approx 10.2°$ (verified, §13) — while the second is the bridge between the $K = 7$ descent and the older residual-module account.

10.2 Threefold numerator (WL-1)

Are there exactly three Killing-orthonormal leakage branches, $\langle \ell\_i, \ell\_j \rangle\_K = \delta\_ij$? If the branches are not orthonormal, or their number is not three, the numerator is not $\sqrt{3}$.

10.3 Orientation multiplicity (WL-3) — the priority check

Two sub-questions, both currently selected by the reactor angle rather than by the support calculus. First, does the weak-neutral quadrature count as two placement dimensions of the support manifold, or is it absorbed as an internal coordinate of the coherent residue (in which case it does not enter the denominator)? Second, is the weak-role complement multiplicity exactly two? The product must be four: $\dim Q = 4$ gives $\beta = \sqrt{3}/20$ ($\theta_{13} \approx 8.6°$), whereas a collapse of either factor gives $\dim Q = 2$, $\beta = \sqrt{3}/10$ ($\theta_{13} \approx 14.2°$, excluded), and a complement multiplicity of three gives $\dim Q = 6$, $\beta = \sqrt{3}/30$ (too small). This is the least secure step in the derivation and should be computed first.

10.3a Carrier-inclusion convention (numerator/denominator overlap)

The coherent numerator carrier is the C-block, which lies inside the denominator's R-factor (§4 note). Does the dilution count include those carrier directions (count all of R → dim R = 5) or exclude them (dilute over R⊖C → dim R = 2)? The benchmark needs the former: exclusion gives denominator 8, $\beta = \sqrt{3/8}$, $\theta_{13} \approx 16.3°$ (verified, excluded). Like the −2 and the ×4, this is at present selected by the reactor angle, and the substrate projection must justify counting a coherently occupied direction as also a dilution slot.

10.4 Commensurability (WL-5)

Does the projected Killing norm admit the unit-slot normalisation, so that the coherent residue and the support count are measured in the same units? If the natural substrate norm assigns unequal weights to slots, the construction does not merely need reweighting — it *loses falsifiability*, because an arbitrary slot-reweighting is a free function able to absorb almost any discrepancy. WL-5 must therefore emerge as a *derived* equal-weighting property of the projected Killing norm on the leakage support, not be stipulated; this is arguably the first owed step, ahead of the orientation count.

10.5 Phase orientation

Is the leakage phase the weak-neutral complement-half-transport branch, $\psi = 3\pi/4$? The square root of complement reversal has two *antipodal* roots — $\psi = 3\pi/4$ and its antipode $7\pi/4 = 3\pi/4 + \pi$ — which carry opposite-sign imaginary parts and hence select opposite CP orientation. (The conjugate $5\pi/4$ is a distinct third point that does not solve the square-root equation; branch selection and CP conjugation are different operations.) The branch is fixed by a phase-branch sign $\sigma_\varphi^\nu = \mathrm{sgn}(\mathrm{Im}\, h_\nu)$, on the same footing as σ_W, not by the word "principal". The absence of a half-transport branch breaks the shared phase grammar.

10.6 Octant sign (LT-2)

Compute $\sigma_W = \mathrm{sgn}(\|P_e H_W P_\tau\|^2 - \|P_e H_W P_\mu\|^2)$. This is the next projected-Hamiltonian number after β.

---

11. Falsification Conditions

The leakage-trace proposal is exposed to clean failure.

1. If the weak-commitment neutrino operator is not of the form $T_\nu = D_0 I + \varepsilon M_\nu$, the whole large-PMNS explanation fails.
2. If first-order diagonal μ–τ splitting appears at the same order as attachment leakage, the β-controlled reactor/atmospheric link loses predictive force.
3. If the leakage numerator is linearly threefold rather than root-normalised (WL-1 misread), β changes to 3/20 and the benchmark is spoiled.

4. If the residual support dimension is not five (WL-2 fails), the denominator 20 is not derived. In particular, if boundary/readout anchoring removes three directions rather than two — dim R = K − 3 = 4 — the denominator is 16, $\beta = \sqrt{3}/16$, and θ_{13} rises to ≈ 10.2° (verified, §13), spoiling the reactor angle. 4a. If the five surviving K = 7 directions do not organise as $C \oplus G$ (dim 3 + 2), the bridge between the closure-carrier descent and the inherited residual module fails, even though the count 5 is correct.
5. If the weak-attachment orientation fibre is not four-dimensional (WL-3 fails) — in particular if the weak-role complement multiplicity is 1 or 3 rather than 2, *or* if the weak-neutral quadrature is absorbed as an internal coordinate of the coherent residue rather than counting as two placement dimensions — the denominator 20 is not derived (it becomes 10 or 30).
6. If Π_leak does not factor as Π_R ⊗ Π_Q (WL-4 fails), the trace calculation must be replaced. 6a. If the dilution count must exclude the coherent carrier (dilute over $R \ominus C$ rather than all of R), dim R_dilution = 2, the denominator is 8, and $\theta_{13} \approx 16.3°$ — so the carrier-inclusion convention, like the −2 and the ×4, must be justified rather than assumed.
7. If the Killing norm admits no unit-slot normalisation (WL-5 fails), the ratio $\sqrt{3}/20$ is not the correct dimensionless amplitude even when the counts are right — and worse, if the natural norm is unequal-weighted, the denominator becomes adjustable by a free reweighting function, so the construction loses falsifiability rather than failing cleanly. WL-5 must be shown to *emerge* as equal-weighting, not be assumed.
8. If β is not close to $\sqrt{3}/20$, the reactor-angle magnitude remains fitted rather than derived.
9. If σ_W = 0 at leading order, the framework does not predict an octant at first order.
10. If σ_W selects the experimentally disfavoured branch after the atmospheric octant stabilises, the weak-attachment orientation rule is falsified.
11. If the PMNS leakage requires rules unrelated to the CKM curvature support rules, the unified flavour grammar is weakened even if the numerical benchmark survives.

---

12. Conclusion

The weak-doublet projection theorem identified the remaining flavour frontier: compute the projected blocks and support traces of

H_W = P_W H_cl P_W.

This paper attacks the weakest exposed PMNS target, the leakage amplitude β, and gives a support-trace derivation:

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

The numerator $\sqrt{3}$ is a root-normalised threefold leakage residue (WL-1, the norm rule). The denominator 20 factors as (K − 2) × 4: a five-dimensional residual support that descends from the K = 7 closure carrier minus two boundary/readout-anchored directions (WL-2), multiplied by a fourfold weak-attachment orientation fibre (WL-3) via an honest tensor factorisation (WL-4), with numerator and denominator made commensurable by a unit-slot Killing norm (WL-5). The descent is asymmetric and the paper says so: the 5 now has a closure pedigree, the 4 does not. The same value produces the intended PMNS scale — $\theta_{13} \approx 8.6°$ and $|\theta_{23} - 45°| \approx 3.4°$ — reproduced here by independent diagonalisation.

The atmospheric octant remains properly separated: β and the inherited phase ψ fix the size of the splitting, σ_W fixes its sign,

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

The advance is that β moves from a named benchmark toward a derived support invariant, while the remaining gap is now a short, named, falsifiable list rather than one opaque number. The honest weak point is explicit. The K = 7 carrier is anchored, but *three* operations on it remain open and are at present selected by the reactor angle, not by the support calculus: the subtraction that turns 7 into 5 (the "exactly two anchor" claim), the orientation multiplicity dim Q = 4 (resting on both the weak-role complement multiplicity and the quadrature counting as placement dimensions), and the carrier-inclusion convention that counts the coherent numerator's own C-directions among the dilution slots (count all of R rather than R⊖C). Each is individually defensible and each lands near $\theta_{13} \approx 8.6°$ only because the data is there to pick it. And the construction stands or falls on WL-5 above all: if equal slot-weighting does not *emerge* from the projected Killing norm, the denominator becomes adjustable by a free reweighting and the result loses falsifiability rather than failing cleanly. If the projected H_cl leakage block carries a 5 × 4 support trace under a genuinely unit-slot norm, with the carrier counted, the weak-commitment PMNS kernel gains its missing reactor-angle derivation; if it does not, each named premise says exactly which joint gave way.

The next technical calculation is therefore completely explicit: compute Π_leak; confirm that the projected Killing norm is equal-weighted (WL-5) before anything else, since an unequal norm would make the rest unfalsifiable; verify or reject Tr_support(Π_leak) = 20 — fixing the subtraction, the orientation count, and the carrier-inclusion convention — and compute σ_W.

---

13. Numerical Note

Using

$$\mathbf{A = 1,\ r\_e = 0,\ B = 6/5,\ r\_s = \sqrt{(3/2)} - 6/5,\ \delta = 0,\ \psi = 3\pi/4,\ \beta = \sqrt{3}/20,}$$

and the electron–tau branch

$$M\_\nu^+ = \begin{pmatrix} 0 & 1 & 1 + \beta\ e\hat{}(-i\psi) \\ 1 & r\_s & 6/5 \\ 1 + \beta\ e\hat{}(+i\psi) & 6/5 & r\_s \end{pmatrix}$$

exact Hermitian diagonalisation, with ν_3 chosen as the column of smallest electron content, gives:

Branch	θ_{12}	θ_{13}	θ_{23}	$|J\_lep|$
electron–tau	33.40°	8.59°	48.37°	2.45×10^{-2}
electron–muon reflection	33.40°	8.59°	41.63°	2.45×10^{-2}

The reflected branches exchange the atmospheric octant and the CP orientation. Exact quoted δ_CP values depend on the branch phase convention; the invariant statements are the octant reflection, the sign flip of the CP orientation, and the common magnitude |J_lep|.

These values were reproduced independently by direct numerical diagonalisation of the benchmark kernel (NumPy `eigh`, standard PDG angle extraction), confirming the §7 audit to the quoted precision.

The reactor angle is a probe of the denominator, but only a *conditional* one, and the condition must be stated. θ_{13} is a function of (β, ψ), not of β alone. Holding $\beta = \sqrt{3}/20$ fixed and sweeping the phase gives $\theta_{13} = 9.75°$ at $\psi = \pi/2$, 8.59° at $\psi = 3\pi/4$, and 8.19° at $\psi = \pi$ (with $|\theta_{23} - 45°| =$ 0.25°, 3.37°, 4.64° respectively) — a range of order 8–10° from the phase alone. So the denominator scan below is read *at the inherited phase $\psi = 3\pi/4$*; the lever is not free-standing, and a (denominator, phase) pair could in principle be traded along. It is usable here only because ψ is pinned independently by the closure-kernel paper, not by θ_{13}.

At fixed $\psi = 3\pi/4$, the denominator does move the angle: $\beta = \sqrt{3}/20$ gives $\theta_{13} = 8.59°$, whereas $\beta = 3/20$ (linear numerator) gives ≈ 12.9°, $\beta = \sqrt{3}/16$ (K − 3 = 4) gives ≈ 10.2°, and $\beta = \sqrt{3}/40$ (dim Q = 8) gives ≈ 4.7° — all verified, all excluded by the observed $\theta_{13} \approx 8.5°$. With ψ fixed, only the (K − 2) × 4 = 20 count lands the angle.

Two cautions on reading the angle table. First, θ_{12} is **inherited, not a β-driven output**: at β = 0 the kernel already gives $\theta_{12} = 33.29°$ (set by ρ_⨀ and B/A), and the benchmark β shifts it by only ≈ 0.1°. Only θ_{13} and $|\theta_{23} - 45°|$ are genuinely β-driven, so the agreement of θ_{12} with global fits should be credited to the solar and pair ratios, not counted again here. Second, the magnitude of the atmospheric deviation is fixed jointly by β *and* ψ, not by β alone (the sweep above); σ_W fixes only its sign. With those caveats, $\theta_{13} \approx 8.6°$ sits close to current global-fit central values and θ_{23} falls within the measured atmospheric range on either branch; the paper makes no claim to a precision fit.

---

14. Source Map

- *The Weak-Doublet Projection Theorem in VERSF* — identifies β as the weakest PMNS target, defines σ_W as the projected octant sign, and reduces the remaining problem to Ω_0^(23), z_C3, M_ν, σ_W, and β.

- *The Projected Weak-Doublet Closure Hamiltonian in VERSF* — establishes H_W = P_W H_cl P_W and the committed / weak-commitment unification of CKM and PMNS.
- *The Weak-Commitment Closure Kernel in VERSF* — supplies the PMNS kernel, ρ_⊙, B/A, ψ, the β benchmark, and the atmospheric branch structure.
- *The Weak-Commitment Neutrino Operator in VERSF* — develops stiffness-gap collapse and the weak-commitment frame theorem.
- *The Electroweak Flavour-Frame Operator in VERSF* — supplies the weak-role decomposition and the role-frame architecture.
- *Deriving Flavour Mixing from Closure Geometry* — supplies the residual support space R = C ⊕ G with dim C = 3, dim G = 2, dim R = 5.
- *A No-Go Theorem for Non-Simplicial Relational Substrates* — derives K = 7 as the minimal admissible encoding via the Hamming-bound / Fano-plane route; the closure-carrier count behind dim R = K − 2.
- *The K = 7 Closure Symmetry* — supplies the closure-counting framework (N_loop = 14 = 2K, the 11 + 2 + 1 channel decomposition) from which the residual support descends.
- *The K = 7 Wilson Limit* — uses N_loop = 14 as the closure-counting scale, fixing the same K used here.
- *Half-Integer Spin as Closure Two-Cycle Structure in VERSF* — supplies the binary closure-phase grammar supporting weak-neutral quadrature as genuine closure-attachment structure.
- *The CKM Curvature Residue in VERSF* — supplies the C_3 curvature target and the root-normalised support grammar used as the quark-side analogue.
- *Substrate Dynamics and the Higgs Ratio* — identifies H_cl as the su(8)-restricted admissibility Hessian.

Appendix A — Minimal Algebra of the Leakage Perturbation

In the μ–τ-symmetric limit, the kernel

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

is diagonalised by the symmetric / antisymmetric basis

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

The antisymmetric state ν_- has no electron component in the symmetric limit. The solar block is

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

with

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

The leakage perturbation on one attachment branch is

$\Delta M = \beta A e^{(-i\psi)} |e\rangle\langle\tau| + \beta A e^{(+i\psi)} |\tau\rangle\langle e|$

(or its μ-reflected counterpart). In the $\nu\pm$ basis, this perturbation contains an electron-to-antisymmetric component of order

$\beta A / \sqrt{2}$,

which opens the reactor angle and simultaneously shifts the atmospheric angle off maximality. Thus the same β controls θ_{13} and $|\theta_{23} - 45°|$. This is why deriving β matters more than fitting a small asymmetry: it locks together two PMNS observables that would otherwise be independently adjustable.

The full numerical values quoted in §13 come from exact Hermitian diagonalisation of the benchmark kernel, not from this first-order sketch.