The Weak-Commitment Neutrino Operator in VERSF

PMNS Mixing from Stiffness-Gap Collapse, μ–τ Breaking, and Completion-Transport Phases

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

Successor to *The Electroweak Flavour-Frame Operator in VERSF*, *The CKM Curvature Residue in VERSF*, and *The Yukawa Operator from Completion-Channel Misalignment in VERSF*.

Summary for the General Reader

The earlier VERSF flavour papers changed the shape of the Standard Model flavour problem. Instead of treating masses, CKM mixing, PMNS mixing, and CP violation as separate numerical mysteries, they proposed that flavour is a *frame* problem. Each particle sector reads the same underlying completion register in its own basis. Masses are the readout values of that sectoral operator; mixing is the mismatch between sectoral frames.

For quarks this is now fairly sharp. The up and down quark frames sit close together because both members of the weak doublet are closure-committed, and their small mismatch gives CKM mixing. The CKM curvature paper reduced the remaining quark problem to a single commutator correction that repairs the long unitarity-triangle side and the Jarlskog invariant — conditional on deriving that common-mode curvature from the weak-doublet closure Hamiltonian.

The lepton side is more open, and for a structural reason. The charged lepton is anchored and mass-resolved, while the neutrino is neutral, weakly committed, and nearly degenerate. If the neutrino stiffness hierarchy collapses, the neutrino frame can rotate by large angles *even when the absolute neutrino masses are tiny*. This paper turns that statement into a concrete operator target.

The central claim:

PMNS mixing is large because the neutrino transport operator loses the stiffness hierarchy that suppresses quark mixing.

In the quark sector, off-diagonal transport is small against large generation gaps, so mixing angles are small. In the neutrino sector, both the residual gaps **and** the off-diagonal transports scale with the same weak-commitment parameter ε, and ε cancels out of the eigenvectors. Small neutrino masses therefore do not force small neutrino mixing.

The paper proposes a minimal weak-commitment operator M_ν. Its symmetric core has approximate e–μ / e–τ equality and approximate μ–τ symmetry. In that limit the atmospheric angle is exactly maximal, the reactor angle vanishes, and the solar angle is large without large couplings — a result confirmed here to machine precision. Small symmetry-breaking terms then generate a nonzero reactor angle, an atmospheric deviation, and a physical leptonic CP phase, also confirmed numerically.

But this draft is deliberate about what it does **not** yet achieve. A parameter count (Section 7) shows that the *symmetric core* is genuinely predictive — μ–τ symmetry fixes two of the three angles by structure, not by fitting — while the *breaking sector* is not yet predictive: it carries exactly as many free parameters as the deviations it must produce. So reproducing the measured PMNS matrix is currently necessary, not impressive. The real lepton-sector task is to reduce that breaking freedom from closure geometry, exactly as the CKM paper reduced its curvature to a single complex number.

Abstract

The VERSF Yukawa-operator paper identified the left flavour operators

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

as sectoral completion operators whose eigenvalues are diagonal mass/readout values and whose eigenvectors are sectoral completion frames. CKM and PMNS mixing are relative left frames,

$$V_{CKM} = U_u^\dagger U_d \, , \quad U_{PMNS} = U_e^\dagger U_\nu \, .$$

The electroweak flavour-frame paper extended this to weak doublets: quarks and leptons are two regimes of one left-handed frame-splitting problem. In the quark regime both doublet members are closure-committed, giving a small CKM split. In the lepton regime the charged lepton is anchored but the neutrino is weakly committed and nearly degenerate, allowing large PMNS mixing.

This paper develops the lepton half. We take the neutrino transport operator in weak-commitment form

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

where $D_0 \cdot I$ is the collapsed common stiffness, ε is the weak-commitment scale, and M_ν is an order-one residual transport-shape operator. The load-bearing assumption is **co-scaling** — the residual gaps and the off-diagonal transport both vanish at rate ε ($D_\nu = D_0 \cdot I + \varepsilon \cdot \Delta$, $\Pi_\nu = \varepsilon \cdot K$) — after which a trivial scale–shift invariance gives that the eigenframe of T_ν equals the eigenframe of M_ν for every $\varepsilon > 0$. PMNS angles are therefore set by the *shape* of M_ν, not by the absolute neutrino mass scale. The neutrino-specific content lives in the co-scaling premise, which is asserted here, not derived.

We propose a minimal Hermitian left-frame target

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

with r_e, r_s, A, B real at the symmetric level, β an e–μ/e–τ asymmetry, δ a μ–τ diagonal splitting, and ψ a weak-commitment transport phase. In the exact $\beta = \delta = 0$ limit the operator is μ–τ symmetric. In the rotated basis

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

the symmetric operator block-diagonalises (verified symbolically) to

$$\begin{bmatrix} r_e & \sqrt{2}\cdot A & 0 \\ \sqrt{2}\cdot A & r_s + B & 0 \\ 0 & 0 & r_s - B \end{bmatrix}$$

giving

$\theta_{23} = 45°$, $\theta_{13} = 0°$, $\tan 2\theta_{12} = 2\sqrt{2}\cdot A / (r_s + B - r_e)$.

The ν_- state decouples exactly. Large solar mixing follows when the residual weak-commitment gap $r_s + B - r_e$ is comparable to the transport A — not when A is large. Small β, δ, ψ then couple ν_- back into the solar block, generating nonzero θ_{13}, an atmospheric deviation, and leptonic CP.

Numerical audit (this draft). With parameters chosen *only* to set $\theta_{12} \approx 33.4°$ via the solar formula, the symmetric core returns $\theta_{23} = 45.00°$, $\theta_{13} = 0.000°$, $J = 0$ exactly. Turning on the breaking terms β and δ (order tens of percent within their block, not parametrically small) generates θ_{13} and an atmospheric deviation; turning on ψ generates a nonzero leptonic Jarlskog invariant that is CP-odd and vanishes at $\psi = 0, \pi$. A coarse scan confirms the operator can reach the observed $(\theta_{12}, \theta_{23}, \theta_{13}) \approx (33.4°, 48.5°, 8.7°)$ with sizeable J — though, as §7 stresses, reaching the data is a re-tuning of free parameters, not a constrained prediction.

Honest status. The symmetric core has real predictive economy *given* μ–τ symmetry: that one assumed symmetry fixes two outputs ($\theta_{23} = 45°$, $\theta_{13} = 0$), and θ_{12} then depends only on the single scale/shift-invariant ratio $(r_s + B - r_e)/A$ — a one-assumption-to-two-outputs compression. But the symmetry is imposed on M_ν, not derived from closure geometry, so "predictive" here means predictive-conditional-on-assumed-symmetry. The breaking sector has no such compression at all: its three parameters (β, δ, ψ) match the three deviations (θ_{13}, $\theta_{23} - 45°$, δ_{CP}) one-to-one, so fitting them is generic. The next lepton task is to derive $A : B : r_e : r_s$ and the breaking scales from weak-commitment closure geometry, reducing that freedom until some PMNS quantity becomes a forced output — the standard the CKM paper already meets.

Table of Contents

---

0. Predictive-Content Ledger

This paper addresses one target.

LF-2: derive a weak-commitment neutrino frame U_ν whose near-degenerate transport operator produces large solar mixing, large atmospheric mixing, small nonzero reactor mixing, and a physical CP phase **without using PMNS angles as inputs**.

Object	Construction	Status in this paper
Lepton frame mismatch	U_PMNS = U_e† U_ν	inherited from Yukawa-frame architecture
Charged-lepton frame	U_e ≈ $U_0\ell$	anchoring premise; still to be derived
Neutrino transport	T_ν = D_0·I + ε·M_ν	weak-commitment form
Co-scaling premise	D_ν = D_0·I + ε·Δ **and** Π_ν = ε·K	**load-bearing physics; asserted, not derived**
Frame theorem	eigenframe(T_ν) = eigenframe(M_ν), ∀ ε > 0	exact, but **trivial once co-scaling is granted**
Symmetric core	M_ν$^{(0)}$ with eμ = eτ, μμ = ττ	proposed minimal core; symmetry **assumed**
Block-diagonal form	(e, +, −) reduction	verified symbolically
Atmospheric angle	$\theta_{23} = 45°$ in exact μ–τ limit	fixed by **assumed** μ–τ symmetry

Object	Construction	Status in this paper
Reactor angle	$\theta_{13} = 0$ in exact μ–τ limit	fixed by **assumed** μ–τ symmetry
Solar angle	$\tan 2\theta_{12} = 2\sqrt{2}\cdot A / (r\_s + B - r_e)$	one invariant ratio
Breaking terms	β, δ, ψ	proposed; must be derived
θ_{13} source	first order in **both** β and δ (spectrum-set), not perturbatively small	both feed θ_{13}
CP phase	weak-commitment transport phase ψ	nonzero J_lep verified; not yet derived
Parameter balance	breaking: 3 params vs 3 deviations	**not yet predictive**
Majorana phases	require full symmetric mass operator	not derived here
PMNS prediction	diagonalise *derived* M_ν	open

Three layers must stay separated.

1. **Exact frame algebra (given co-scaling).** *If* the gaps and the off-diagonal transport co-scale — $D\_\nu = D_0\cdot I + \varepsilon\cdot\Delta$ and $\Pi\_\nu = \varepsilon\cdot K$ — then $T\_\nu = D_0\cdot I + \varepsilon\cdot M\_\nu$ and the neutrino eigenframe is M_ν's eigenframe for all $\varepsilon > 0$. The eigenframe step is then trivial linear algebra; the physics is entirely in the co-scaling premise, which is asserted (and made a falsifier in §12.2), not derived. If M_ν has exact μ–τ symmetry, it gives $\theta_{23} = 45°$ and $\theta_{13} = 0$.
2. **Operator target.** A minimal weak-commitment M_ν should have approximate eμ/eτ equality, approximate μ–τ symmetry, and small complex breaking. Proposed, not derived.
3. **Physical derivation, still owed.** Closure geometry must derive r_e, r_s, A, B, β, δ, ψ — and, prior to those, the co-scaling premise and the μ–τ symmetry that the predictive content rests on. Choosing parameters to match PMNS is fitting, not derivation.

The discipline is unchanged: **a good-looking neutrino matrix is not a derived neutrino operator.**

---

1. The Inherited PMNS Problem

The Yukawa-operator paper made the central flavour move,

$$Y\_S = U\_S \cdot \Lambda\_S^2 \cdot U\_S^\dagger ,$$

where the diagonal eigenvalues Λ_S belong to the mass-trace and completion-density programme, and the sector frames U_S belong to completion-channel transport. Mixing appears because weak charged currents couple sectors before both frames are diagonalised:

$$V\_CKM = U\_u^\dagger U\_d , \quad U\_PMNS = U\_e^\dagger U\_\nu .$$

For quarks the programme has advanced through three steps: the CKM matrix was reinterpreted as a relative exponential rather than the diagonalisation of one absolute Hermitian transport Hamiltonian; the electroweak frame paper identified the weak-doublet split and its controlling commutator $\kappa = [\Omega_0, \Omega_q]$; and the curvature-residue paper proposed a concrete common-mode curvature lifting $|V_{td}|$ and $|J|$ while leaving a sharply bounded $|V_{ub}|$–ratio wall.

The lepton side stated the structure $U_{PMNS} = U_e^\dagger U_\nu$ and proposed that U_e is anchored while U_ν is weakly committed — but did not derive U_ν. The question for this paper is therefore:

What weak-commitment neutrino operator produces the PMNS frame?

The answer cannot be "choose a matrix that fits PMNS." It must come from the same closure-geometry discipline as the CKM programme. This draft does not start from measured PMNS angles; it starts from the weak-commitment co-scaling premise and the scale–shape separation it produces, and asks what minimal M_ν structure follows.

---

2. The Weak-Commitment Neutrino Regime

The lepton doublet $L_L = (\nu_L, e_L)$ is not symmetric the way the quark doublet is. The charged lepton is mass-resolved and anchored; the neutrino is neutral, weakly anchored, and weakly committed. The lepton-sector split should therefore **not** copy the quark pattern

$$U_\nu = U_{0\ell} \cdot e^{(-\Omega_\ell/2)} , \quad U_e = U_{0\ell} \cdot e^{(+\Omega_\ell/2)} ,$$

which would make PMNS CKM-like and small. The proposed lepton regime is instead

$$U_e \approx U_{0\ell} , \quad U_\nu = U_{0\ell} \cdot U_{wc} , \quad \text{so } U_{PMNS} = U_e^\dagger U_\nu \approx U_{wc} .$$

PMNS is then not a small opposite tilt inside a closure-committed doublet, but the frame of a weak-commitment neutrino transport operator viewed relative to an anchored charged-lepton frame.

In a closure-committed sector, generation stiffness gaps suppress rotations:

$$\theta_{ij} \sim \Pi_{ij} / (D_j - D_i) .$$

In the weak-commitment neutrino sector the stiffness gaps collapse,

$$D_\nu = D_0 \cdot I + \varepsilon \cdot \Delta ,$$

and consistency requires the off-diagonal transport to co-scale,

$$\Pi_\nu = \varepsilon \cdot K .$$

The total neutrino transport operator is therefore

$$T_\nu = D_\nu + \Pi_\nu = D_0 \cdot I + \varepsilon \cdot (\Delta + K) ,$$

and defining $M_\nu = \Delta + K$,

$$T_\nu(\varepsilon) = D_0 \cdot I + \varepsilon \cdot M_\nu .$$

The small parameter ε sets the *scale* of the weak commitment but not the *frame*. The contrast with quarks is exact:

$\theta_{ij}^{CKM} \sim \Pi_{ij}^{q} / (D_j^{q} - D_i^{q})$ (order-one gaps), $\theta_{ij}^{PMNS} \sim \varepsilon \cdot K_{ij} / (\varepsilon \cdot (\Delta_j - \Delta_i)) = K_{ij} / (\Delta_j - \Delta_i)$ (ε cancels).

The weak-commitment scale cancels out of the mixing.

---

3. The Scale–Shape Separation: the Co-Scaling Premise and a Trivial Theorem

The load-bearing physical assumption of this paper is **co-scaling**: in the weak-commitment regime the residual stiffness gaps and the off-diagonal transport vanish at the *same* rate in ε,

$D_\nu = D_0 \cdot I + \varepsilon \cdot \Delta$ **and** $\Pi_\nu = \varepsilon \cdot K$,

so that $T_\nu = D_0 \cdot I + \varepsilon \cdot M_\nu$ with $M_\nu = \Delta + K$. Everything below follows from this; it is asserted on physical grounds (the same weak commitment that flattens the spectrum also weakens the transport) and is the right thing to attack — §12.2 makes its failure a falsifier.

Granting co-scaling, the frame statement is then a one-line fact, true of *any* matrix, with nothing neutrino-specific in it:

Lemma (scale–shift invariance). For $T_\nu = D_0 \cdot I + \varepsilon \cdot M_\nu$ with $\varepsilon > 0$, the eigenvectors of T_ν are exactly those of M_ν. *Proof.* $D_0 \cdot I$ shifts every eigenvalue equally and fixes every eigenvector; multiplying by the positive scalar ε rescales eigenvalues and again fixes eigenvectors. ■

It is worth being blunt that this lemma is not where the content lives. The work is done by co-scaling: it is what makes ε cancel between numerator and denominator of the mixing ratio,

$$\theta_{ij}^{PMNS} \sim \varepsilon \cdot K_{ij} / (\varepsilon \cdot (\Delta_j - \Delta_i)) = K_{ij} / (\Delta_j - \Delta_i) ,$$

so that the absolute weak-commitment scale drops out of the angles. The scale fact (small neutrino masses) lives in ε; the shape fact (PMNS angles) lives in M_ν. A neutrino can be extremely light and still mix maximally — not a paradox, but exactly what co-scaling plus the trivial lemma deliver.

If co-scaling is only approximate, $T\_\nu(\varepsilon) = D_0 \cdot I + \varepsilon \cdot M\_\nu + O(\varepsilon^2)$, the separation holds at leading order with corrections set by the $O(\varepsilon^2)$ residue. If those higher-order terms instead dominate the eigenframe, the weak-commitment explanation fails (§12.1).

---

4. The Minimal Symmetric Neutrino Operator

The first target is not the full PMNS matrix; it is the symmetric core that explains why two angles are large and one is small. The minimal real symmetric core is

$$M\_\nu^{(0)} = \begin{bmatrix} r_e & A & A \\ A & r\_s & B \\ A & B & r\_s \end{bmatrix}$$

with r_e the residual electron-flavour stiffness, r_s the shared μ/τ residual stiffness, A the equal e↔μ and e↔τ weak-commitment transport, and B the μ↔τ transport inside the heavy-flavour lepton pair. It has exact μ–τ symmetry, $M\_e\mu = M\_e\tau$ and $M\_\mu\mu = M\_\tau\tau$.

That symmetry is not decoration. In VERSF language it states that weak-commitment neutrino transport does not strongly distinguish the two non-electron completion channels at leading order: the neutrino is not closure-anchored enough to keep the quark-like hierarchy, but its residual transport still remembers a paired μ/τ block.

The explicit broken operator adds three small parameters:

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

with β breaking eμ/eτ equality, δ breaking μμ/ττ equality, ψ a weak-commitment transport phase, and $\beta = \delta = 0$ returning the symmetric core. This is the minimal Hermitian left-frame target — enough to study PMNS angles. If neutrinos are Majorana, the full mass operator is complex-symmetric and diagonalised by Takagi factorisation; that is a later layer (§9).

---

5. Exact Diagonalisation in the μ–τ Symmetric Limit

Introduce the rotated states

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

In the basis $(|\nu_e\rangle, |\nu_+\rangle, |\nu_-\rangle)$, the symmetric operator becomes — confirmed by exact symbolic computation —

$$M\_\nu^{(0)} = \begin{bmatrix} r_e & \sqrt{2}\cdot A & 0 \\ \sqrt{2}\cdot A & r\_s + B & 0 \\ 0 & 0 & r\_s - B \end{bmatrix}$$

This is the key structural result. The $|\nu_-\rangle$ state **decouples exactly**: the third eigenstate carries equal-and-opposite μ and τ components and no electron component. In PMNS language, assigning that decoupled antisymmetric state to ν_3,

$\theta_{23} = 45°$, $\theta_{13} = 0°$.

The remaining solar sector is the 2×2 block

$$\begin{bmatrix} r_e & \sqrt{2}\cdot A \\ \sqrt{2}\cdot A & r\_s + B \end{bmatrix}$$

whose mixing angle satisfies

$\tan 2\theta_{12} = 2\sqrt{2}\cdot A / (r\_s + B - r_e)$.

This proves the first LF-2 structural statement:

Large solar mixing does not require large absolute neutrino mass. It requires the residual weak-commitment gap $(r\_s + B - r_e)$ to be comparable to the transport A.

In a closure-committed quark sector the corresponding denominator is stiff and order-one; in the weak-commitment neutrino sector it is a collapsed residual gap that can be small enough to produce large mixing without making the coupling large.

Numerical confirmation. Choosing parameters *only* to set $\theta_{12} \approx 33.4°$ through the solar formula — $r_e = 0$, $A = 0.30$, $B = 0.10$, hence $r\_s = 2\sqrt{2}\cdot A/\tan(66.8°) - B + r_e \approx 0.264$ — the exact diagonalisation of the full 3×3 core returns

$\theta_{12} = 33.40°$, $\theta_{23} = 45.00°$, $\theta_{13} = 0.000°$, $J = 0$,

to machine precision. Moreover θ_{12} is invariant under $M\_\nu \rightarrow \alpha\cdot M\_\nu + \beta_0\cdot I$ (verified): it depends only on the single ratio $(r\_s + B - r_e)/A$, so the symmetric core has **one** physical solar parameter, with θ_{23} and θ_{13} fixed by the assumed μ–τ symmetry. The qualifier "assumed" matters: this is a one-symmetry-to-two-outputs compression, which is real predictive economy, but it is conditional on a symmetry imposed on M_ν rather than derived from closure geometry (§7).

---

6. Small Breaking Terms, the Reactor Angle, and What Feeds θ_{13}

The exact μ–τ limit is too symmetric: it gives $\theta_{13} = 0$, but observation requires a nonzero reactor angle. So the operator needs μ–τ breaking. Transforming the broken M_ν into the (e, +, −) basis, the couplings that re-link the decoupled ν_- state are, at leading order,

$M_{e-} \approx -\sqrt{2}\cdot A\cdot\beta\cdot e^{(-i\psi)}/2$, $M_{+-} \approx \delta$

(the conjugate entry M_{-e} carries $e^{(+i\psi)}$; the Hermitian pair is verified symbolically).

These two couplings do not map one-to-one onto θ_{13} and the atmospheric deviation. Because ν_+ carries a large electron component (the solar angle is ~33°), the δ-driven (+,−) coupling feeds back into the electron line and generates θ_{13} as well as shifting θ_{23}. Concretely, with the §5 spectrum:

β	δ	θ_{12}	θ_{23}	θ_{13}
0.10	0.00	33.92°	47.34°	1.05°
0.00	0.10	33.02°	40.79°	11.27°
0.10	0.10	33.74°	43.25°	11.90°
0.15	0.08	34.16°	45.22°	10.04°

So **both** breaking terms feed θ_{13}, and in this spectrum δ is the larger contributor. The split between β and δ is set by the spectral denominators (how close $r_s - B$ sits to a solar eigenvalue), not by a fixed rule. The leading-order statement is:

θ_{13} and (θ_{23} − 45°) both turn on with the breaking (β, δ); their separate coefficients are fixed by the spectrum, with θ_{13} enhanced when the antisymmetric eigenvalue $r_s - B$ approaches a solar-block eigenvalue.

A caveat on "small" and "leading order." The breaking that produces the observed reactor angle is not parametrically small. To reach $\theta_{13} \approx 10°$ one needs β, δ ≈ 0.1–0.15 against block entries A ≈ 0.2–0.3 and B ≈ 0.1–0.2 — a 50–75% perturbation *within the block it acts on*. And the response is visibly non-linear: at fixed δ = 0, θ_{13} runs 0.55° → 1.05° → 1.92° → 2.63° as β goes 0.05 → 0.10 → 0.20 → 0.30, already saturating below linear. So the first-order expressions above are illustrative bookkeeping, not a controlled small-parameter expansion; the exact diagonalisation is what the numbers come from. The honest reading is that θ_{13}'s "smallness" is *relative* — it is the largest of the symmetry-breaking effects, not a small one in absolute terms.

What is robust is the *qualitative* mechanism:

θ_{13} is small not because neutrino mixing is generically suppressed, but because it is **symmetry-forbidden at exact μ–τ symmetry** and appears only once that symmetry is broken — so it is naturally the smallest of the three angles even though the breaking itself is order tens of percent.

This names two lepton-frame sub-targets:

- **LF-2a (symmetric core).** Derive a weak-commitment $M\_\nu^{(0)}$ with approximate $e\mu = e\tau$ and $\mu\mu = \tau\tau$, giving large θ_{12}, maximal θ_{23}, zero θ_{13}.
- **LF-2b (controlled breaking).** Derive small β, δ, ψ producing nonzero θ_{13}, an atmospheric deviation, and leptonic CP — and, critically, derive enough *relations* among them that the breaking sector becomes predictive (see §7).

---

7. Parameter Counting: Where the Operator Predicts and Where It Only Fits

This section is the analogue of the CKM paper's parameter-count argument, and it is where this draft is most candid.

Physical PMNS observables (Dirac case). Ignoring Majorana phases, the frame carries four physical numbers: θ_{12}, θ_{23}, θ_{13}, δ_CP.

Operator parameters. The ansatz M_ν has seven real parameters: r_e, r_s, A, B, β, δ, ψ. The eigenvectors are invariant under $M\_\nu \to \alpha \cdot M\_\nu + \beta_0 \cdot I$ ($\alpha > 0$, β_0 real), which removes two. So **five effective real parameters** control the frame.

Five effective parameters against four observables means the frame is **not over-determined**: there is a one-parameter family of operators reproducing any achievable (θ_{12}, θ_{23}, θ_{13}, δ_CP). A coarse scan confirms reachability. The full parameter tuple is

$r_e = 0$, $r\_s = 0.0061$, $A = 0.17$, $B = 0.20$, $\beta = 0.075$, $\delta = 0$, $\psi = 135°$,

which gives $(\theta_{12}, \theta_{23}, \theta_{13}) = (33.6°, 48.5°, 8.7°)$ with a sizeable leptonic Jarlskog ($J\_lep \approx 0.025$). Two points must be printed alongside it, or the scan looks like cherry-picking. First, the solar value here is re-fixed by a *different* invariant ratio than §5: the §5 spectrum used $r\_s \approx 0.264$, and substituting that value into this (A, B) gives (22.7°, 44.0°, 2.6°) — wildly off — so the scan has silently shrunk the solar denominator ($r\_s + B - r_e$) by ~40× to hold θ_{12}. (Equivalently, $r_e = 0.165$, $r\_s = 0.171$ lands the same angles; only the invariant combination matters.) Second, and consequently, reaching the data this way is *generic* — a re-tuning of free parameters, not a constrained prediction. **Fitting PMNS with this ansatz is necessary, not impressive.**

The predictive content is unevenly distributed, and it matters where it lives — but both parts rest on *assumed* structure, so the labels need care:

- **Symmetric core (predictive given μ–τ symmetry).** Four parameters minus scale + shift leave two; of those, θ_{12} depends only on the single invariant ratio $(r\_s + B - r_e)/A$. So the core spends one parameter on θ_{12} and returns $\theta_{23} = 45°$, $\theta_{13} = 0$ as outputs. This is a genuine one-assumption-to-two-outputs compression — but the assumption is the μ–τ symmetry imposed on M_ν, which is itself not derived from closure geometry. So "predictive" means *predictive conditional on an assumed symmetry*, not predictive from first principles.

- **Breaking sector (not predictive).** Three parameters (β, δ, ψ) produce exactly three deviations (θ_{13}, θ_{23} − 45°, δ_CP). Three for three is a parametrisation with no quantity left over.

The real distinction between the two is therefore one of *compression*, not of "structure vs. fitting": the core assumes a symmetry and gets a 1→2 reduction; the breaking assumes values and gets none. Both await closure geometry. This sets the precise bar:

The lepton sector reaches CKM-level rigour when closure geometry supplies (i) the μ–τ symmetry of the core and the ratio (r_s + B − r_e)/A, and (ii) at least one **relation** among (β, δ, ψ) — a fixed ratio or a determined phase — so that one PMNS quantity becomes a forced output rather than a fitted input.

By contrast, the CKM curvature paper already cleared an analogous bar: two parameters (|z|, phase) pinned by two observables (J, |V_td|), leaving |V_ub| and the ratio as genuine forced predictions that hit a structural wall. The neutrino operator is at least one derived relation short of the same standing.

---

8. Leptonic CP as a Weak-Commitment Transport Phase

The CKM papers treat CP violation as a curvature/loop effect in generation transport. The neutrino sector should not simply copy the CKM phase, but the same lesson applies: a physical phase must survive rephasing and appear in a closed invariant.

In the minimal operator, ψ enters through the asymmetric e–τ channel, M_eτ = A(1 + β·e^(−iψ)). When β = 0 the phase is absent. When β ≠ 0 it enters the e ↔ ν_- coupling, M_{e−} ≈ −√2·A·β·e^(−iψ)/2 — the first place a physical leptonic phase can reach the PMNS frame.

Verification. Diagonalising the full complex operator and computing the leptonic Jarlskog invariant J_lep = Im(U_e1·U_μ2·U_e2*·U_μ1*) over ψ (at β = 0.15, δ = 0.08):

ψ	**θ_{12}**	**θ_{23}**	**θ_{13}**	**J_lep**
0°	34.16°	45.22°	10.04°	0
45°	33.91°	44.31°	9.89°	-4.0×10^{-3}
90°	33.25°	41.84°	9.30°	-6.5×10^{-3}
135°	32.46°	38.91°	8.20°	-5.3×10^{-3}
180°	32.10°	37.52°	7.46°	0
270°	33.25°	41.84°	9.30°	$+6.5\times10^{-3}$

So ψ produces a genuine nonzero J_{lep} that is CP-odd (sign-flipping under $\psi \to -\psi$) and vanishes at the CP-conserving points $\psi = 0, \pi$. This is the required behaviour: ψ is a candidate physical Dirac-type phase, not a removable rephasing artefact, **provided** it is derived rather than chosen.

A standing caution remains. A phase written in a matrix is not automatically physical; the invariant test above is necessary but the *value* must still come from a weak-commitment transport loop, not from matching δ_{CP}. The target:

LF-CP. Derive a weak-commitment transport loop whose residual phase survives charged-lepton and neutrino rephasings and appears as the PMNS Dirac phase. In a Majorana treatment, additional phases can be physical, but those belong to the full symmetric mass operator, not the left-square frame.

---

9. Relation to Majorana Mass and Takagi Factorisation

The left-frame analysis suffices for the PMNS angles because charged-current mixing depends on $U_e^\dagger U_\nu$. But neutrinos may be Majorana, and a Majorana mass matrix is not diagonalised by ordinary Hermitian similarity.

- **Dirac case.** Form the left-square operator $Y_\nu = Y_\nu Y_\nu^\dagger = U_\nu \Lambda_\nu^2 U_\nu^\dagger$, parallel to the quark and charged-lepton sectors.
- **Majorana case.** The mass matrix is complex-symmetric and diagonalised by Takagi factorisation, $M_\nu^{Maj} = U_\nu \Lambda_\nu U_\nu^T$.

The PMNS matrix still depends on the left neutrino frame U_ν, so the frame architecture and the mixing angles are unchanged either way. But the two Majorana phases cannot be recovered from the left-square operator alone — they require the full complex-symmetric mass operator. This draft therefore separates three objects: the weak-commitment frame operator M_ν (controls U_ν, treated here); the left-square flavour operator Y_ν (parallel to other sectors); and the full Majorana mass operator (carries Majorana phases and neutrinoless-double-beta information, not addressed here).

---

10. Relation to CKM and the Curvature Programme

The two sectors are now parallel in architecture but different in mechanism.

CKM: $V_{CKM} = U_u^\dagger U_d$. Both frames are closure-committed and nearly aligned. The relative generator Ω_q gives the small hierarchy; the curvature residue $\kappa = [\Omega_0, \Omega_q]$ supplies the CP and unitarity-triangle correction. The question is *why the relative split is small but curved*.

PMNS: $U\_PMNS = U\_e^\dagger U\_\nu$. The charged-lepton frame is anchored; the neutrino frame is weakly committed. The leading effect is not a small residue around an aligned pair but the *collapse of the neutrino stiffness hierarchy itself*. The question is *why the residual frame is large, and what symmetry makes one angle small*.

This paper answers the second structurally: the frame is large because stiffness gaps collapse; θ_{23} is large by approximate μ–τ symmetry; θ_{12} is large because A competes with a collapsed solar gap; θ_{13} is small because it vanishes in the exact symmetry limit and appears only through breaking; CP enters through complex weak-commitment transport. The shared principle is one ratio,

mixing angle ~ (off-diagonal transport) / (stiffness gap) ,

evaluated in the closure-committed regime for CKM and in the collapsed-gap regime for PMNS.

---

11. What Is Proved, Conditional, and Open

11.1 Proved here, given definitions.

- *Given the co-scaling premise* ($D\_\nu = D_0 \cdot I + \varepsilon \cdot \Delta$, $\Pi\_\nu = \varepsilon \cdot K$, so $T\_\nu = D_0 \cdot I + \varepsilon \cdot M\_\nu$), T_ν and M_ν share eigenvectors for every $\varepsilon > 0$ — a one-line scale/shift fact, not a deep theorem. Hence a small neutrino mass scale does not imply small PMNS angles. The content is in the co-scaling premise, which is asserted.
- The symmetric core block-diagonalises in the (e, +, −) basis to diag-plus-2×2 form (verified symbolically), giving $\theta_{23} = 45°$, $\theta_{13} = 0$, and $\tan 2\theta_{12} = 2\sqrt{2} \cdot A/(r\_s + B - r_e)$.
- θ_{12} depends only on the scale/shift-invariant ratio $(r\_s + B - r_e)/A$ (verified): the symmetric core has one physical solar parameter and returns the other two angles as outputs of the *assumed* μ–τ symmetry.
- β and δ break the symmetry and generate nonzero θ_{13} and atmospheric deviation; **both** terms feed θ_{13}, with the split spectrum-determined (verified, §6). The breaking needed for $\theta_{13} \approx 10°$ is order tens of percent within its block and the response is non-linear, so the leading-order formulae are illustrative, not a controlled expansion.
- ψ generates a nonzero leptonic Jarlskog invariant that is CP-odd and vanishes at $\psi = 0, \pi$ (verified, §8).

11.2 Structural content. The paper turns the PMNS half of the programme from a qualitative weak-commitment statement into a concrete operator target. PMNS largeness is the collapse of stiffness denominators, not large couplings; the symmetric core explains the large–large–small pattern; the small reactor angle is the symmetry-breaking signal. The parameter count (§7) is the load-bearing analytic point: it isolates *where* the construction compresses (the symmetric core, one assumed symmetry yielding two outputs) and *where* it merely parametrises (the breaking sector, three parameters for three deviations), and states exactly what closure geometry must supply — the core symmetry and ratio, plus at least one relation among the breaking parameters — to close the gap.

11.3 Conditional premises.

- The charged-lepton frame is approximately anchored, $U_e \approx U_0\ell$.
- The neutrino operator has the weak-commitment co-scaling form $T_\nu = D_0 \cdot I + \varepsilon \cdot M_\nu$.
- The leading operator has approximate $e\mu = e\tau$ and approximate μ–τ symmetry.
- The breaking terms β, δ, ψ are subdominant to the symmetric core but not parametrically small (order tens of percent within their block; see §6).
- The substrate supports bath-type continuous transport.

11.4 Open problems.

- Derive r_e, r_s, A, B from weak-commitment closure geometry — in practice, derive the single ratio $(r_s + B - r_e)/A$ that sets θ_{12}.
- Derive the breaking parameters β, δ, **and at least one relation among them** (a ratio or a phase), so the breaking sector predicts rather than fits (§7).
- Derive the weak-commitment phase ψ from a transport loop (LF-CP).
- Diagonalise the *derived* operator and compare with PMNS without using PMNS angles as inputs.
- Decide whether the spectrum prefers normal ordering, inverted ordering, or leaves ordering to the mass-eigenvalue layer.
- Extend the left-frame operator to the full Dirac or Majorana mass operator; derive or reject the Majorana phases.

Honest status: a first weak-commitment neutrino-operator draft. The scale–shape separation (given the asserted co-scaling premise) and the symmetric-limit angles are established; the breaking and phase behaviour are verified numerically. The symmetric core is predictive *conditional on its assumed μ–τ symmetry*; the breaking sector is, as written, a parametrisation. It is therefore an operator target with exact consequences and a clear bar to clear — not yet a derived PMNS matrix.

---

12. Falsification Conditions

12.1 No stiffness-gap collapse. If the neutrino closure calculation gives order-one generation stiffness gaps rather than $D_\nu = D_0 \cdot I + \varepsilon \cdot \Delta$, PMNS should be CKM-like and small, and the weak-commitment explanation fails.

12.2 Wrong scaling of off-diagonal transport. If the off-diagonal transport does not co-scale, $\Pi_\nu \neq \varepsilon \cdot K$, then ε no longer cancels from the mixing ratio: the weak-commitment limit becomes singular or trivial and the scale–shape separation supplies no finite nontrivial PMNS frame. This is the premise the whole construction rests on (§3), so it is the primary falsifier.

12.3 No approximate μ–τ symmetry. If the derived M_ν lacks approximate $e\mu = e\tau$ and $\mu\mu = \tau\tau$ structure, the simple account of large atmospheric and small reactor mixing fails.

12.4 Reactor angle requires fitting. If θ_{13} can only be obtained by choosing β (or δ) by hand from the measured reactor angle, LF-2 is unmet. Since §6 shows both terms feed θ_{13}, this falsifier now reads: closure geometry must fix the (β, δ) combination, not just their existence.

12.5 CP phase requires fitting. If ψ is selected from PMNS data rather than derived from a weak-commitment transport loop, the CP part stays phenomenological.

12.6 Parameter count never closes. If closure geometry cannot supply at least one relation among (β, δ, ψ), the breaking sector remains a 3-for-3 parametrisation and the lepton side never reaches the predictive standing the CKM sector already has. (This is the sharpened, §7 form of "it only fits.")

12.7 Charged-lepton frame not anchored. If U_e is not approximately completion-diagonal, PMNS is not simply the neutrino weak-commitment frame; part of it moves into the charged-lepton sector and requires a separate derivation of U_e.

12.8 Majorana mismatch. If the left-frame operator gives a good PMNS pattern but cannot be embedded in a consistent Dirac or Majorana mass operator, the frame result is incomplete.

12.9 Ledger verdict. If the substrate is ledger-like at the relevant pre-commitment grain, continuous off-diagonal neutrino transport is forbidden; then M_ν must be diagonal up to permutations and the construction fails at the level of principle.

---

13. Conclusion

The VERSF flavour programme has reached the point where PMNS can no longer remain an architectural promise. The Yukawa paper made CKM and PMNS relative sector frames; the electroweak frame paper explained why the lepton doublet is not quark-like — the charged lepton anchored, the neutrino weakly committed and nearly degenerate. This paper writes the first weak-commitment neutrino operator target and tests its consequences numerically.

The central result is the scale–shape separation, $T\_\nu(\varepsilon) = D_0 \cdot I + \varepsilon \cdot M\_\nu$: granting the co-scaling premise, ε controls the weak-commitment scale while M_ν controls the frame, so small neutrino masses are compatible with large mixing. The minimal symmetric operator has exact μ–τ symmetry; in the (e, +, −) basis it block-diagonalises, giving maximal atmospheric mixing, zero reactor mixing, and a solar angle set by one invariant ratio. Symmetry-breaking terms — order tens of percent, not parametrically small — then generate the realistic deviations $\theta_{13} \neq 0$, $\theta_{23} \neq 45°$, $\delta\_CP \neq 0$, all confirmed here, including a nonzero, CP-odd leptonic Jarlskog invariant.

The draft is also clear about its ceiling. By a direct parameter count, the symmetric core is genuinely predictive (μ–τ symmetry fixes two angles), but the breaking sector currently carries one free parameter per deviation and so only parametrises. Reproducing the measured PMNS matrix is therefore not yet evidence; it is generic. The lepton sector reaches the standing the CKM sector already holds precisely when closure geometry supplies one or more **relations** among the breaking parameters, converting at least one PMNS quantity from a fitted input into a forced output.

That is the right kind of draft result: not a fit dressed as a derivation, and not a vague promise, but a precise operator target with proved symmetric-limit consequences, a verified breaking-and-phase response, and a sharp, quantitative statement of exactly what remains to be derived.

Next calculation. Derive the ratio $(r_s + B - r_e)/A$ and at least one breaking relation ($\beta : \delta$, or ψ) from weak-commitment closure geometry — then diagonalise and compare with PMNS without using PMNS angles as inputs.

References

VERSF flavour architecture

- Keith Taylor, *The Electroweak Flavour-Frame Operator in VERSF.* Weak-doublet frame splitting; QF-1 and QF-2; weak-commitment PMNS frame theorem; minimal PMNS operator target.
- Keith Taylor, *The CKM Curvature Residue in VERSF.* Common-mode curvature; z_C3; CKM residue; Jarlskog and |V_td| repair; parameter-count and the |V_ub|–ratio wall.
- Keith Taylor, *The Yukawa Operator from Completion-Channel Misalignment in VERSF.* Left Yukawa-square operator; CKM/PMNS as relative sector frames; CKM relative exponential; PMNS weak-commitment architecture.
- Keith Taylor, *Deriving Flavour Mixing from Closure Geometry: CKM and PMNS Structure in the VERSF Framework.* Generation-space transport; stiffness gaps; CKM–PMNS dichotomy.
- Keith Taylor, *Completion Density and Flavour Mixing from Refinement Geometry.* Completion basis; sector bases; frame-misalignment formulation.

VERSF Standard Model core

- Keith Taylor, *The Standard Model from Hexagonal Geometry.* Gauge–Higgs–confinement core; weak mixing; Higgs relation; gauge-group uniqueness; flavour left open.
- Keith Taylor, *A Unified Derivation of Closure Geometry, Gauge Redundancy, and Mass Structure in the Hexagonal Framework.* Closure-count theorem; gauge redundancy; mass-structure programme.

VERSF mass-trace programme

- Keith Taylor, *The Strange Current-Mass Operator in VERSF.* Strange support trace; Johnson-scheme commutant; twentyfold down-to-strange ratio.
- Keith Taylor, *The Strange Confinement Operator in VERSF.* Boundary capacity; complement role record; canonical-invariant self-weight.

VERSF transport foundations

- Keith Taylor, *Bath or Ledger.* Mixing iff shared pre-factual conserved bath.
- Keith Taylor, *The Bath Criterion.* Continuous mixing from shared closure transport; relation to non-abelian structure and generation transport.

Numerical note. The block-diagonalisation of the μ–τ symmetric core was verified symbolically. All PMNS angles were obtained by exact Hermitian diagonalisation of M_ν, with eigenvectors permuted to the standard convention (ν_3 = the column of smallest electron content) before reading θ_{12} = atan2(|U_e2|, |U_e1|), θ_{23} = atan2(|U_μ3|, |U_τ3|), θ_{13} = asin(|U_e3|), and J_lep = Im(U_e1·U_μ2·U_e2*·U_μ1*). The symmetric-limit values θ_{23} = 45.00°, θ_{13} = 0.000°, J = 0 are exact to machine precision; θ_{12}'s dependence on the single invariant ratio (r_s + B − r_e)/A was confirmed by invariance under M_ν → α·M_ν + β_0·I. The breaking and phase tables (§6, §8) and the reachability scan (§7) were produced by direct diagonalisation over the stated parameter grids. Empirical PMNS comparison values are taken at the global-fit level (normal ordering) and are used only as audit targets, never as operator inputs.*