The Weak-Commitment Closure Kernel in VERSF

Deriving the Solar Ratio, the μ–τ Breaking Relation, the Octant–Phase Lock, and the Leptonic Phase of the PMNS Operator

Keith Taylor VERSF Theoretical Physics Programme — Standard Model Flavour Series
Working manuscript, v0.4

Successor to The Weak-Commitment Neutrino Operator in VERSF, 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 previous neutrino paper explained why neutrino mixing can be large even when neutrino masses are tiny. Mass size and mixing shape are not the same quantity. A neutrino can carry a very small mass scale while still having a large internal mixing frame, because the stiffness that normally separates the generations has collapsed.

That paper proposed a simple weak-commitment neutrino operator. In its symmetric form it naturally produced a large solar angle, a maximal atmospheric angle, and a zero reactor angle. Small symmetry-breaking terms then produced the real-world deviations: a nonzero reactor angle, an atmospheric angle slightly away from 45°, and leptonic CP violation.

But that paper was honest about its weakness. The symmetric core predicted; the breaking sector still had too many free knobs. The basic shape was meaningful, yet the detailed corrections could still be dialled to fit data.

This paper moves past that. Instead of asking whether a weak-commitment matrix *can* fit PMNS, it asks why the matrix should have its particular shape. The target is the **weak-commitment closure kernel**: the completion-geometry rule that fixes the ratios inside the neutrino operator.

The main proposal is that the symmetric core is controlled by one closure ratio,

$\rho\odot = (r_s + B - r_e) / A,$

and that weak-commitment geometry fixes it to

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

This gives a large solar angle without ever using the solar angle as an input. The neutrino's residual completion cloud has a threefold structure, but its coherent μ/τ transport is paired, so the ratio is a three-over-two root ratio.

There is one structural result that strengthens the framework. When the full broken operator is diagonalised exactly, the atmospheric octant is not free. It is locked to the quadrant of the leptonic phase, and the proposed phase ($\psi = 3\pi/4$) puts it in the **upper** octant, $\theta_{23} \approx 48.4°$, with a definite sign of leptonic CP violation — and this holds whatever the small leakage amplitude is. So the kernel does not merely accommodate the atmospheric angle; it predicts which side of 45° it must fall on and ties that choice to the sign of leptonic CP. The measured reactor angle θ_{13} is reproduced at the benchmark, but it is a consistency check, not the thing that picks the octant. This lands the benchmark essentially on the upper local minimum of the current global fit.

Together these rules reduce the PMNS operator from a flexible fitting matrix to a sharp closure-kernel candidate that reproduces the measured θ_{12} and θ_{13} to within a fraction of a degree, predicts a specific atmospheric octant, and predicts a specific sign and size of leptonic CP violation. The result is not yet a completed derivation, because the kernel rules must still be forced from the full weak-commitment closure Hamiltonian. But the lepton-sector problem is no longer "fit PMNS." It is now:

derive or reject the weak-commitment closure kernel that fixes the solar ratio, the leading μ–τ breaking relation, the octant–phase lock, and the leptonic phase.

Abstract

The previous weak-commitment neutrino paper proposed a Hermitian left-frame operator

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

arising from the scale-separated weak-commitment form

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

The Weak-Commitment Frame Theorem shows that the PMNS frame is the eigenframe of M_ν, independent of the small scale ε. In the exact μ–τ symmetric limit ($\beta = \delta = 0$),

$M_\nu^{(0)}$ has entries $M_{e\mu} = M_{e\tau} = A$, $M_{\mu\mu} = M_{\tau\tau} = r_s$, $M_{\mu\tau} = B$,

the rotated basis

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

block-diagonalises the operator and gives

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

The previous draft identified the **solar ratio** $\rho_\odot = (r_s + B - r_e)/A$ as the first lepton-sector derivation target, and the breaking parameters (β, δ, ψ) as the second.

This paper proposes a weak-commitment closure kernel that reduces those freedoms.

The symmetric kernel rule is

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

read as the root ratio of a threefold weak-commitment residual cloud to a coherent two-channel μ/τ pair. This gives

$\theta_{12} = \frac{1}{2}\cdot\tan^{-1}(2\sqrt{2} / \sqrt{(3/2)}) \approx 33.3°$,

without using θ_{12} as an input. A second kernel rule fixes the internal pair ratio

$\sigma\_\nu = B/A = 6/5$,

which does not affect the symmetric-limit angles but controls the breaking response through the eigenvalue denominators. The leading weak-commitment breaking is proposed to be off-diagonal rather than diagonal,

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

because weak neutrality forbids a first-order μ/τ self-stiffness split while allowing an e–τ transport asymmetry. The leptonic phase is proposed to be the weak-commitment bisector

$\psi\_wc = (\pi/2 + \pi)/2 = 3\pi/4$,

between weak-neutral quadrature and complement reversal.

With the benchmark leakage amplitude $\beta\star = \sqrt{3}/20$, the no-continuous-fit kernel point

$\rho_\odot = \sqrt{(3/2)}$, $\sigma\_\nu = 6/5$, $\delta = 0$, $\psi = 3\pi/4$, $\beta = \sqrt{3}/20$

gives, under **exact Hermitian diagonalisation**,

$\theta_{12} \approx 33.4°$, $\theta_{23} \approx 48.4°$, $\theta_{13} \approx 8.6°$, $|J\_lep| \approx 2.4 \times 10^{-2}$, $\delta\_CP \approx 227°$.

A new structural result is reported: the atmospheric octant and the sign of leptonic CP violation are **not independent outputs** of the kernel. They are locked to the quadrant of ψ, β-independently, so the phase rule $\psi = 3\pi/4$ fixes the **upper** octant on its own; the measured θ_{13} enters only as a benchmark consistency check, not as the octant selector. The benchmark $\sin^2\theta_{23} = 0.559$ then sits essentially on the upper local minimum of the current global fit.

The result is conditional. The ratio $\rho_\odot = \sqrt{(3/2)}$, the pair ratio $B/A = 6/5$, the relation $\delta = 0$, the phase $\psi = 3\pi/4$, and especially the leakage amplitude $\beta_\star = \sqrt{3}/20$ must be derived from the full weak-commitment closure Hamiltonian. This paper is therefore not a completed PMNS derivation. Its contribution is to reduce the previous PMNS ansatz to a sharply testable closure-kernel candidate with a fixed octant and a fixed CP sign.

Table of Contents

---

0. Predictive-Content Ledger

This paper addresses one target.

LF-3 — derive enough of the weak-commitment closure kernel to turn the PMNS operator from a flexible ansatz into a constrained prediction problem.

Object	Construction	Status in this paper
Weak-commitment transport	$T\_\nu = D_0 I + \varepsilon M\_\nu$	inherited
PMNS frame	eigenframe of M_ν	inherited exact theorem
Symmetric core	$M\_\nu^{(0)}$ with $e\mu = e\tau$, $\mu\mu = \tau\tau$	inherited
Solar ratio	$\rho_\odot = (r_s + B - r_e)/A$	first kernel target
Kernel solar rule	$\rho_\odot = \sqrt{(3/2)}$	proposed derivation target
Solar prediction	$\theta_{12} \approx 33.3°$	follows from kernel rule

Object	Construction	Status in this paper
Pair ratio	$\sigma_\nu = B/A$	controls breaking response
Pair kernel rule	$\sigma_\nu = 6/5$	proposed
Diagonal μ–τ breaking	δ	set to $0 + O(\beta^2)$ at leading order
Off-diagonal breaking	β	leading reactor / asymmetry source
Phase	$\psi_{wc} = 3\pi/4$	proposed weak-commitment bisector
Octant–phase lock	$\mathrm{sgn}(\theta_{23} - 45^\circ) \leftrightarrow \mathrm{quadrant}(\psi)$	**new exact result**
Predicted octant	$\theta_{23} > 45^\circ$ (upper), $\sin^2\theta_{23} \approx 0.559$	**forced by the $\psi = 3\pi/4$ quadrant (LF-3d)**
Benchmark leakage	$\beta^\star = \sqrt{3}/20$	candidate, least secure
PMNS audit	$(33.4^\circ, 48.4^\circ, 8.6^\circ)$, $\delta_{CP} \approx 227^\circ$	benchmark result, not final derivation
Proof obligation	derive kernel rules from closure Hamiltonian	open

Three levels must stay separated.

1. **Exact algebra.** Given M_ν, the scale ε does not affect the eigenframe. Given μ–τ symmetry, $\theta_{23} = 45^\circ$, $\theta_{13} = 0$, and the solar angle depends on one ratio. Given the full broken operator, the octant and CP sign are fixed functions of ψ — this is exact, not a proposal.
2. **Kernel proposal.** The weak-commitment closure kernel fixes $\rho_\odot = \sqrt{3/2}$, $\sigma_\nu = 6/5$, $\delta = 0$, and $\psi = 3\pi/4$.
3. **Physical derivation still owed.** The closure Hamiltonian must force these rules. If they are only chosen because they work, the paper is a fit in disguise.

The discipline remains:

A constrained-looking kernel is not a derivation until the closure geometry forces it.

---

1. The Inherited LF-2 Bottleneck

The previous weak-commitment neutrino paper established the operator target M_ν above. It also proved the clean symmetric result. When $\beta = \delta = 0$, the operator has exact μ–τ symmetry,

$M_{e\mu} = M_{e\tau}$, $M_{\mu\mu} = M_{\tau\tau}$.

In the $(e, +, -)$ basis the symmetric core becomes block form

$\lceil r_e\ \sqrt{2}\cdot A\ 0 \rceil | \ | \sqrt{2}\cdot A\ r_s + B\ 0 \ | \ | \lfloor 0\ 0\ r_s - B \rfloor$

which gives

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

That result was already meaningful. The large–large–small PMNS pattern arises structurally: large solar mixing comes from a collapsed residual gap, maximal atmospheric mixing comes from μ–τ symmetry, and small reactor mixing appears only when that symmetry is broken.

But the same paper identified the bottleneck. After removing an irrelevant global scale and a common shift, the full ansatz still has enough freedom to fit the observed PMNS frame. The symmetric core predicts; the breaking sector parametrises.

The next task is therefore not another scan but the derivation of relations among the parameters. The minimum useful derivations are the solar ratio $\rho_\odot = (r_s + B - r_e)/A$ and at least one relation among (β, δ, ψ). This paper attempts exactly that — and finds a third relation, the octant–phase lock, that was hiding in the exact algebra.

---

2. From Neutrino Operator to Closure Kernel

The operator M_ν is a frame-shape object. The closure kernel is the rule that generates its entries.

In this paper, **kernel** means the residual weak-commitment completion rule that survives after the scalar stiffness has collapsed:

$T\_\nu = D_0 I + \varepsilon M\_\nu$.

The large common part $D_0 I$ fixes the overall weak-commitment scale but carries no frame information. The frame information lives entirely in M_ν.

The kernel must explain four things.

- **First**, why the leading core has approximate μ–τ symmetry, eμ ≈ eτ and μμ ≈ ττ.
- **Second**, why the solar ratio sits near the value required for large but non-maximal solar mixing.
- **Third**, why the leading breaking is small and structured rather than arbitrary.
- **Fourth**, why the breaking phase is physical and not removable by rephasing.

The kernel is not another matrix. It is the closure rule behind the matrix.

---

3. Kernel Invariants: Scale, Shift, Solar Ratio, and Pair Ratio

The eigenvectors of M_ν are invariant under

$M_\nu \to \alpha M_\nu + \lambda I, \alpha > 0$.

The scale α rescales eigenvalues but not eigenvectors; the shift λI moves all eigenvalues equally but also leaves eigenvectors untouched. The symmetric core therefore has fewer physical frame parameters than it appears to have.

Starting from the symmetric core with entries (r_e, A, A, r_s, B, r_s), and removing scale and shift, the frame depends on exactly two invariant ratios:

$\rho\odot = (r_s + B - r_e) / A, \sigma_\nu = B / A$.

The first ratio, $\rho\odot$, fixes the solar angle through

$\tan 2\theta_{12} = 2\sqrt{2} / \rho\odot$.

The second ratio, σ_ν, does not affect the symmetric-limit mixing angles directly. It sets the eigenvalue spacing — in particular the distance between the antisymmetric state ν_- and the solar block. That distance controls how strongly the breaking terms feed reactor mixing and atmospheric deviation.

So the kernel has two symmetric-core jobs: **$\rho\odot$ sets the solar angle, σ_ν sets the breaking response.** The previous paper identified $\rho\odot$ as the main target. This paper adds that σ_ν is the hidden second target: invisible in the exact-symmetry angles, but decisive once breaking is turned on.

---

4. The Symmetric Kernel and the Solar Ratio

The proposed weak-commitment solar rule is

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

The intended closure reading: the weak-commitment neutrino residual cloud has a **threefold** completion spread, while the coherent non-electron transport is a **two-channel** μ/τ pair. Using root-normalised amplitudes, a threefold residual gap against a twofold coherent pair gives

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

This is the simplest possible weak-commitment ratio: three residual completion directions read against a paired μ/τ transport channel. Substituting into the solar formula,

$\tan 2\theta_{12} = 2\sqrt{2} / \sqrt{(3/2)} = 2.309\ldots$, $\theta_{12} = \tfrac{1}{2}\cdot\tan^{-1}(2.309\ldots) \approx 33.29°$.

This is the first real gain over the previous draft. The solar angle is no longer a fitted ratio inside the symmetric core; it is a candidate output of the closure kernel, and it lands on the measured value $\sin^2\theta_{12} \approx 0.31$ essentially exactly.

The result is not magic. It follows from the weak-commitment thesis itself: if the stiffness hierarchy collapses, the relevant denominator is not large. A root three-over-two denominator is naturally comparable to the off-diagonal transport, so the solar angle is large but not maximal.

LF-3a — Solar Kernel Rule. The weak-commitment closure kernel fixes the scale/shift-invariant solar ratio to $\rho_\odot = (r_s + B - r_e)/A = \sqrt{(3/2)}$. If true, the solar angle follows without using PMNS as input.

This rule is proposed here, not fully derived. The full closure Hamiltonian must show why the residual count is three and why the coherent transport normalisation is two.

5. The Pair-Block Ratio and the Breaking Denominators

The second symmetric invariant is $\sigma\_\nu = B/A$. In the exact μ–τ limit it does not change $\theta_{23} = 45°$ or $\theta_{13} = 0$, and it does not appear by itself in the solar formula except through the combination $r_s + B - r_e$. Once $\rho_\odot$ is fixed, changing σ_ν can be compensated by changing $r_s - r_e$ without changing θ_{12}.

So why does σ_ν matter? Because the breaking terms couple the antisymmetric state $|\nu_-\rangle = (|\nu\_\mu\rangle - |\nu\_\tau\rangle)/\sqrt{2}$ back into the solar block, and the strength of that response depends on the spectral distance between

$\lambda_- = r_s - B$

and the two solar-block eigenvalues. The proposed pair-kernel rule is

$\sigma\_\nu = B/A = 6/5$.

The intended closure reading: the internal μ/τ pair transport carries a six-to-five enhancement relative to the electron-to-pair transport. Equivalently, the non-electron pair has one extra continuation channel over a five-channel weak-commitment support unit,

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

This is deliberately weaker than the solar rule. The solar rule has a clear root-count interpretation; the 6/5 pair rule is a candidate continuation count that must be checked against the full closure Hamiltonian. Its role, however, is precise: it fixes the breaking denominators. Without it, β and δ can be retuned by moving λ_-. With it, the breaking response becomes constrained.

LF-3b — Pair-Block Kernel Rule. The weak-commitment closure kernel fixes $\sigma_\nu = B/A = 6/5$, setting the spectral denominator through which μ–τ breaking feeds θ_{13}, θ_{23}, and leptonic CP.

Again: candidate, not completed derivation.

---

6. The Leading μ–τ Breaking Relation

The previous neutrino paper used two real breaking parameters: β for the $e\mu/e\tau$ asymmetry, and δ for the $\mu\mu/\tau\tau$ diagonal split. In the (e, +, −) basis the leading couplings into the antisymmetric state are

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

The previous paper correctly warned that both can feed θ_{13}, and that the effect is spectrum-dependent: if ν_- sits close to a solar eigenvalue, even a small δ can strongly affect the reactor angle.

The closure-kernel proposal here is that the leading weak-commitment breaking is **off-diagonal, not diagonal**:

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

The physical reading:

- weak neutrality does not distinguish μ and τ by a first-order self-stiffness;
- the first asymmetry appears in how the electron channel attaches to the paired μ/τ cloud;
- diagonal μ/τ stiffness splitting requires a second-order anchoring effect.

So the leading broken operator keeps the diagonal degenerate ($M_{\mu\mu} = M_{\tau\tau} = r_s$) and carries the asymmetry only in the $e\tau$ entry, $A(1 + \beta\, e^{-i\psi})$, up to $O(\beta^2)$ corrections.

This is a significant reduction. The breaking sector loses one free parameter,

$(\beta, \delta, \psi) \longrightarrow (\beta, \psi)$,

at leading order.

LF-3c — Leading Breaking Relation. The weak-commitment closure kernel forbids a first-order diagonal μ/τ stiffness split: $\delta = 0 + O(\beta^2)$. The leading breaking is the off-diagonal $e\tau$ transport asymmetry β.

This is exactly the kind of relation the previous paper said was needed. It turns the breaking sector from a three-for-three parametrisation into a correlated prediction problem.

---

7. The Weak-Commitment Phase and the Octant–Phase Lock

7.1 The proposed phase

The previous operator allowed a phase ψ in $M_{e\tau} = A(1 + \beta e^{-i\psi})$. A phase is not physical merely because it appears in a matrix; it must survive rephasing and appear in a CP-odd invariant such as the leptonic Jarlskog. The previous paper already checked that ψ behaves correctly: J_{lep} is nonzero for generic ψ, flips sign under CP reversal, and vanishes at $\psi = 0, \pi$. The question now is not whether ψ can create CP — it can — but whether weak-commitment geometry fixes it.

The proposed weak-commitment phase is the bisector

$$\psi_{wc} = (\pi/2 + \pi)/2 = 3\pi/4,$$

read in parallel with the CKM curvature logic, but in the lepton regime:

- **$\pi/2$** is weak-neutral quadrature: the neutrino is not charge-anchored, so it carries a quadrature transport relation to the charged-lepton anchor;
- **π** is complement reversal: moving from one side of the weak doublet to the other reverses the role reading;
- the weak-commitment phase sits at the bisector, $\psi_{wc} = 3\pi/4$.

LF-3d — Weak-Commitment Phase Rule. The leading lepton-frame CP phase is the bisector between weak-neutral quadrature and complement reversal: $\psi_{wc} = 3\pi/4$.

7.2 The octant–phase lock (new exact result)

Exact diagonalisation of the full broken operator reveals a structural fact that the symmetric-limit analysis cannot see: **the atmospheric octant is not a free output.** It is a fixed function of the quadrant of ψ. Reading θ_{23} in the standard PDG/NuFIT convention, $\tan\theta_{23} = |U_{\mu 3}|/|U_{\tau 3}|$, and scanning the four bisector candidates at fixed ($\rho_\odot$, σ_ν, β) gives:

ψ	θ_{23}	θ_{13}	sgn(J_{lep})	octant
$\pi/4$	40.0°	11.2°	−	lower

ψ	θ_{23}	θ_{13}	sgn(J_lep)	octant
$3\pi/4$	48.4°	8.6°	−	upper
$5\pi/4$	48.4°	8.6°	+	upper
$7\pi/4$	40.0°	11.2°	+	lower

Two things follow, and both are exact consequences of the operator, not further proposals.

1. **Octant ↔ CP-sign correlation.** The quadrant of ψ simultaneously fixes the atmospheric octant *and* the sign of leptonic CP violation. The kernel cannot deliver an upper-octant θ_{23} with one CP sign and a lower-octant θ_{23} with the same sign freely; they are tied.
2. **The phase quadrant fixes the octant, β-independently.** Within the physical regime the octant does not depend on the leakage amplitude: $\psi = 3\pi/4$ yields the upper octant at every β, and $\psi = \pi/4$ yields the lower octant at every β. The reactor angle does *not* discriminate between them — the lower-octant family ($\psi = \pi/4$) reproduces the measured $\theta_{13} \approx 8.59°$ just as well, at a slightly smaller leakage amplitude $\beta \approx 0.067$ (where it lands at $\theta_{23} \approx 41.2°$). So θ_{13} is a consistency check at the benchmark β, not the discriminator. **What forces the upper octant is the phase rule LF-3d ($\psi = 3\pi/4$); the octant descends from the phase quadrant, not from the reactor angle.**

LF-3e — Octant–Phase Lock. In the broken weak-commitment operator the atmospheric octant and sgn(J_lep) are locked to the quadrant of ψ, independently of the leakage amplitude β. The phase rule LF-3d ($\psi = 3\pi/4$, or its CP-conjugate $5\pi/4$) therefore predicts $\theta_{23} > 45°$ (upper octant, $\sin^2\theta_{23} \approx 0.559$) together with a definite, nonzero leptonic CP violation. The observed θ_{13} does not select the octant — both octant families can reproduce it at their own β — so it enters only as a consistency check fixing the benchmark β within the already-chosen quadrant.

Security of the octant prediction. Because the octant rides entirely on the phase quadrant, it is exactly as secure as the phase rule LF-3d — no more. The lock (LF-3e) is an exact algebraic fact, but it only converts an *assumed* phase into a *predicted* octant; it adds no independent support. Since LF-3d is at present a heuristic bisector argument ($\pi/2$ quadrature bisected with π complement reversal), not a derivation, the upper-octant prediction inherits that status: it stands or falls with $\psi = 3\pi/4$. The earlier impression that θ_{13} supplies a second, independent leg was mistaken — there is one leg, the phase rule.

A note on convention, since this is exactly where octant errors hide. The atmospheric octant is a physical statement about whether $|U_{\mu3}|$ exceeds $|U_{\tau3}|$ or the reverse, and it must be read with the standard parametrization $U_{\mu3} = s_{23}c_{13}$, $U_{\tau3} = c_{23}c_{13}$, i.e. $\tan\theta_{23} = |U_{\mu3}|/|U_{\tau3}|$. For the benchmark operator $|U_{\mu3}| = 0.739 > |U_{\tau3}| = 0.657$, so $\theta_{23} = 48.4°$ lies in the upper octant. The complement 41.6° is the reciprocal readout $|U_{\tau3}|/|U_{\mu3}|$ and does not correspond to the standard angle. The asymmetric complex coupling sits on the e–τ entry, $A(1 + \beta e^{-i\psi})$; with that placement the operator selects the upper octant, and a lower-octant prediction would require relocating the asymmetry to the e–μ entry and justifying from closure geometry why the electron channel attaches preferentially to μ rather than τ. That is not what the present kernel does.

8. Benchmark PMNS Audit

This section is an audit, not a proof.

Convention. $A = 1$, $r_e = 0$.

Kernel rules. $\rho\odot = \sqrt{(3/2)}$, σ_ν = 6/5, hence

$B = 6/5$, $r_s + B - r_e = \rho\odot \cdot A$, so $r_s = \sqrt{(3/2)} - 6/5 = 0.02475\ldots$

Breaking. $\delta = 0$ (LF-3c), $\psi = 3\pi/4$ (LF-3d).

Leakage amplitude. $\beta\star = \sqrt{3}/20 \approx 0.0866$. This is the least secure rule, included as a minimal closure-leakage candidate: a threefold weak-commitment leakage amplitude distributed over a twentyfold support unit. It must not be treated as established.

The full benchmark operator is the Hermitian matrix

$$\begin{bmatrix} 0 & 1 & 1 + (\sqrt{3}/20)\cdot e^{+3\pi i/4} \\ 1 & 0.02475 & 6/5 \\ 1 + (\sqrt{3}/20)\cdot e^{-3\pi i/4} & 6/5 & 0.02475 \end{bmatrix}$$

Exact Hermitian diagonalisation (eigenvalues $\{-1.1845, -0.8808, +2.1148\}$, with ν_3 identified as the eigenvector of smallest electron content) gives:

Quantity	Kernel benchmark	NuFIT 6.0 (NO)
θ_{12}	**33.40°** ($\sin^2\theta_{12} = 0.303$)	33.68°, $\sin^2\theta_{12} \approx 0.307$
θ_{13}	**8.59°** ($\sin^2\theta_{13} = 0.0223$)	≈ 8.5°, $\sin^2\theta_{13} \approx 0.022$
θ_{23}	**48.37°** ($\sin^2\theta_{23} = 0.559$), upper octant	octant unresolved; upper local min ≈ 0.56
	J_lep	
δ_CP	**≈ 227°**	CP-conserving within 1σ

This is the correct qualitative and near-quantitative PMNS regime: a large solar angle on the measured value, a small reactor angle on the measured value, a definite (upper) atmospheric octant, and nonzero CP violation. Per NuFIT 6.0 the atmospheric octant remains experimentally unresolved (octant degeneracy with $\Delta\chi^2 < 4$, local minima at $\sin^2\theta_{23} \approx 0.56$ and 0.47), so the upper-octant prediction is viable and falsifiable rather than excluded; in normal ordering the global fit without Super-Kamiokande atmospheric data favours the upper octant, while including that data pulls toward the lower. The benchmark $\sin^2\theta_{23} = 0.559$ sits essentially on the upper local minimum (0.56), and θ_{12} and θ_{13} agree with NuFIT 6.0 to better than half a degree.

The important point is the parameter discipline. Once the kernel rules are accepted, the benchmark has **no continuous PMNS-angle fit** — the three angles, the CP sign, and the octant follow from the closure-kernel ratios and the phase rule. The leakage amplitude β⋆ sets the *magnitudes* of θ_{13} and J_lep; it does **not** set the octant or the CP sign, which are fixed β-independently by the phase quadrant. So β is the one continuous knob still in play, and it controls magnitudes only.

The draft must remain honest. The kernel rules have not been derived from the closure Hamiltonian, so the benchmark is a candidate prediction, not a completed VERSF result.

---

9. What This Fixes, and What It Still Does Not Fix

This paper improves the previous neutrino operator in four ways.

1. It turns the solar angle from a fitted ratio into a kernel target, $\rho_\odot = \sqrt{3/2}$, landing on $\sin^2\theta_{12} \approx 0.31$.
2. It identifies the hidden symmetric-core parameter σ_ν = B/A as the control on the breaking response, and proposes σ_ν = 6/5.
3. It reduces the breaking sector by proposing $\delta = 0 + O(\beta^2)$ and $\psi = 3\pi/4$.
4. It establishes the **octant–phase lock** as an exact consequence of the operator, converting the atmospheric octant from a free output into a forced prediction (upper octant, $\sin^2\theta_{23} \approx 0.559$) once the phase rule $\psi = 3\pi/4$ is adopted — β-independently, and without leaning on θ_{13}.

That is the right type of progress. The earlier paper had three breaking parameters for three deviations. This paper replaces two of those freedoms with closure-kernel rules and ties the octant to the phase rule rather than leaving it free.

What remains weakest is the leakage amplitude $\beta\star = \sqrt{3}/20$. The benchmark value works well but is not derived with the clarity of the solar root ratio. This is the main vulnerability of the draft.

A hostile reader should be answered directly here, because the "no continuous angle fit" claim can be oversold. Removing the continuous fit does not by itself amount to a derivation: it moves the fit into the **discrete choice of small integers**. The kernel ratios — $\rho_\odot = \sqrt{3/2}$, B/A = 6/5, and especially $\beta = \sqrt{3}/20$ — are values that reproduce the data, each dressed with a mnemonic reading ("threefold residual cloud over a twofold μ/τ pair," "five-channel support with one extra continuation," "threefold leakage over a twentyfold support unit"). The √3/20 reading is the most transparently reverse-engineered: the integer 20 appears nowhere else in the kernel, so its only present justification is that it produces the observed θ_{13}. Until the closure Hamiltonian forces these ratios, a skeptic is entitled to see **five fitted small-integer ratios with labels, not five derivations** — and that skeptic would be making a fair point. The labels are hypotheses about where the integers come from; they are not yet the source of the integers. The whole programme value of this paper rests on later turning those labels into counts the Hamiltonian is obliged to produce.

There are three possible outcomes.

- **Strong outcome.** The full closure Hamiltonian derives all rules: $\rho\odot = \sqrt{(3/2)}$, B/A = 6/5, $\delta = 0$, $\psi = 3\pi/4$, $\beta = \sqrt{3}/20$. Then PMNS — angles, octant, and CP sign — becomes a genuine VERSF frame prediction.
- **Moderate outcome.** The Hamiltonian derives the solar ratio, the phase, the lock, and $\delta = 0$, but leaves β as a small leakage amplitude. Then the framework predicts the solar angle, the octant, the CP sign, and the correlations among θ_{13}, θ_{23}, and J_lep, but not the absolute magnitude of θ_{13}.
- **Weak outcome.** The kernel rules are not derived. Then this paper is only a good ansatz with mnemonic labels, and should not be presented as a Standard Model derivation.

The honest default is the **moderate outcome** until the closure Hamiltonian is written explicitly. Note that even the moderate outcome retains the octant and CP sign as predictions independent of β — but only because they ride on the phase rule LF-3d. They are therefore exactly as secure as that rule, no more (see §7.2): if LF-3d falls, so does the octant.

---

10. Relation to CKM Curvature and the Standard Model Programme

The CKM and PMNS sides now have parallel structures.

CKM. V_CKM = U_u† U_d. The leading small hierarchy comes from the role-odd relative split Ω_q. The CP and long-triangle correction comes from the common-mode curvature commutator κ = [Ω_0, Ω_q]. The latest CKM paper reduced that correction to a specific C_3-shaped common-mode rotation.

PMNS. U_PMNS = U_e† U_ν. The large frame comes from weak-commitment stiffness collapse, T_ν = D_0 I + ε M_ν. The large–large–small pattern comes from the symmetric weak-commitment kernel; the deviations come from the breaking kernel; and the octant and CP sign come from the phase quadrant.

The flavour programme now has a clean division:

- **CKM:** small split plus curvature residue;
- **PMNS:** collapsed stiffness plus controlled symmetry breaking, with octant and CP sign locked to one phase.

The shared principle is the same,

mixing ~ (off-diagonal transport) / (stiffness gap).

Quarks have large gaps, so the ratio is small. Neutrinos have collapsed gaps, so the ratio can be large. What this paper adds is a proposed origin for the particular large PMNS frame and a forced atmospheric octant.

11. What Is Proved, Conditional, and Open

11.1 Proved or algebraically inherited

- The PMNS frame is the eigenframe of M_ν when $T_\nu = D_0 I + \varepsilon M_\nu$.
- The symmetric core gives $\theta_{23} = 45°$, $\theta_{13} = 0$, and $\tan 2\theta_{12} = 2\sqrt{2} / \rho_\odot$.
- The symmetric core depends on two scale/shift-invariant ratios, $\rho_\odot = (r_s + B - r_e)/A$ and $\sigma_\nu = B/A$; $\rho_\odot$ controls the solar angle and σ_ν controls the breaking response.
- **(New, exact.)** In the broken operator the atmospheric octant and $\mathrm{sgn}(J_{lep})$ are fixed functions of the quadrant of ψ. This is an algebraic property of M_ν, not a proposal.

11.2 New structural content

This paper proposes that weak-commitment closure geometry fixes $\rho_\odot = \sqrt{3/2}$ (giving $\theta_{12} \approx 33.3°$), a pair-block rule $B/A = 6/5$, a leading breaking relation $\delta = 0 + O(\beta^2)$, and a phase $\psi = 3\pi/4$. The exact octant–phase lock then turns that phase into a definite prediction: the **upper** atmospheric octant and a definite nonzero leptonic CP sign, both fixed β-independently by the phase quadrant (the measured θ_{13} is a consistency check, not the selector). These rules reduce the PMNS operator from a flexible ansatz to a constrained kernel candidate with no continuous angle fit.

11.3 Conditional premises

- The weak-commitment residual cloud has a threefold root-normalised structure.
- The coherent non-electron channel is a twofold μ/τ pair.
- The internal μ/τ transport carries a six-to-five continuation enhancement.
- Weak neutrality forbids first-order diagonal μ/τ splitting.
- The CP phase is the bisector between weak-neutral quadrature and complement reversal.
- The benchmark leakage amplitude is $\beta_\star = \sqrt{3}/20$.

11.4 Open problems

1. Derive the three-over-two solar root ratio from the explicit weak-commitment closure Hamiltonian.
2. Derive or reject $B/A = 6/5$.
3. Prove that $\delta = 0$ at first order, rather than assuming it.
4. Derive $\psi = 3\pi/4$ from a physical weak-commitment transport loop, including which CP-conjugate quadrant ($3\pi/4$ vs $5\pi/4$) is selected.
5. Derive the leakage amplitude β, or show that it remains a free small parameter.
6. Embed the resulting frame operator into a full Dirac or Majorana neutrino mass operator.
7. Decide whether the kernel predicts mass ordering or only frame mixing.

Honest status: a closure-kernel candidate, not a completed PMNS derivation. Its value is that it replaces several PMNS-fitting freedoms with explicit kernel rules — and removes the octant freedom outright — so the remaining content can be derived or falsified.

---

12. Falsification Conditions

12.1 Wrong solar ratio. If the closure Hamiltonian does not give $\rho_\odot \approx \sqrt{3/2}$, the solar prediction fails. A significantly larger ratio makes θ_{12} too small; a smaller ratio pushes it toward maximal mixing.

12.2 No paired μ/τ kernel. If the leading neutrino operator does not have approximate $e\mu = e\tau$ and $\mu\mu = \tau\tau$, the entire symmetric-core explanation of maximal atmospheric and zero reactor mixing fails.

12.3 Wrong pair ratio. If B/A is not near the value needed to place the antisymmetric eigenvalue in the correct spectral position, the breaking response generates the wrong θ_{13}–θ_{23} correlation.

12.4 First-order diagonal μ–τ splitting. If the closure kernel gives a first-order δ unrelated to β, the breaking sector regains its lost freedom and the paper falls back to the previous three-parameter ansatz.

12.5 Wrong phase. If weak-commitment transport does not select $\psi = 3\pi/4$ (or its CP-conjugate $5\pi/4$), leptonic CP remains fitted, not derived.

12.6 Wrong octant. *(New, sharp.)* The kernel predicts the **upper** atmospheric octant, $\theta_{23} > 45°$ ($\sin^2\theta_{23} \approx 0.559$), fixed β-independently by the phase quadrant $\psi = 3\pi/4$ (LF-3d). If future long-baseline and atmospheric data (DUNE, Hyper-Kamiokande, JUNO) establish the **lower** octant, the phase rule is falsified: the lower octant corresponds to the $\pi/4$ quadrant, which carries the opposite leptonic-CP sign, so an octant measurement is simultaneously a test of LF-3d and of the predicted CP sign. (A lower octant could still be accommodated only by relocating the asymmetric coupling from the e–τ to the e–μ entry — a different operator needing its own closure-geometry justification.)

12.7 Wrong CP sign. *(New.)* The lock ties sgn(J_lep) to the phase quadrant. If the measured δ_CP lands with the opposite sign to the kernel benchmark (δ_CP ≈ 227°) once the ordering is fixed, the chosen quadrant is wrong, and the phase rule must be revised.

12.8 Free leakage amplitude. If β cannot be derived or bounded, the paper predicts the solar angle, the octant, the CP sign, and correlations, but not the absolute magnitude of θ_{13}. This is not fatal, but it means LF-3 is incomplete.

12.9 No Majorana or Dirac embedding. If the frame operator cannot be embedded into a consistent neutrino mass operator, the PMNS frame result remains incomplete.

12.10 Ledger verdict. If the substrate is ledger-like rather than bath-like at the weak-commitment grain, continuous off-diagonal transport is forbidden, and the entire PMNS frame construction fails.

---

13. Conclusion

The previous weak-commitment neutrino paper moved PMNS from a qualitative idea to an operator target. It showed that small neutrino masses do not imply small mixing, because the weak-commitment scale cancels out of the eigenvectors, and that a μ–τ symmetric core naturally gives the large–large–small pattern (θ_{12} large, $\theta_{23} = 45°$, $\theta_{13} = 0$). But it left the solar ratio underived and the breaking sector over-free.

This paper proposes the next layer: the weak-commitment closure kernel. The central rule is

$\rho_\odot = (r_s + B - r_e)/A = \sqrt{3/2} \Longrightarrow \theta_{12} \approx 33.3°$.

The pair-block rule $B/A = 6/5$ sets the breaking denominators. The leading breaking relation $\delta = 0 + O(\beta^2)$ removes the first diagonal μ/τ split. The phase rule $\psi = 3\pi/4$ turns leptonic CP from a matrix phase into a transport-loop target, and through the exact octant–phase lock it fixes the **upper** atmospheric octant with a definite CP sign — β-independently, with θ_{13} serving as a consistency check rather than the octant selector.

A benchmark leakage value $\beta\star = \sqrt{3}/20$ then gives a PMNS-like frame,

$\theta_{12} \approx 33.4°$, $\theta_{23} \approx 48.4°$, $\theta_{13} \approx 8.6°$, $\delta\_CP \approx 227°$,

with θ_{12} and θ_{13} on the measured values to better than half a degree.

This is not the end of the PMNS problem. The closure Hamiltonian must still derive the kernel rules. But the question is no longer whether a weak-commitment matrix can fit PMNS. It can. The question is now precise:

Does VERSF weak-commitment closure force the kernel $\rho_\odot = \sqrt{3/2}$, $B/A = 6/5$, $\delta = 0$, $\psi = 3\pi/4$, $\beta \approx \sqrt{3}/20$ — and with it the upper atmospheric octant and a fixed leptonic CP sign?

If yes, the PMNS side becomes a genuine Standard Model flavour result with a falsifiable octant. If no, the weak-commitment operator remains a useful ansatz but not a derivation. That is the correct next state of the programme: a sharp kernel, a clear audit, a forced octant, and clean failure conditions.

---

14. References

VERSF flavour architecture

- Keith Taylor, *The Weak-Commitment Neutrino Operator in VERSF.* PMNS mixing from stiffness-gap collapse; μ–τ symmetric core; breaking-sector parameter count; LF-2 target.
- Keith Taylor, *The Electroweak Flavour-Frame Operator in VERSF.* Weak-doublet frame splitting; QF-1/QF-2; weak-commitment frame theorem.
- Keith Taylor, *The CKM Curvature Residue in VERSF.* CKM BCH residue; common-mode curvature; Jarlskog and $|V_{td}|$ repair.
- Keith Taylor, *The Yukawa Operator from Completion-Channel Misalignment in VERSF.* Left Yukawa-square operator; CKM/PMNS as relative sector frames.
- Keith Taylor, *Deriving Flavour Mixing from Closure Geometry.* Generation transport; stiffness gaps; CKM–PMNS contrast.
- Keith Taylor, *Completion Density and Flavour Mixing from Refinement Geometry.* Completion basis; sectoral frame misalignment.

VERSF Standard Model core

- Keith Taylor, *The Standard Model from Hexagonal Geometry.*
- Keith Taylor, *A Unified Derivation of Closure Geometry, Gauge Redundancy, and Mass Structure in the Hexagonal Framework.*

VERSF transport foundations

- Keith Taylor, *Bath or Ledger.*
- Keith Taylor, *The Bath Criterion.*

External comparison data

- I. Esteban, M. C. González-García, M. Maltoni, et al., *NuFit-6.0: Updated global analysis of three-flavor neutrino oscillations*, JHEP 12 (2024) 216. Used only as the external benchmark for θ_{12}, θ_{13}, the unresolved θ_{23} octant, and the δ_{CP} determination.

---

15. Numerical Note

The benchmark values use the scale/shift convention $A = 1$, $r_e = 0$, with

$B/A = 6/5$, $r_s + B - r_e = \sqrt{3/2}\cdot A$, so $r_s = \sqrt{3/2} - 6/5 = 0.024745\ldots$

The broken operator uses $\delta = 0$, $\psi = 3\pi/4$, $\beta = \sqrt{3}/20$.

The PMNS angles are obtained by **exact Hermitian diagonalisation** of M_ν, with the ν_3 column identified as the eigenvector of smallest electron content, followed by the standard readings

$\theta_{13} = \sin^{-1}|U_{e3}|$, $\theta_{12} = \tan^{-1}(|U_{e2}| / |U_{e1}|)$, $\theta_{23} = \tan^{-1}(|U_{\mu3}| / |U_{\tau3}|)$.

Under this convention and the benchmark inputs, the eigenvalues are $\{-1.1845, -0.8808, +2.1148\}$, and the readouts are $\theta_{12} = 33.40°$, $\theta_{23} = 48.37°$, $\theta_{13} = 8.59°$, with $J\_lep = -2.45 \times 10^{-2}$ and $\delta\_CP \approx 227°$. The benchmark magnitudes are $|U\_e3| = 0.149$, $|U\_\mu3| = 0.739$, $|U\_\tau3| = 0.657$; since $|U\_\mu3| > |U\_\tau3|$, θ_{23} lies in the upper octant. The atmospheric octant follows the quadrant of ψ, β-independently within the physical regime: the bisector quadrants $3\pi/4$ and $5\pi/4$ give the upper octant ($\theta_{23} = 48.37°$ at the benchmark β) and the quadrants $\pi/4$ and $7\pi/4$ give the lower octant. The phase quadrant — not θ_{13} — selects the octant; at a fixed quadrant, θ_{13} depends on β (e.g. $\psi = 3\pi/4$ gives $\theta_{13} = 8.59°$ at $\beta = \sqrt{3}/20$ and $\theta_{13} = 11.2°$ at $\beta = 0.12$), and the lower-octant quadrant $\psi = \pi/4$ also reaches $\theta_{13} = 8.59°$ at $\beta \approx 0.067$. The δ_CP determination is invariant under octant reflection, since the Jarlskog factor $s_{23}c_{23} = ½ \sin 2\theta_{23}$ is unchanged by $\theta_{23} \to 90° - \theta_{23}$; $\delta\_CP \approx 227°$ therefore holds independently of the octant reading.

Convention note and correction (v0.2 → v0.4). A v0.2 draft read θ_{23} with the reciprocal ratio $|U\_\tau3|/|U\_\mu3|$ and consequently reported the complement, $\theta_{23} \approx 41.6°$ (lower octant). That readout is non-standard and reflects the angle across 45°. With the standard PDG/NuFIT convention $\tan \theta_{23} = |U\_\mu3|/|U\_\tau3|$ — equivalently $U\_\mu3 = s_{23}c_{13}$, $U\_\tau3 = c_{23}c_{13}$ — the same operator gives $\theta_{23} = 48.37°$, $\sin^2\theta_{23} = 0.559$, the **upper** octant. The operator was never in question; only the readout was. A subsequent v0.3 draft then mis-attributed the octant to the measured θ_{13} ("imposing θ_{13} forces the upper octant"); v0.4 corrects this. The octant is fixed β-independently by the phase quadrant (LF-3d), and θ_{13} does not discriminate the octant — the lower-octant quadrant $\psi = \pi/4$ reproduces $\theta_{13} = 8.59°$ at $\beta \approx 0.067$. The upper-octant prediction therefore inherits exactly the (heuristic) security of the phase rule, no more. The corrected value still matches the original neutrino-operator paper's upper-octant result and NuFIT 6.0's upper local minimum ($\sin^2\theta_{23} \approx 0.56$), and inverts falsifier 12.6: a **lower** octant would falsify the phase rule and the predicted CP sign together.

The benchmark is included only as an audit of the kernel proposal. It is not a substitute for deriving the kernel from the weak-commitment closure Hamiltonian.