The Electroweak Flavour-Frame Operator in VERSF

Part I — Weak-Doublet Frame Splitting and CKM Curvature Part II — Neutrino Weak Commitment and the PMNS Regime

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

Successor to *The Yukawa Operator from Completion-Channel Misalignment in VERSF*, *Deriving Flavour Mixing from Closure Geometry*, *Completion Density and Flavour Mixing from Refinement Geometry*, and the strange current-mass / confinement-operator manuscripts.

Summary for the General Reader

The previous Yukawa-operator paper changed the target of the VERSF flavour programme. Instead of treating quark masses and mixing angles as separate numerical puzzles, it proposed one object — a sectoral Yukawa operator whose eigenvalues are the masses and whose eigenvectors are the way each sector reads the common completion register. CKM mixing is then the mismatch between the up and down sector frames; PMNS mixing is the mismatch between the charged-lepton and neutrino frames.

That paper made real progress but left two doors open. The CKM result depended on a sharp statement about weak doublets — that the up and down frames are equal-and-opposite tilts about one shared closure frame, named **QF-1**. If QF-1 holds, the common frame cancels and CKM is a clean relative exponential; if it fails, extra Baker–Campbell–Hausdorff (BCH) curvature enters and the Cabibbo repair may not survive. The PMNS half was only architecture: neutrinos were called weakly committed and near-degenerate, but that was never turned into a frame theorem.

This paper attempts both tasks in one two-part electroweak frame paper. The theme is single:

The weak interaction reads a fermion doublet through a shared left-handed completion frame, but the two members of the doublet can split from that frame in different regimes.

For quarks the doublet is $Q_L = (u_L, d_L)$. Both members are closure-committed, so the proposal is that weak-isospin complementarity makes the two sector frames opposite readings of one common completion frame, $U_u = U_0 \exp(-\Omega_q/2)$, $U_d = U_0 \exp(+\Omega_q/2)$, giving $V_{CKM} = U_u^\dagger U_d = \exp(\Omega_q)$. This paper writes the weak-doublet operator whose kernel would force that split, and — this is the central sharpening of Part I — shows that the split has **one** controlling quantity: the commutator between the common-mode transport and the role-odd split. That single commutator decides whether the clean exponential holds, sets the size of the leftover

BCH curvature, and is the only object that can repair the two quantities the previous draft missed: the Jarlskog invariant and the long unitarity-triangle side $|V_{td}|$.

For leptons the doublet is $L_L = (\nu_L, e_L)$, and the situation is not symmetric. The charged lepton is anchored and mass-resolved; the neutrino is weakly committed, neutral, and nearly degenerate. So the neutrino frame can rotate far from the charged-lepton frame. CKM and PMNS then stop being unrelated: CKM is the **small-angle, closure-committed** regime of electroweak frame splitting, PMNS the **large-angle, weak-commitment** regime. Part II proves a frame theorem making this precise — large neutrino mixing follows from the *collapse of the stiffness hierarchy*, not from large couplings, and is compatible with arbitrarily small neutrino masses.

The claim is bounded. This paper does not compute the final CKM correction, does not fit PMNS, and does not derive the full Yukawa sector. It states the electroweak frame operator connecting both sides, gives a theorem-level target for QF-1, derives the curvature residue and isolates the single commutator that controls it, and proves the weak-commitment frame theorem the neutrino sector must satisfy.

Table of Contents

Abstract

The VERSF Yukawa-operator draft proposed the left flavour operator $\mathcal{Y}_S = Y_S Y_S^\dagger = U_S \Lambda_S^2 U_S^\dagger$, with diagonal eigenvalues Λ_S from the mass-trace / completion-density programme and sector frames U_S from completion-channel transport, and identified CKM and PMNS as relative sector frames, $V_{CKM} = U_u^\dagger U_d$ and $U_{PMNS} = U_e^\dagger U_\nu$. The quark result was promising — the inherited transport entries 9/40, 81/2000, 243/100000 placed into an anti-Hermitian relative generator Ω_q and exponentiated gave $|V_{us}| \approx 0.2231$, $|V_{cb}| \approx 0.0403$, $|V_{ub}| \approx 0.00393$ at the C_3 holonomy $\varphi = 2\pi/3$ — but it was explicitly conditional on **QF-1**, the

symmetric up/down split $U_u = U_0 \exp(-\Omega_q/2)$, $U_d = U_0 \exp(+\Omega_q/2)$, and it underproduced CP violation and $|V_{td}|$. The PMNS half was only a weak-commitment architecture.

This paper develops the electroweak flavour-frame operator behind both statements.

Part I (quarks). On the doublet space $\mathcal{H}_{EW} = \mathcal{H}_C \otimes \mathcal{H}_W$ we define a weak-role involution and decompose the doublet generator into role-even (common-mode), role-odd (CKM), and role-changing pieces. The clean QF-1 result $V_{CKM} = \exp(\Omega_q)$ requires the doublet frame to be a **product** — a common left factor times the role-odd split — and we show this is strictly stronger than writing the two sector frames as single exponentials of summed generators. The two readings coincide **if and only if** the common-mode and role-odd generators commute; otherwise

$$\log(U_u^\dagger U_d) = \Omega_q - \tfrac{1}{2}[\Omega_0, \Omega_q] + \tfrac{1}{6}[\Omega_0, [\Omega_0, \Omega_q]] + \cdots ,$$

so the entire departure from the clean exponential is governed by the single commutator $\kappa \equiv [\Omega_0, \Omega_q]$. This curvature residue, not a rescaling of the direct 1↔3 entry, is the structurally correct place to lift J and $|V_{td}|$ without disturbing the leading magnitudes (QF-2). All three orders of the residue are verified numerically.

Part II (leptons). For $L_L = (\nu_L, e_L)$ the charged-lepton frame is anchored and near completion-diagonal while the neutrino frame is weakly committed. We prove the **Weak-Commitment Frame Theorem**: if the neutrino transport operator is $T_\nu = D_0 \cdot I + \varepsilon M_\nu$ with ε a small commitment parameter, then for *every* $\varepsilon > 0$ the neutrino eigenframe is exactly the eigenframe of M_ν — independent of ε. Large PMNS angles therefore arise from the loss of the stiffness hierarchy, not from large couplings, and are compatible with arbitrarily small absolute neutrino masses. Approximate e–μ/e–τ symmetry gives large solar mixing; approximate μ–τ symmetry gives near-maximal atmospheric mixing; small breaking gives a small nonzero reactor angle.

The contribution is conditional but specific: it turns the next flavour task into two computable closure-geometry problems — derive the weak-doublet frame-splitting operator and its single controlling commutator κ; derive the weak-commitment neutrino frame and test whether it predicts PMNS without using PMNS angles as inputs.

---

0. Predictive-Content Ledger

EW-FRAME: derive the electroweak left-handed flavour-frame operator whose closure-committed quark regime gives small CKM mixing and whose weak-commitment neutrino regime gives large PMNS mixing.

Object	Construction	Status in this paper
Weak doublets	$Q_L = (u_L, d_L)$, $L_L = (\nu_L, e_L)$	standard EW input, read through VERSF completion frames
Completion space	$\mathscr{H}_C \cong \mathbb{C}^3$	inherited from the Yukawa / refinement programme
Doublet space	$\mathscr{H}_{EW} = \mathscr{H}_C \otimes \mathscr{H}_W$, $\mathscr{H}_W \cong \mathbb{C}^2$	defined here
Left sector frames	U_u, U_d, U_e, U_ν	inherited architecture; electroweak origin sought here
Role decomposition	$\Omega_{EW} = \Omega_{even}\otimes I + \Omega_{odd}\otimes\tau_3 + \Omega_{mix}$	defined; QF-1 ⇔ role-diagonal + product form
Quark symmetric split	$U_u = U_0 e^{(-\Omega_q/2)}$, $U_d = U_0 e^{(+\Omega_q/2)}$	**theorem target QF-1 (load-bearing)**
CKM matrix	$V_{CKM} = U_u^\dagger U_d$	inherited frame-mismatch theorem
CKM leading result	$V_{CKM} = \exp(\Omega_q)$	exact under product-form QF-1
Controlling commutator	$\kappa \equiv [\Omega_0, \Omega_q] = 2[\Omega_{even}, \Omega_{odd}]$	**single object governing QF-1, residue, and $J/\lvert V_{td}\rvert$**
Curvature residue	$\Delta_{BCH} = -\frac{1}{2}\kappa + \frac{1}{6}[\Omega_0, \kappa] + \cdots$	derived; verified to third order
CKM magnitudes (consumed)	0.2231, 0.0403, 0.00393 at $\varphi = 2\pi/3$	$\lvert V_{us}\rvert\approx a$ inherited, $\lvert V_{cb}\rvert\approx b$ consistency check, genuine outputs ($\lvert V_{ub}\rvert$, $\lvert V_{td}\rvert$, ratio, J) 8–41 % off
CP / $\lvert V_{td}\rvert$ repair	via κ (QF-2)	open target
Lepton frame split	$U_{PMNS} = U_e^\dagger U_\nu$	inherited frame-mismatch theorem
Charged-lepton frame	$U_e \approx$ completion-diagonal	structural premise, to be derived from anchoring
Neutrino frame	weak-commitment near-degenerate frame	**theorem target LF-1 / LF-2**
Weak-commitment theorem	$\text{eigenframe}(T_\nu) = \text{eigenframe}(M_\nu)$, $\forall\ \varepsilon>0$	**proved (exact)**
PMNS prediction	large angles from gap collapse	structural theorem target, not numerical fit

Three layers, as in the mass-trace series:

1. **Frame algebra (exact).** Given sector frames, CKM and PMNS are relative frames; given product-form QF-1, CKM = $\exp(\Omega_q)$; given $T_\nu = D_0 I + \varepsilon M_\nu$, the neutrino eigenframe is M_ν's.
2. **Operator-construction claim (proposed).** Weak doublets produce either a symmetric role split (quarks) or a weak-commitment degenerate frame (neutrinos).
3. **Physical derivation (owed).** Closure geometry must derive the role-splitting operator, the leading Ω_q entries, the commutator κ, and the neutrino frame eigenmodes.

The discipline is unchanged: **a matrix with right-looking entries is not a derived operator, and reproducing an entry that was fed in is not predicting it.**

---

1. The Inherited Problem from the Yukawa-Operator Paper

The previous draft established the matrix architecture $\mathcal{Y}_S = U_S \Lambda_S^2 U_S^\dagger$ with the two charged-current mixing matrices as relative frames, $V_{CKM} = U_u^\dagger U_d$ and $U_{PMNS} = U_e^\dagger U_\nu$. Its central advance was the CKM relative exponential: instead of diagonalising one Hermitian generation-transport Hamiltonian, it placed the inherited hierarchy into an anti-Hermitian relative generator and computed $V_q = \exp(\Omega_q)$. For

$a = 9/40$, $b = 81/2000$, $c = 243/100000$, $\varphi = 2\pi/3$,

it returned $|V_{us}| \approx 0.2231$, $|V_{cb}| \approx 0.0403$, $|V_{ub}| \approx 0.00393$ with exact unitarity — but conditional on QF-1, $U_u = U_0 \exp(-\Omega_q/2)$, $U_d = U_0 \exp(+\Omega_q/2)$, which it motivated from weak-isospin complementarity without deriving. It left two residues: the Jarlskog invariant J about 40 % low, and $|V_{td}|$ underproduced. The PMNS half was open — neutrinos were called weakly committed and near-degenerate, but U_ν was never derived.

This paper asks whether the weak doublet itself supplies the missing frame-splitting operator, in both the quark and lepton regimes.

---

2. The Electroweak Flavour-Frame Viewpoint

The electroweak theory organises left-handed fermions into doublets, $Q_L = (u_L, d_L)$ and $L_L = (\nu_L, e_L)$. The weak charged current couples the two members before the mass matrices are diagonalised; after diagonalisation the relative left frames become CKM or PMNS. VERSF reads this as a statement about completion geometry: the weak current sees a **doublet completion frame**, while mass readout sees **sectoral completion frames**.

The CKM/PMNS contrast then has a structural cause rather than a coincidence:

- in quark doublets both members are closure-committed and confinement-filtered, so the frame split is small;
- in lepton doublets the charged lepton is anchored but the neutrino is weakly committed and nearly degenerate, so the relative frame can be large.

The single electroweak problem is therefore: *given a left-handed weak doublet, derive the two sectoral completion frames it induces.* The quark answer should be a small antisymmetric split about a common frame; the lepton answer should be an anchored charged-lepton frame and a weak-commitment neutrino frame.

3. Weak-Doublet Completion Space and Role Operators

Let $\mathscr{H}_C \cong \mathbb{C}^3$ be the generation completion space and $\mathscr{H}_W \cong \mathbb{C}^2$ the weak-role space of a left-handed doublet. The electroweak doublet frame lives on

$$\mathscr{H}_{EW} = \mathscr{H}_C \otimes \mathscr{H}_W.$$

Choose role basis vectors $|+\rangle$, $|-\rangle$. For quarks read $|+\rangle$ as up-type and $|-\rangle$ as down-type; for leptons $|+\rangle$ as neutrino and $|-\rangle$ as charged-lepton. The role operator is $\tau_3|\pm\rangle = \pm|\pm\rangle$, with the role-flip operators τ_1, τ_2 exchanging the two members.

A general anti-Hermitian doublet generator decomposes against the role algebra as

$$\Omega_{EW} = \Omega_{even} \otimes I_W + \Omega_{odd} \otimes \tau_3 + \Omega_{mix}, \quad \Omega_{mix} = \Omega_{(1)} \otimes \tau_1 + \Omega_{(2)} \otimes \tau_2.$$

The three pieces have distinct physical roles:

- **Ω_{even}** — common-mode completion transport, seen identically by both members;
- **Ω_{odd}** — the antisymmetric (role-odd) frame split between the members;
- **Ω_{mix}** — role-changing transport that mixes the up-type and down-type readouts.

For charged-current mixing we want the sector frames induced after projecting onto the two role components. The basic electroweak frame principle is then:

Common-mode completion transport should cancel from charged-current mixing; role-odd completion transport should survive as CKM or PMNS; role-changing transport must be gapped from the mass-readout frame for the role basis to be a clean readout basis.

The quark sector is the clean test. Note already that this is *two* conditions, not one: role-diagonality ($\Omega_{mix} \approx 0$) makes the role basis clean, and — as §6 shows — a further factorisation condition is needed for the common piece to cancel exactly.

Part I — Quark Weak-Doublet Frame Splitting

4. Quark Frame Splitting and the Product Form

Quark doublets are closure-committed, colour-confined, and mass-resolved, so their frame splitting should be small and controlled by the inherited generation-transport hierarchy. The proposed leading doublet operator is the **product form**

$$U_Q = (U_0 \otimes I_W) \cdot \exp\left(\tfrac{1}{2} \Omega_q \otimes \tau_3 \right),$$

with U_0 acting on $\mathscr{H}$_C and common to both roles. Because τ_3 = diag(+1, −1) is role-diagonal, $\exp(\tfrac{1}{2}\,\Omega\_q \otimes \tau_3)$ is block-diagonal with blocks $\exp(\pm\tfrac{1}{2}\,\Omega\_q)$, and $U_0 \otimes I\_W$ multiplies both blocks identically. Projecting onto the two roles gives

$U\_d = U_0 \exp(+\Omega\_q/2)$, $U\_u = U_0 \exp(-\Omega\_q/2)$,

(up to the convention assigning which role sign is up). This is QF-1 in operator form: the up and down frames are not independent, but two opposite readings of one weak-doublet completion frame. The word "leading" should be read with §6 in mind: this product form is not the generic dynamical output but the special $\kappa = 0$ slice of the natural additive split a single weak-doublet Hamiltonian would generate — a point §6 makes precise rather than assumes.

4.1 Physical reading

The weak doublet carries one shared left-handed closure object. The up/down distinction is not two unrelated completion registers but a role record inside one weak-isospin pair. Common frame curvature belongs to the doublet as a whole and should cancel from the charged current; antisymmetric role curvature belongs to the *difference* of the two readings and becomes CKM. The quark doublet is thus a test of complementarity:

shared doublet completion + opposite role readouts ⇒ small relative CKM frame.

The proof target is not merely that U_u and U_d are related; it is that the relation is the *exact symmetric product split* needed for common-mode cancellation. As §6 makes precise, that exactness is a genuine and separable assumption.

---

5. The Symmetric-Split Theorem Target (QF-1)

Theorem target QF-1 (Weak-Doublet Symmetric Split). Let Q_L = (u_L, d_L) be a closure-committed weak doublet. Suppose:

1. the doublet has a shared completion frame U_0 on $\mathscr{H}$_C;
2. the up/down role record is a role-odd completion generator $\Omega\_q \otimes \tau_3$;
3. role-changing components Ω_mix are gapped from the mass-readout frame;
4. the common transport enters as a **left factor** $U_0 \otimes I\_W$, not as an additive generator term (equivalently, $[\Omega\_even, \Omega\_q] = 0$; see §6).

Then the induced sector frames have the product form $U\_u = U_0 \exp(-\Omega\_q/2)$, $U\_d = U_0 \exp(+\Omega\_q/2)$, and the charged-current relative frame is $V\_CKM = U\_u^\dagger U\_d = \exp(\Omega\_q)$.

Proof, conditional on the premises. By (2)–(3) the doublet generator is role-diagonal, so the role basis block-diagonalises it; by (4) the common piece factors as $U_0 \otimes I\_W$. The doublet operator is then $U\_Q = (U_0 \otimes I\_W)\cdot\exp(\tfrac{1}{2}\,\Omega\_q \otimes \tau_3)$. Projecting onto the τ_3 eigenvectors gives sector factors $\exp(\mp\Omega\_q/2)$, each left-multiplied by the same U_0. Hence

$U\_u^\dagger U\_d = \exp(+\Omega\_q/2)\cdot U_0^\dagger U_0\cdot\exp(+\Omega\_q/2) = \exp(+\Omega\_q/2)\cdot\exp(+\Omega\_q/2) = \exp(\Omega\_q)$,

using anti-Hermiticity $\exp(-\Omega\_q/2)^\dagger = \exp(+\Omega\_q/2)$. ■

This is not yet first-principles: premises (1)–(4) must come from the explicit weak-doublet closure Hamiltonian. But the theorem isolates exactly what must be proved. The previous draft folded all four into "symmetric split"; the sharpening here is that **(4) is the genuinely strong premise** — the others fix the *basis*, but (4) fixes how the common transport *composes*, and §6 shows the natural single-Hamiltonian reading violates it.

---

6. Common-Mode Cancellation and the Product-vs-Additive Distinction

The product form is stronger than the additive notation

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

In general $\exp(\Omega_0 - \Omega\_q/2) \neq \exp(\Omega_0)\exp(-\Omega\_q/2)$: the equality holds only when $[\Omega_0, \Omega\_q] = 0$. So the clean CKM cancellation is exact for the **product** split but not, in general, for the **additive** one.

This distinction is not cosmetic — it is where the natural physics and the clean result diverge. A single weak-doublet Hamiltonian generates the doublet frame as **one exponential of a summed generator**:

$U\_Q = \exp(\Omega\_even \otimes I\_W + \Omega\_odd \otimes \tau_3)$, $\Omega\_odd = \Omega\_q/2$.

Because both terms are role-diagonal, this is block-diagonal with blocks $\exp(\Omega\_even \pm \Omega\_odd)$ — i.e. it produces the **additive** sector frames $U\_u = \exp(\Omega_0 - \Omega\_q/2)$, $U\_d = \exp(\Omega_0 + \Omega\_q/2)$, *not* the product frames. (Verified directly: projecting $\exp(\Omega\_even\otimes I + \Omega\_odd\otimes\tau_3)$ reproduces $\exp(\Omega_0 \pm \Omega\_q/2)$ to machine precision.) The product form of §4 is a *different* object, $(U_0\otimes I)\cdot\exp(\Omega\_odd\otimes\tau_3)$, and the two agree only when the common-mode and role-odd generators commute.

Theorem 1 (Common-Mode Cancellation). If $U\_u = U_0 A^\dagger$ and $U\_d = U_0 A$ with $A = \exp(\Omega\_q/2)$, then $V\_CKM = A^2 = \exp(\Omega\_q)$, independent of U_0.

Proof. $V\_CKM = (U_0A^\dagger)^\dagger(U_0A) = A\, U_0^\dagger U_0\, A = A^2 = \exp(\Omega\_q)$. ■

Theorem 1 is algebraic and exact; its content is the factorisation, *U_0 must be a genuine left-common factor*. Any effect acting identically on both doublet members is invisible to CKM — but only if it composes multiplicatively, not additively.

The controlling commutator. Write $\kappa \equiv [\Omega_0, \Omega_q] = 2[\Omega_{even}, \Omega_{odd}]$. Then product-form QF-1 holds exactly ⟺ $\kappa = 0$, and to leading order the departure of V_CKM from $\exp(\Omega_q)$ is set by κ (§7). More precisely: the residue is $-\tfrac{1}{2}\kappa + \tfrac{1}{6}[\Omega_0, \kappa] + \cdots$, so κ alone controls it only when $\|\Omega_0\|$ is small enough that the cubic and higher terms — which require Ω_0 itself, not just κ — are negligible. §8 shows that the window where the J / |V_td| repair actually lives is exactly such a small-$\|\Omega_0\|$ regime, so there the single object κ does govern the residue; for a larger common generator the full pair (Ω_0, Ω_q) is needed. With that smallness understood, κ — the failure of the common closure frame to commute with the role-odd split — is the spine of Part I: it decides whether the Cabibbo repair is exact, it *is* the leading curvature residue, and it is the lever for the J / |V_td| repair (§8).

So QF-1's clean exponential is conditional in a precise and computable way: it is the $\kappa = 0$ slice of the weak-doublet frame algebra.

---

7. CKM Curvature Residue and the BCH Correction

If the sector frames are the additive exponentials of noncommuting absolute generators, a curvature residue appears. With Ω_0, Ω_q anti-Hermitian and

$$U_u = \exp(\Omega_0 - \Omega_q/2), \quad U_d = \exp(\Omega_0 + \Omega_q/2),$$

write $A = -\Omega_0 + \Omega_q/2$ (so $U_u^\dagger = \exp A$) and $B = \Omega_0 + \Omega_q/2$ (so $U_d = \exp B$). Then $A + B = \Omega_q$ and $[A, B] = -[\Omega_0, \Omega_q] = -\kappa$, and the BCH expansion

$$\log(e^A e^B) = A + B + \tfrac{1}{2}[A,B] + 1/12([A,[A,B]] + [B,[B,A]]) + \cdots$$

gives

$$\log(U_u^\dagger U_d) = \Omega_q - \tfrac{1}{2}\kappa + \tfrac{1}{6}[\Omega_0, \kappa] + \cdots, \quad \kappa = [\Omega_0, \Omega_q].$$

Defining the curvature residue $\Delta_{BCH} \equiv \log(U_u^\dagger U_d) - \Omega_q$,

$$\boldsymbol{\Delta_{BCH} = -\tfrac{1}{2}[\Omega_0, \Omega_q] + \tfrac{1}{6}[\Omega_0, [\Omega_0, \Omega_q]] + \cdots,}$$

which vanishes iff $\kappa = 0$ (product-form QF-1). Both displayed coefficients are confirmed numerically: in the perturbative regime the residual after subtracting $-\tfrac{1}{2}\kappa$ is $O(t^3)$, and after also subtracting $\tfrac{1}{6}[\Omega_0,\kappa]$ it is $O(t^4)$, where t scales the generators.

7.1 Why this residue is the right correction

Adjusting the direct 1↔3 generator entry c or the holonomy φ disturbs |V_ub| (next section). The curvature residue is structurally different: it enters distinct matrix elements with signs and phases fixed by commutators, not by a uniform rescaling. The leading exponential already places |V_us|, |V_cb|, |V_ub| well; the deficits are J and |V_td|. A commutator correction can in principle lengthen the unitarity triangle and raise CP violation *without* simply inflating |V_ub|, precisely because it is not a rescaling of any single entry. This is the calculation the paper hands forward, not one it closes.

---

8. Can the Residue Repair J and |V_td|?

First, the discipline carried from the mass-trace series and the Yukawa draft: **the CKM magnitudes this construction reproduces are not all predictions.** Reading the consumed numbers honestly:

quantity	class	value ($\varphi = 2\pi/3$)	measured	deviation
|V_us|	inherited (≈ a)	0.2231	0.2243	−0.5 % (was +0.3 % at a = 0.225)
|V_cb|	consistency (≈ b)	0.0403	0.0408	−1.2 %
|V_ub|	**output**	0.00393	0.00382 excl / ~0.0043 incl	+2.9 % / −8.6 %
|V_td|	**output**	0.0061	0.0086	**−29 %**
|V_ub|/|V_cb|	**output**	0.098	0.083 ± 0.004	**+18 % (3.6σ)**
J	**output**	1.84×10^{-5}	3.12×10^{-5}	**−41 %**

The construction is most accurate where it inherits (|V_us| ≈ a, with the exponential nudging it *away* from data) and least accurate where it genuinely predicts. Two of the four genuine outputs are also branch-dependent: against the inclusive |V_ub| ≈ 0.0043 the sign of the |V_ub| deviation flips, and the $\varphi = 2\pi/3$ / |V_ub| "coincidence" is exclusive-branch only — on the inclusive branch the preferred phase drifts toward $\pi/2$, which itself raises J toward 2.1×10^{-5}. The cleanest single output, |V_ub|/|V_cb| = 0.098 against the tight 0.083 ± 0.004, is a 3.6σ statement that the construction underproduces the 1↔3 sector relative to the 2↔3 sector.

This is exactly why the curvature residue, not a knob, is the right target. The Yukawa draft showed the knobs are coupled: raising c lifts J but drives |V_ub|/|V_cb| from ≈ 0.10 toward ≈ 0.14; moving φ raises J near $\pi/2$ but overshoots |V_ub|. The genuine deficits (J, |V_td|) are precisely the entries the leading exponential cannot place, and they must be supplied by κ without disturbing the inherited entries.

CKM curvature target QF-2. Compute $\kappa = [\Omega_0, \Omega_q]$ from the weak-doublet closure operator and test whether

$\exp(\Omega_q + \Delta_{BCH})$, $\Delta_{BCH} = -\frac{1}{2}\kappa + \frac{1}{6}[\Omega_0, \kappa] + \cdots$, raises J and |V_td| toward measurement while holding |V_us| ≈ 0.224, |V_cb| ≈ 0.041, and |V_ub| within its inclusive/exclusive band. The four constraints are severe: Δ_BCH cannot be a generic correction, only a weak-doublet curvature commutator. If QF-2 succeeds, the CKM mechanism is much stronger — leading hierarchy from Ω_q, triangle and CP geometry from κ. If it fails, the relative exponential remains a good leading approximation but not a full CKM derivation.

Because Part I has shown κ is a *single* object (the common/role-odd commutator), QF-2 is not an open-ended search over corrections — it is the computation of one commutator and a four-constraint test. That is the sharpest form the CKM completion problem has yet taken.

The repair window (a result, not a hope). A direct five-target solve confirms a clean repair already exists at the C_3 value φ = 2π/3, with no need to move the holonomy. A small common generator, $\|\Omega_0\| \approx 0.035$ (hence ‖κ‖ ≈ 0.007), brings J to $\approx 3.1\times10^{-5}$ and |V_td| to ≈ 0.0083 while |V_us|, |V_cb|, |V_ub| stay within ≈ 1σ of measurement. Two facts make this sharp:

- *In this window −½κ carries the entire correction.* At the repair point ‖−½κ‖ equals the full residue ‖Δ_BCH‖ to four figures, while the cubic term $\|\tfrac{1}{6}[\Omega_0, \kappa]\| \approx 5\times10^{-5}$ is about a hundred times smaller. So QF-2 collapses, in the regime the repair lives in, to *compute κ and take −½κ* — true to ≈ 1 %. This collapse is contingent on the smallness of $\|\Omega_0\|$ (§6); a larger common generator would require the full pair (Ω_0, Ω_q), not κ alone.
- *The repair is viable because the sensitivities are asymmetric.* A ≈ 3 % perturbation (‖κ‖ ≈ 0.007 against ‖Ω_q‖ ≈ 0.32) moves J by a factor of roughly three and |V_td| by ≈ 25 %, while |V_us| stays rigid below 0.5 %. This is the mechanism that lets QF-1 (wants κ ≈ 0) and QF-2 (wants κ ≠ 0) live on the same κ: the Cabibbo-scale magnitudes the leading exponential already saturates have little slack and barely move, whereas J and |V_td| are exactly the entries it leaves underdetermined, so they have maximal leverage per unit κ. |V_cb| is slightly less rigid than |V_us| but not a binding constraint: a generic κ at the repair norm shifts it by only a few percent (median ≈ 2 %, tail rarely past ≈ 6 %), against the solution's own ≈ 0.3 %, and the direction that *maximally* lifts J already disturbs |V_cb| by under 1 %. So a J-raising, |V_cb|-sparing κ is easy to find, not narrowly threaded — the repair needs no fine alignment, which is a friendlier fact for the construction than a large cancellation would be. The genuine stringency of the four-constraint test is in lifting J and |V_td| *together* while holding |V_us| and |V_ub|, not in any competition from |V_cb|.

The holonomy is held at φ = 2π/3 throughout this repair; the "derive φ" obligation (§15.4) is separate and is not doing the work here.

Part II — Lepton Weak-Doublet Frame Splitting

9. The Asymmetric Lepton Doublet

The lepton doublet $L_L = (\nu_L, e_L)$ should not be forced into the quark pattern. Its members do not share confinement or anchoring: the charged lepton is mass-resolved and anchored, the neutrino is neutral, weakly anchored, and weakly committed. The lepton analogue of QF-1 is therefore **not** a symmetric split $U_\nu = U_0 \exp(-\Omega_l/2)$, $U_e = U_0 \exp(+\Omega_l/2)$ — that would give CKM-like small mixing, contrary to observation. The proposed regime is instead

$$U_e \approx U_0^l,\ U_\nu = U_0^l\, U_{wc},\ \Rightarrow U_{PMNS} = U_e^\dagger\, U_\nu \approx U_{wc},$$

with U_{wc} a large weak-commitment neutrino frame. PMNS is thus the relative frame between an anchored charged-lepton basis and a weak-commitment neutrino transport basis — the same electroweak frame principle (relative frames produce mixing), in a different regime because the two doublet members have radically different commitment status.

---

10. Weak Commitment and the Collapse of Neutrino Stiffness Gaps

In the CKM construction the generation stiffness hierarchy is $D_q = \mathrm{diag}(1, 2, 4)$, with order-unity gaps; off-diagonal transport is small relative to those gaps, so mixing angles are small. The VERSF weak-commitment claim is that the mechanism producing this hierarchy is largely absent for neutrinos — a neutrino does not create the local closure-commitment events a confined quark does — so the gaps **collapse**:

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

with ε a small commitment parameter. Consistency requires the off-diagonal neutrino transport to co-scale with the same parameter, $\Pi_\nu = \varepsilon K$, so the weak-commitment limit is finite rather than singular. The transport operator is then

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

The term $D_0 \cdot I$ shifts all eigenvalues equally and the overall ε is a positive scalar — *neither* affects eigenvectors. So the neutrino frame is set by the $O(1)$ operator $\Delta + K$, not by any hierarchy of gaps. That is the structural reason PMNS mixing can be large while neutrino masses are small.

---

11. The Weak-Commitment Frame Theorem

Theorem 2 (Weak-Commitment Frame). Let the neutrino transport operator be

$$T_\nu(\varepsilon) = D_0 \cdot I + \varepsilon\, M_\nu,\ M_\nu \text{ Hermitian (or complex-symmetric) on } \mathscr{H}_C,\ \varepsilon > 0.$$

Then for **every** $\varepsilon > 0$ the eigenvectors of T_ν are exactly the eigenvectors of M_ν. Consequently the neutrino frame U_ν is independent of ε, and approaches a finite nontrivial unitary even as the absolute neutrino mass scale $\varepsilon \to 0$.

Proof. $D_0 \cdot I$ is a scalar shift and commutes with everything, so it leaves eigenvectors unchanged; multiplying M_ν by the positive scalar ε rescales eigenvalues but fixes eigenvectors. Hence diagonalising T_ν is diagonalising M_ν, exactly, for any $\varepsilon > 0$. ■

This is stronger than a leading-order statement: when the transport operator has the exact co-scaling form $T_\nu = D_0 \cdot I + \varepsilon M_\nu$, the eigenframe equals M_ν's eigenframe identically, not merely as $\varepsilon \to 0$. (Numerically confirmed: the eigenframes of T_ν and M_ν coincide to machine precision across ε from 10^{-1} to 10^{-6}, with only conditioning-level drift.) If the form is only $T_\nu = D_0 \cdot I + \varepsilon M_\nu + O(\varepsilon^2)$, the equality holds to leading order in ε, with corrections suppressed by the residual $O(\varepsilon)$ splittings. Either way, **small neutrino masses do not imply small PMNS angles**: the eigenvectors are set by the *shape* of M_ν, not the *scale* of the masses.

11.1 CKM–PMNS contrast proposition

Proposition. CKM is small because closure-committed quark transport has non-degenerate stiffness gaps and filtered off-diagonal transport; PMNS can be large because weak-commitment neutrino transport has near-degenerate stiffness and off-diagonal transport of comparable order to the residual gaps.

In symbols, the quark mixing angle scales as $\theta_{ij}^{CKM} \approx \Pi_{ij}^q / (D_j^q - D_i^q)$ with $O(1)$ gaps and filtered Π^q, hence small; while the neutrino angle scales as

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

in which ε cancels, leaving an $O(1)$ angle. The contrast is therefore not two unrelated facts but one principle in two regimes — the **same** ratio "off-diagonal transport over stiffness gap," evaluated where the gap is $O(1)$ (quarks) versus where the gap has collapsed (neutrinos). This is a structural theorem target, not a PMNS fit.

---

12. A Minimal PMNS Operator Target

The frame theorem needs a concrete M_ν. The earlier transport work suggests a minimal μ–τ symmetric form,

$$M_\nu = \begin{bmatrix} r_e & A & A \\ A & r_s & B \\ A & B & r_s \end{bmatrix}$$

in which the e–μ and e–τ entries are approximately equal, the μ–τ block is nearly symmetric, the reactor angle vanishes in the exact μ–τ limit, a small breaking parameter generates θ_{13}, and phases in the weak-commitment transport loop generate leptonic CP violation. A more explicit target separates the breakings:

$M\_\nu = \lceil 0 \; A \; A(1+\beta) \rceil | \; A \; 0 \; B \; | \lfloor A(1+\beta) \; B \; \delta \rfloor$

with small β (e–τ asymmetry) and δ (μ–τ breaking). This is a target, not a derived matrix: the closure-geometry calculation must supply A, B, β, δ, and the CP phase from weak-commitment transport, not from PMNS data.

PMNS theorem target LF-2. Derive a weak-commitment neutrino frame U_ν whose near-degenerate transport operator produces large θ_{12}, large θ_{23}, small nonzero θ_{13}, and a physical CP phase, *without* fitting PMNS angles.

This is deliberately less numerical than QF-2: the present paper establishes the electroweak frame operator and the regime, not the completed neutrino sector.

---

13. Majorana Caveat and Neutrino Phases

If neutrinos are Dirac, the left-square framework is parallel to quarks and charged leptons, $\mathcal{Y}\_\nu = Y\_\nu Y\_\nu^\dagger = U\_\nu \Lambda\_\nu^2 U\_\nu^\dagger$. If neutrinos are Majorana, the mass matrix is complex symmetric and is diagonalised by a Takagi (congruence) factorisation,

$M\_\nu = U\_\nu \Lambda\_\nu U\_\nu^T$, U_ν unitary, $\Lambda\_\nu \geq 0$,

not by a Dirac similarity transform. The charged-current PMNS matrix still depends on the left neutrino frame U_ν exactly as above, so the frame-mismatch architecture and Theorem 2 are unaffected at the level of mixing angles. But the two Majorana phases are physical and cannot be recovered from the left-square operator alone — they require the symmetric mass operator itself. This paper therefore makes no claim about Majorana phases or neutrinoless double-beta decay; those belong to the neutrino-mass operator paper that follows LF-2.

---

14. Relation to Bath/Ledger and Continuous Mixing

Both halves assume continuous off-diagonal frame transport — bath-like, not ledger-like, pre-commitment completion weight. Under a **ledger** reading each completion alternative keeps a separately conserved account; admissible transformations are phases and relabellings, giving diagonal Yukawa operators up to permutations and no CKM or PMNS mixing. Under a **bath** reading only the total weight is conserved, weight moves continuously between alternatives, and off-diagonal generators (Ω_q, M_ν) are admissible. The electroweak frame programme depends on the same bath-side choice as the flavour-transport and non-abelian gauge programmes:

bath transport ⇒ nontrivial sector frames ⇒ CKM and PMNS.

This paper consumes the bath as the transport condition needed for any Yukawa-frame theory with continuous mixing; it does not prove it.

---

15. What Is Proved, Conditional, and Open

15.1 Proved here, given definitions.

- CKM and PMNS are relative sector frames once the left frames are defined.
- Product-form QF-1 ⇒ exact common-mode cancellation and $V_{CKM} = \exp(\Omega_q)$ (Theorem 1).
- The natural single-exponential doublet generator gives the *additive* sector frames, and product = additive **iff** $\kappa = [\Omega_0, \Omega_q] = 0$; otherwise $\Delta_{BCH} = -\frac{1}{2}\kappa + \frac{1}{6}[\Omega_0, \kappa] + \cdots$ (coefficients verified to third order).
- For $T_\nu = D_0 \cdot I + \varepsilon M_\nu$ the neutrino eigenframe equals M_ν's eigenframe for every $\varepsilon > 0$ (Theorem 2), so large mixing is compatible with arbitrarily small neutrino masses.

15.2 New structural content. Part I is unified by a single object: the commutator $\kappa = [\Omega_0, \Omega_q]$ simultaneously (i) measures the failure of product-form QF-1, (ii) is the curvature residue Δ_{BCH}, and (iii) is the only admissible lever for the J / $|V_{td}|$ repair. Part II reduces the CKM/PMNS dichotomy to one ratio (transport / gap) evaluated in two regimes. CKM and PMNS become two regimes of one electroweak frame problem rather than separate mechanisms.

15.3 Conditional premises.

- The weak doublet has a shared completion frame at the left-handed level.
- (**QF-1, load-bearing**) the quark split is role-diagonal *and* product-form ($\kappa = 0$) at leading order; §§4–6 motivate this from weak-isospin complementarity but do not derive it, and §6 shows the natural single-Hamiltonian reading instead gives the additive form.
- The leading Ω_q entries are the inherited closure-transport values 9/40, 81/2000, 243/100000.
- The weak-commitment neutrino operator has the co-scaling form $D_\nu = D_0 \cdot I + \varepsilon \Delta$, $\Pi_\nu = \varepsilon K$.
- The substrate supports bath-type continuous mixing.

15.4 Open problems.

- Derive product-form QF-1 ($\kappa = 0$ at leading order) from the weak-doublet closure Hamiltonian.
- Compute κ and Δ_{BCH} from closure geometry and test the four-constraint QF-2 repair of J and $|V_{td}|$.
- Derive the C_3 holonomy $\varphi = 2\pi/3$ rather than selecting it (and note its $|V_{ub}|$ attractiveness is exclusive-branch only).
- Derive the minimal M_ν entries and phases from weak-commitment geometry (LF-2).
- Extend the left-frame analysis to full Dirac / Majorana mass operators.

- Integrate the diagonal eigenvalue programme, including charm and bottom.

Honest status: a candidate electroweak flavour-frame operator with exact algebraic consequences and sharply named proof targets — not a completed Standard Model flavour derivation. Its value is that the next two calculations are now single, well-posed objects: one commutator (κ, for CKM) and one near-degenerate operator (M_ν, for PMNS).

---

16. Falsification Conditions

16.1 QF-1 failure (product form). If the quark weak doublet does not produce a product-form split at leading order — i.e. if $\kappa = [\Omega_0, \Omega_q]$ is not parametrically small — the clean CKM relative exponential is not derived, and V_{CKM} carries an uncontrolled Δ_{BCH}.

16.2 Common-mode leakage. If common-mode transport does not cancel from V_{CKM} (factorisation fails), the Cabibbo repair is unstable and the previous $|V_{us}|$ over-suppression can return.

16.3 Wrong curvature residue. If the computed κ is too large it spoils $|V_{us}|$ and $|V_{cb}|$; if too small or of the wrong phase structure it cannot repair J or $|V_{td}|$. QF-2 is the operational test, and it is two-sided.

16.4 Weak-commitment failure. If neutrino stiffness gaps do not collapse, or if off-diagonal neutrino transport is suppressed relative to the gaps (i.e. $K/\Delta \ll 1$), PMNS becomes CKM-like and the large-mixing explanation fails.

16.5 PMNS fitting failure. If M_ν must be chosen to match PMNS angles rather than derived from weak-commitment geometry, the PMNS half is Standard-Model-style fitting (LF-2 unmet).

16.6 Ledger verdict. If the substrate is ledger-like at the relevant pre-commitment transport grain, continuous off-diagonal electroweak frame mixing is forbidden and both CKM and PMNS frame constructions fail at the level of principle.

---

17. Conclusion

The previous Yukawa-operator paper showed that VERSF flavour is a matrix problem, not a collection of isolated mass ratios, and identified CKM and PMNS as relative sector frames. This paper asks what the weak doublet itself contributes to those frames, and answers with a two-regime electroweak frame operator.

In the quark regime both members of $Q_L = (u_L, d_L)$ are closure-committed, so the doublet should split symmetrically about a common completion frame, $U_u = U_0 \exp(-\Omega_q/2)$, $U_d = U_0$

exp(+Ω_q/2), giving V_CKM = exp(Ω_q). The central result of Part I is that this clean exponential is the $\kappa = 0$ slice of the weak-doublet frame algebra, where $\kappa = [\Omega_0, \Omega_q]$ is the commutator of the common-mode transport with the role-odd split. The same κ is the curvature residue Δ_BCH = $-\frac{1}{2}\kappa + \frac{1}{6}[\Omega_0, \kappa] + \cdots$ and the only structurally admissible correction to the two quantities the leading exponential misses — the Jarlskog invariant and |V_td| — without disturbing the inherited Cabibbo-scale entries.

In the lepton regime L_L = (ν_L, e_L) is asymmetric: the charged lepton is anchored, the neutrino weakly committed. If $T_\nu = D_0 \cdot I + \varepsilon M_\nu$, the neutrino frame is set by M_ν for every $\varepsilon > 0$, so large PMNS angles follow from the collapse of the stiffness hierarchy and are compatible with arbitrarily small neutrino masses. CKM and PMNS are then the small-gap and collapsed-gap regimes of one ratio.

So CKM and PMNS are two regimes of a single question: *how does a left-handed weak doublet split a common completion frame into sectoral mass-readout frames?* This paper does not finish that derivation. It does something more useful at this stage: it reduces the quark task to computing one commutator κ and the neutrino task to deriving one near-degenerate operator M_ν, and it names exactly how each must succeed or fail. That is the right next Standard Model step for VERSF — not another isolated number, but the electroweak operator that turns completion geometry into flavour frames.

References

VERSF flavour architecture

- Keith Taylor, *The Yukawa Operator from Completion-Channel Misalignment in VERSF.* Left Yukawa-square operator; CKM/PMNS as relative sector frames; CKM relative exponential; QF-1 premise; CP and PMNS open problems.
- Keith Taylor, *Deriving Flavour Mixing from Closure Geometry: CKM and PMNS Structure in the VERSF Framework.* Generation-space transport; $D = \mathrm{diag}(1, 2, 4)$; $\eta = 3/5$; $\chi = 2/5$; Cabibbo input 9/40; 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.* K = 7; 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.

Dynamical realisation context

- Keith Taylor, *The Twisted-Light Void Anchoring Framework.* Dynamical candidate realisation of geometric mass and mixing structures.

Numerical note. The exactness of the product-form cancellation $V\_CKM = \exp(\Omega\_q)$; the fact that the single-exponential doublet generator yields the additive sector frames $\exp(\Omega_0 \pm \Omega\_q/2)$; the equivalence of the two readings precisely at $[\Omega_0, \Omega\_q] = 0$; the curvature residue $\Delta\_BCH = -\frac{1}{2}[\Omega_0, \Omega\_q] + \frac{1}{6}[\Omega_0, [\Omega_0, \Omega\_q]] + \cdots$ verified to third order by perturbative scaling; and the ε-independence of the neutrino eigenframe for $T\_\nu = D_0 \cdot I + \varepsilon M\_\nu$ — were all checked by direct matrix computation. The §8 repair window was checked by a five-target least-squares solve over the common generator Ω_0 at fixed $\varphi = 2\pi/3$: a clean repair exists with $\|\Omega_0\| \approx 0.035$, $\|\kappa\| \approx 0.007$, $J \approx 3.1\times10^{-5}$ and $|V\_td| \approx 0.0083$ while the Cabibbo-scale entries hold to $\approx 1\sigma$; in that window $\|-\frac{1}{2}\kappa\|$ matches the full residue $\|\Delta\_BCH\|$ to four figures and the cubic term is $\approx 100\times$ smaller. Freeing φ leaves the small-$\|\Omega_0\|$ solution at the C_3 value, so the repair does not require moving the holonomy. The consumed CKM magnitudes and their inherited / consistency / output classification, the $|V\_ub|$ inclusive/exclusive split, and $|V\_ub|/|V\_cb| = 0.083 \pm 0.004$ follow the Yukawa-operator draft and PDG 2024 (Navas et al., Phys. Rev. D 110, 030001). The Jarlskog sign convention is held consistent with the flavour series; magnitudes are reported throughout.