The First Weak-Doublet Projection Audit in VERSF

Reducing P_W H_cl P_W to a Finite List of CKM Curvature, PMNS Support-Trace, Phase-Branch, and Octant Outputs

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

Successor to *The Weak-Doublet Projection Theorem in VERSF*.

Summary for the General Reader

Some of the deepest open questions in particle physics are about *mixing*. Quarks come in families, and when they feel the weak force they refuse to stay in their own family — they blend into one another by fixed, very specific amounts. Neutrinos do the same thing, but with a completely different blending pattern. In the standard textbook account nobody explains either pattern; they are simply measured in experiments and written into the theory by hand, as two separate tables of numbers.

The research programme behind this paper makes a bolder bet: that the quark blend and the neutrino blend are not two unrelated accidents but two views of a single deeper structure. Get that one structure right, and both tables should fall out of it on their own, instead of being inserted by hand.

This paper does **not** claim to have finished that calculation. It does something more disciplined, and more honest. It draws up a precise checklist — a short, fixed list of exactly what the deeper theory must produce by itself for the claim to count — and it imposes a strict rule while doing so: no peeking at the measured answers. A theory that quietly tunes itself to match the data isn't explaining anything, it's just copying. So every item has to be *earned* from the structure, not borrowed from experiment, and each is graded honestly: already proven, provable if one stated condition holds, or still owed.

The sharpest results concern the quark sector's fingerprint in the matter–antimatter imbalance of the universe — a tiny, *oriented* asymmetry that physicists package as the CKM "triangle." The paper shows the entire effect collapses to a single quantity the deeper theory must deliver: one number, with a definite size and a definite direction (the direction being which way the asymmetry points — the difference between our universe and its mirror image). It then goes further and pins down *where* that number comes from. The **size** is now reduced to a single clean

structural question: *if* the deep structure treats the three particle families even-handedly before their masses are read off, the size follows automatically from sharing one inherited quantity evenly across them. (That even-handedness is the one thing still to be proven, but it is now a sharp yes/no property rather than a guess.) The **direction** is reduced to a single clean principle about how the deepest dynamics pick the shorter of two possible paths. If the theory produces that number, the asymmetry is explained; if it produces the mirror-image number, the prediction is simply wrong — there is no middle ground to hide in.

One honest caveat runs underneath all of this. The paper has *not* computed that number from scratch. It worked backwards from the measured pattern to find the exact target the deep theory must hit — which is real progress, because the target is now precise and falsifiable — but the task ahead is to produce the same number *forwards*, from first principles, without ever looking at the experimental data. Knowing exactly what to aim for is not the same as having hit it, and the paper is careful never to confuse the two.

That is the spirit of the paper. It trades vague plausibility for a hard scorecard. If the deeper theory returns the listed quantities, the programme graduates from an appealing picture to a genuine derivation. If it doesn't, the failure is clean and local: the one item that breaks tells you exactly where the picture is wrong.

Contents

Abstract

The weak-doublet projection theorem reduced the residual Standard Model flavour problem to the projected closure Hamiltonian **$H\_W = P\_W\ H\_cl\ P\_W$**, where H_cl is the inherited su(8)-restricted admissibility Hessian and P_W is the combined generation, weak-role, and mass/readout projection. The theorem fixed the admissible block structure that could unify the committed-quark CKM regime with the weak-commitment neutrino PMNS regime. It did not evaluate H_W.

This paper defines the first **projection-audit protocol** for H_W. It decomposes the weak-doublet frame generator by weak role,

$$\Omega\_W = \Omega\_even \otimes I + \Omega\_odd \otimes \tau_3 + \Omega\_mix,$$

and then names the finite set of outputs a successful microscopic calculation must return. Each output is assigned exactly one grade — **exact**, **audit condition**, **conditional target**, or **owed projected output** — and the paper's discipline is that nothing is allowed to sit between them.

Two results are **exact** once the relevant blocks are present. First, with the first shared curvature in the 2↔3 channel, $\Omega_0(z) = z\, E_{23} - z\, E_{32}^*$, the triangle-generating matrix element of the commutator with the Cabibbo-dominated role-odd block Ω_q is fixed by one matrix-element sum:

$$[\Omega_0(z), \Omega\_q]_{13} = -a\, z.$$

(For generic Ω_q and complex z the commutator carries other nonzero entries; the *exact* claim is this single triangle element, confirmed to machine precision in Appendix C.) Second, under the committed doublet-split convention the Baker–Campbell–Hausdorff expansion gives the leading triangle correction

$$\Delta_{13} = \tfrac{1}{4}\, a\, z,$$

with the coefficient ¼ established symbolically and checked by a scaled-generator test (Appendix C): the ratio of the full closed-form residue to ¼ a z tends to 1 as the generators are scaled down, isolating it as the genuine leading term. (¼ is the half-split value; convention-independently $\Delta_{13} = \tfrac{1}{2} a \zeta$ with the invariant curvature $\zeta = \lambda z$, §4.2.)

The remaining quantities are **conditional targets** or **owed projected outputs**: the C_3-normalised amplitude $|\zeta| = b/\sqrt{3}$; the weak-commitment neutrino form $T\_\nu = D_0\, I + \varepsilon\, M\_\nu$; the residual-support ratios $\rho\_\odot = \sqrt{3/2}$, $B/A = 6/5$, and $\delta = 0 + O(\beta^2)$; the leakage support trace $\mathrm{Tr\_support}(\Pi\_leak) = (K - 2) \times 4 = 20$ with $K = 7$, which fixes $\beta = \sqrt{3}/20$; the atmospheric octant sign σ_W; and the two CP phase-branch signs $\sigma\_\varphi\hat{}q$ and $\sigma\_\varphi\hat{}\nu$.

For the quark sector these collapse to a single statement: the entire CKM curvature candidate, written through the convention-invariant curvature $\zeta = \lambda\, z$ as $\zeta\_C3 = (b/\sqrt{3})\, e\hat{}(5\pi i/6)$, is equivalent to one complex entry of the projected role-even block, **$(H\_even\hat{}q)_{23} = (b/(\lambda\ \varepsilon\_W$**

$\sqrt{3}$)) e^(iπ/3), with λ- and ε_W-independent phase π/3. The CKM curvature claim therefore reduces to whether the projected Hamiltonian returns this one matrix element with the correct modulus, phase, and sign.

A delimiting caveat runs through the paper and is stated here at the outset: **CKM measures the output, not H_cl.** This paper reconstructs the projected-Hamiltonian target that a successful calculation must return, and proves which parts of the CKM/PMNS structure are exact, conditional, or still owed; it does not evaluate the substrate Hessian, and it names the independent first-principles calculation that derivation would require (§4.6).

The contribution of this paper is methodological force, not a new ansatz. It converts the residual flavour problem into a finite, falsifiable list of projected-Hamiltonian outputs. If any output fails, the failed invariant localises the revision.

---

0. Predictive-Content Ledger

The previous theorem named the weak-doublet object. This paper names the first audit outputs and grades each one.

Object or claim	Required output	Grade
Weak-doublet Hamiltonian	H_W = P_W H_cl P_W	inherited object
Role decomposition	Ω_W = Ω_even ⊗ I + Ω_odd ⊗ τ_3 + Ω_mix	**exact**
Role-changing suppression	‖Ω_mix‖ ≪ ‖Ω_even‖ + ‖Ω_odd‖	**audit condition**
Quark role-odd block	Cabibbo-dominated Ω_q (a ≫ b ≫ c)	**owed projected output**
First shared curvature	role-even 2↔3 channel Ω_0(z)	**owed projected output**
CKM residue	[Ω_0, Ω_q]$_{13}$ = −a z	**exact**
CKM correction	Δ_{13} = ½ a ζ (= ¼ a z)	**exact at leading order** (coefficient λ/2 exact; tail ~0.1%)
CKM amplitude		ζ
C_3 covariance (amplitude)	[P_C H_cl P_C, R_3] = 0 ⟹	ζ
CKM phase branch	σ_φ^q selects 5π/6 vs 11π/6	**owed projected output** (discrete)
Minimal Hermitian lift	h² ∝ O_q + shortest lift ⟹ σ_φ^q = +1	**owed admissibility axiom** (§4.5.1)
Neutrino weak block	T_ν = D_0 I + ε M_ν	**owed projected output**
Eigenframe invariance	T_ν, M_ν share eigenvectors ∀ ε > 0	**exact**

Object or claim	Required output	Grade
Solar ratio	$\rho_\odot = \sqrt{3/2}$	**conditional target**
Pair ratio	$B/A = 6/5$	**conditional target**
Diagonal μ–τ split	$\delta = 0 + O(\beta^2)$	**conditional target**
PMNS phase branch	σ_φ^ν selects $3\pi/4$ vs $7\pi/4$	**owed projected output** (discrete)
Leakage support trace	$\mathrm{Tr\_support}(\Pi\_leak) = 20$	**conditional target** (central)
Leakage amplitude	$\beta = \sqrt{3}/20$	derived only if trace = 20
Octant sign	σ_W	**owed projected output** (discrete)

The grades are used strictly:

1. **Exact** — follows from matrix algebra once the stated blocks are present, independent of any data.
2. **Audit condition** — an admissibility inequality that must hold for the clean weak-doublet regime (e.g. the role gap). It is a precondition, not an output.
3. **Conditional target** — a finite projected quantity that holds *only under a stated structural condition* (a C_3 orbit, a support split, a trace factorisation). Not yet derived; falsified if the condition is not met. Where the conditional implication itself is proven and only its structural antecedent is owed, the entry is sharpened to a **conditional theorem** (e.g. the amplitude $|\zeta| = b/\sqrt{3}$, §4.3, where the norm algebra is exact and only the C_3-orbit structure of the projected block remains to be shown).
4. **Owed projected output** — a block or discrete sign that must be *returned* by evaluating H_cl on the relevant projected subspace. Cannot be upgraded to a derivation until that evaluation is done.

This paper succeeds only if it makes every owed output and conditional target finite, explicit, and falsifiable.

---

1. Purpose and the Standard of Proof

The weak-doublet projection theorem established the admissible structure of P_W H_cl P_W. That was a theorem about **shape**. It was not a calculation of the substrate Hessian.

This paper therefore asks the stricter question:

What must P_W H_cl P_W actually return for the VERSF flavour programme to count as a derivation rather than a fit?

The answer is a finite list:

Ω_q, $\Omega_0^{(23)}$, z, M_ν, Π_{leak}, $\mathrm{Tr\_support}(\Pi_{leak})$, σ_W, σ_φ^q, σ_φ^ν.

These are not optional refinements; they are the decisive outputs.

The purpose of the paper is **not** to introduce a new flavour ansatz. It is to *prevent* new ansätze from entering. Every surviving quantity must trace to exactly one of three sources: exact algebra, an inherited projection theorem, or a named substrate computation still owed. If a matrix entry is introduced because it improves agreement with CKM or PMNS data, the audit has failed *methodologically* — even if the number it produces looks good.

The audit does not succeed merely by reproducing familiar numerical targets. It succeeds only if the projected Hamiltonian returns the structures *before* any comparison with CKM or PMNS data is made. The decisive question is therefore not "can VERSF write a matrix that fits flavour?" but "does the independently defined block P_W H_cl P_W contain the required role-odd split, role-even curvature, weak-commitment kernel, leakage support trace, and branch signs?" This distinction is the difference between a parametrisation and a derivation.

1.1 What "calculation" means here

The word is used in a restricted sense. This paper does not claim the microscopic evaluation of H_cl is complete. A quantity counts as progress only if it can be obtained **without using the observed CKM or PMNS matrices as input**.

By that standard:

- $\mathbf{[\Omega_0(z), \Omega\_q]_{13} = -a\, z}$ is an exact calculation once $\Omega_0(z)$ and Ω_q are specified (§4.2).
- $\mathbf{|\zeta| = b/\sqrt{3}}$ (the convention-invariant amplitude, §4.2) is a *conditional theorem* (§4.3): the norm algebra $3|\zeta|^2 = b^2$ is exact, so the proof debt is not "why b/√3?" but only "does the projected block return the democratic, Killing-orthonormal C_3 orbit?"
- $\mathbf{\beta = \sqrt{3}/20}$ is *not* exact algebra. It is derived only if the leakage support trace is 20 (§6).

An output is therefore exact, conditional, or owed. It is not permitted to hide in between.

---

2. The Projected Weak-Doublet Object

Let H_cl be the su(8)-restricted admissibility Hessian inherited from the substrate free-energy functional:

H_cl = G | su(8).

Let the weak-doublet projection factor as

P_W = P_Y P_R P_C,

where P_C projects generation completion, P_R projects weak role, and P_Y projects mass/readout. The weak-doublet Hamiltonian is

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

and the associated anti-Hermitian frame generator is

$$\Omega_W = i\, \varepsilon_W H_W.$$

On the two-dimensional weak-role space the Pauli decomposition is exact:

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

The component blocks are recovered as partial traces over the role space (using $\mathrm{Tr}_{role}\, I = 2$, $\mathrm{Tr}_{role}\, \tau_i = 0$, $\mathrm{Tr}_{role}\, \tau_i \tau_j = 2\delta_{ij}$):

$$\Omega_{even} = \tfrac{1}{2} \mathrm{Tr}_{role}(\Omega_W),\ \Omega_{odd} = \tfrac{1}{2} \mathrm{Tr}_{role}(\Omega_W \tau_3),\ \Omega_1 = \tfrac{1}{2} \mathrm{Tr}_{role}(\Omega_W \tau_1),\ \Omega_2 = \tfrac{1}{2} \mathrm{Tr}_{role}(\Omega_W \tau_2),$$

with the role-changing part

$$\Omega_{mix} = \Omega_1 \otimes \tau_1 + \Omega_2 \otimes \tau_2.$$

This trace definition is what makes the later audit conditions precise: "Ω_{even} contains …" and "$\Omega_{odd} \to$ …" refer to these partial-trace blocks, not to a heuristic sandwich with role projectors. The clean weak-doublet flavour regime requires the role-changing part to be gapped:

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

If this gap fails, CKM and PMNS cannot be read as clean sectoral frames of the *same* projected block, and the unification the previous theorem aimed at is lost at leading order.

---

3. Audit Rules

The projection audit is governed by five rules. They exist to keep the construction honest.

Rule 1 — No observed flavour matrix as input. The CKM and PMNS matrices may be used only as comparison targets, never as construction data. If P_W, H_{cl}, or any support projector is *chosen* to reproduce observed entries, the audit has failed. A clarification this rule needs to carry: the prohibition is on inserting *new* flavour matrices by hand, not a claim that every number is already derived from first principles. Some numerical targets — in particular the inherited quark generator scales a, b, c (with $b = 81/2000$) — are carried over from earlier sector audits. They are not freshly fitted here, but neither are they yet substrate-derived: they remain first-principles outputs still owed by the H_{cl} calculation. So "no data input" means "no new flavour matrix inserted," not "every target derived."

Rule 2 — Charge-blind mass readout. Mass/readout may distinguish depth and stiffness, but not electric charge:

$$[P_Y, \chi] = 0 \Rightarrow [P_Y H_{cl} P_Y, \chi] = 0.$$

This forbids flavour generation from changing gauge representation.

Rule 3 — Role-gap admissibility. The role-changing part must be gapped or subleading:

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

This is what lets the weak-doublet partners define clean flavour frames at leading order.

Rule 4 — Separate continuous amplitudes from discrete signs. The amplitudes z and β do not fix every physical branch. The CP phases and the octant require discrete projected signs:

$$\sigma_\varphi^q = \mathrm{sgn}(\mathrm{Im}\, z),\ \sigma_\varphi^\nu = \mathrm{sgn}(\mathrm{Im}\, h_\nu),\ \sigma_W = \mathrm{sgn}(\|P_e H_W P_\tau\|^2 - \|P_e H_W P_\mu\|^2).$$

These must be *computed* from the projection. They are not conventions.

Rule 5 — Support traces must be projected traces. The denominator 20 in $\beta = \sqrt{3}/20$ is acceptable only if it is returned as the support trace of the leakage projector:

$$\mathrm{Tr}_{support}(\Pi_{leak}) = 20.$$

If 20 is chosen because it gives the right reactor angle, the derivation fails.

---

4. Quark-Sector Audit

In the committed quark regime both weak partners are closure-committed:

$$\gamma_u \simeq \gamma_d \simeq 1.$$

The leading CKM split is therefore expected in the **role-odd** block. The inherited quark generator has the anti-Hermitian form

$$\Omega_q = \begin{pmatrix} 0 & a & c \cdot e^{i\varphi} \\ -a & 0 & b \\ -c \cdot e^{-i\varphi} & -b & 0 \end{pmatrix}$$

with hierarchy **a ≫ b ≫ c**. The values of a, b, c (with b = 81/2000) are inherited from the earlier Yukawa / electroweak-frame audits, not fitted in this paper; per Rule 1 they are inherited targets that the substrate calculation still owes, not yet derived substrate numbers. The first quark audit condition is

$$\Omega\_odd \to \Omega\_q,$$

with the Cabibbo-dominated hierarchy intact. This is a *structural* test, not a numerical fit: the projected role-odd block must produce the correct ordering of generation channels without being told the observed CKM matrix.

4.1 The first shared curvature

The role-even block supplies common-mode curvature. The previous theorem selected the 2↔3 channel as the first admissible, Cabibbo-protected curvature:

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

The audit condition is

Ω_even contains the leading curvature $\Omega_0\text{^}(23)(z)$,

with no leading 1↔2 shared curvature and no directly inserted 1↔3 shared curvature. The reason is geometric:

- a 1↔2 shared curvature directly disturbs the Cabibbo doorway;
- a 1↔3 shared curvature inserts the missing triangle entry by hand;
- a 2↔3 shared curvature preserves the Cabibbo doorway at first order, yet still *generates* a 1↔3 residue by failing to commute with the large 1↔2 role-odd component.

This is the first point where the audit bites. If the projected role-even block does not return the 2↔3 channel as the leading shared curvature, the CKM curvature mechanism is not derived.

4.2 Exact CKM residue

For $\Omega_0(z) = z\, E_{23} - z\, E_{32}^*$, the (1,3) entry of the commutator follows from a single matrix-element sum. Because row 1 of Ω_0 vanishes, $(\Omega_0\, \Omega\_q)_{13} = 0$, while

$$(\Omega\_q\ \Omega_0)_{13} = \Sigma\_k\ (\Omega\_q)_{1k}\ (\Omega_0)_{k3} = (\Omega\_q)_{12}\ (\Omega_0)_{23} = a \cdot z.$$

Hence

$$[\Omega_0(z), \Omega\_q]_{13} = (\Omega_0\ \Omega\_q - \Omega\_q\ \Omega_0)_{13} = -a\ z.$$

For generic Ω_q (nonzero $c\cdot e\text{^}(i\varphi)$) and complex z the full commutator carries other nonzero entries; the exact, data-independent claim is this single **triangle-generating** element. It is reproduced to machine precision for generic parameters in Appendix C.

Under the committed doublet-split convention,

$$U\_u = \exp((\Omega_0 - \Omega\_q)/2),\ U\_d = \exp((\Omega_0 + \Omega\_q)/2),\ V\_CKM = U\_u^\dagger\ U\_d,$$

write $X = (\Omega\_q - \Omega_0)/2$ and $Y = (\Omega_0 + \Omega\_q)/2$, so that $V\_CKM = \exp(X) \exp(Y)$. Then $X + Y = \Omega\_q$, and the first Baker–Campbell–Hausdorff term is

$$\tfrac{1}{2} [X, Y] = \tfrac{1}{2} \cdot \tfrac{1}{4} [\Omega\_q - \Omega_0, \Omega_0 + \Omega\_q] = -\tfrac{1}{4} [\Omega_0, \Omega\_q].$$

Therefore

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

and the leading triangle correction is

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

This part is exact once the 2↔3 shared curvature exists. The algebra does not know the observed CKM triangle; it knows only the non-commutation of the role-even and role-odd blocks. The proof is the matrix-element sum above together with the BCH expansion; the scaled-generator test of Appendix C verifies the implementation, not the result — the exact algebra is not discovered numerically.

The residue depends only on $(\Omega\_q)_{12} = a$ and the vanishing of row 1 of Ω_0: it is independent of b, c, and φ — it holds even at b = c = 0 (Appendix C). The triangle mechanism therefore never touches Ω_q's own 1↔3 entry; the missing (1,3) element is *generated* by non-commutation with the Cabibbo doorway, never inserted by hand, and the result is robust to the rest of the role-odd block. This is the algebraic form of the §4.1 channel argument, and a direct rebuttal to any charge that the 1↔3 entry was put in by hand.

Convention note. The coefficient is convention-dependent. The earlier *CKM Curvature Residue* paper builds the curved matrix as $V\_curv = \exp(-\Omega_0 + \tfrac{1}{2}\Omega\_q)\cdot\exp(\Omega_0 + \tfrac{1}{2}\Omega\_q)$, with the **full** Ω_0 in each exponent, and records $\Delta_{13} = \tfrac{1}{2}\, a\, z$. The present paper uses the half-split $U\_u = \exp((\Omega_0 - \Omega\_q)/2)$, $U\_d = \exp((\Omega_0 + \Omega\_q)/2)$, with $\mathbf{\Omega_0/2}$ in each exponent, giving $\Delta_{13} = \tfrac{1}{4}\, a\, z$. Both are internally correct, but importing an amplitude target from one convention into the other unchanged would mis-state the triangle correction by a factor of two. The next paragraph removes the ambiguity at the root.

Convention-invariant curvature variable

To fix the normalisation once and for all, introduce a curvature-split parameter λ and write the factorisation as

$$V\_CKM = \exp(-\lambda\, \Omega_0 + \tfrac{1}{2}\, \Omega\_q) \cdot \exp(\lambda\, \Omega_0 + \tfrac{1}{2}\, \Omega\_q).$$

The leading BCH correction is then

$$\log V\_CKM = \Omega\_q - (\lambda/2) [\Omega_0, \Omega\_q] + \cdots, \text{ so } \Delta_{13} = (\lambda/2)\, a\, z.$$

The half-split of this paper is **$\lambda = ½$** (recovering $\Delta_{13} = ¼\ a\ z$); the earlier *CKM Curvature Residue* paper is **$\lambda = 1$** (recovering $\Delta_{13} = ½\ a\ z$). The convention-independent quantity is the **invariant curvature**

$\zeta = \lambda\ z$, so that $\Delta_{13} = ½\ a\ \zeta$ in every convention.

Physical targets are therefore stated for ζ, not for the convention-dependent z. The convention-invariance is structural, not coincidental: because λ is real, $\lambda\ \Omega_0(z) = z\ \lambda\ E_{23} - z^*\ \lambda\ E_{32} = \Omega_0(\lambda z) = \Omega_0(\zeta)$, so the two exponentiated factors $\exp(\mp\lambda\Omega_0(z) + ½\Omega\_q) = \exp(\mp\Omega_0(\zeta) + ½\Omega\_q)$ are *literally the same matrices* in both conventions. V_CKM is thus a function of ζ alone, and identical physics in the two conventions is manifest rather than a numerical accident (Appendix C.4 confirms it, including the shared higher-order tail). Throughout this paper, any amplitude written as "$z = b/\sqrt{3}$" is shorthand for the $\lambda = 1$ normalisation; the convention-invariant content is **$|\zeta| = b/\sqrt{3}$**, with $z = \zeta/\lambda$ — so $z = b/\sqrt{3}$ at $\lambda = 1$ and $z = 2b/\sqrt{3}$ at the $\lambda = ½$ used here. The same physical residue ζ_C3 thus corresponds to $z = \zeta\_C3$ in the full-Ω_0 convention and $z = 2\ \zeta\_C3$ in the half-split convention, closing the factor-of-two gap between the two papers.

4.3 Amplitude audit: the C_3 Killing-norm-sharing theorem, $|\zeta| = b/\sqrt{3}$

The amplitude part of the CKM curvature target is stronger than the phase part: it is not obtained by fitting the triangle. It follows from an inherited 2↔3 transport norm distributed over a threefold C_3 generation orbit, provided the projected role-even curvature decomposes into Killing-orthonormal branches. Stated for the convention-invariant curvature $\zeta = \lambda z$ (§4.2), the amplitude target is **$|\zeta| = b/\sqrt{3}$**.

Assume the role-even 2↔3 curvature returned by H_W is the common-mode projection of the committed 2↔3 transport norm b. In the C_3 generation loop this shared curvature decomposes into three branch components,

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

with total inherited norm $\|Z\|\_K = b$. By Killing-orthonormality,

$$\|Z\|^2\_K = \|Z_1\|^2\_K + \|Z_2\|^2\_K + \|Z_3\|^2\_K = 3\ |\zeta|^2 \Longrightarrow 3\ |\zeta|^2 = b^2 \Longrightarrow |\zeta| = b/\sqrt{3}.$$

The square root is the content of the theorem: the branches share a *norm*, not a linear amplitude. A coherent linear split would give $|\zeta| = b/3$; no split would give $|\zeta| = b$; only the C_3-democratic Killing-orthonormal projection gives the root-normalised amplitude $b/\sqrt{3}$. The factor is therefore structural, not adjustable.

The theorem is conditional on three projected-Hamiltonian facts: (i) the leading shared curvature is the 2↔3 role-even channel; (ii) that curvature is the common-mode projection of the inherited committed 2↔3 transport norm b; (iii) the C_3 branches are Killing-orthonormal and democratic. Under (i)–(iii) the amplitude is *derived* — the norm algebra is not in question. What remains owed is not "why $b/\sqrt{3}$?" but the substrate fact "does P_W H_cl P_W actually return this democratic Killing-orthonormal C_3 orbit?" That relocates the proof debt from a numerical guess to one structural property of the projected Hamiltonian. If the C_3 branches are absent, non-orthonormal, or non-democratic, the theorem does not apply and ζ is again an owed number.

4.3.1 The C_3 orbit exists; only its democratic population is owed

The theorem rests on a Killing-orthonormal C_3 orbit, and that orbit is not hypothetical — it exists explicitly in the three-generation algebra. The elementary anti-Hermitian branch generators

$$\mathbf{X_{12} = E_{12} - E_{21},\ X_{23} = E_{23} - E_{32},\ X_{31} = E_{31} - E_{13}}$$

are cycled by the C_3 generation rotation ($X_{12} \to X_{23} \to X_{31} \to X_{12}$) and, under the compact Killing form $\langle A, B\rangle_K \propto -\mathrm{Tr}(A B)$, have Gram matrix

$$\boldsymbol{\langle X_i, X_j\rangle_K \propto 2\, \delta_{ij},}$$

so they are mutually orthogonal with equal norm (Appendix C.5). Three claims must then be separated:

- **(A) the orbit exists** — the generation algebra carries this C_3 triple of off-diagonal branch generators. *Proven.*
- **(B) the orbit is Killing-orthonormal** — the three branches are orthogonal and equal-norm under the trace form. *Proven.*
- **(C) H_W populates the orbit democratically** — the role-even quark block places equal Killing weight on the three branches, $\Pi_{C_3}(\Omega_{even}^q) = \zeta_1 Z_1 + \zeta_2 Z_2 + \zeta_3 Z_3$ with $|\zeta_1| = |\zeta_2| = |\zeta_3|$, before any branch is selected. *Owed.*

Only (C) is outstanding; (A) and (B) are settled. The amplitude theorem therefore needs no new algebra — it needs one equivariance fact about the projected Hamiltonian.

Support level, not visible level. A caution prevents a clash with the channel rule of §4.1. The democratic orbit must *not* mean Ω_even literally shows equal 1↔2, 2↔3, and 3↔1 entries — that would reinstate a leading 1↔2 common-mode curvature and disturb the Cabibbo doorway. The democracy lives at the *support/norm* level of the projected closure geometry, while the weak-doublet readout selects the Cabibbo-protected 2↔3 branch as the first *visible* common-mode curvature. In one line: **the C_3 orbit supplies the root-normalisation (the ÷√3); Cabibbo protection supplies the visible channel.** The √3 thus appears without forcing all three matrix entries into the visible Ω_even block.

The real test for (C). Discharging the owed democratic population reduces to three checks on P_W H_cl P_W:

1. **C_3 covariance** — the generation-completion part of the projection is equivariant under the C_3 generation rotation.
2. **Killing orthogonality** — the three branch supports are mutually orthogonal under ⟨·,·⟩_K (held at the algebra level by (B); required to survive projection).

3. **Democratic projection** — the role-even common-mode support carries equal branch weight before the readout selects the visible 2↔3 branch (made operational as a pre-readout weight test in §4.3.2).

If all three hold, $|\zeta| = b/\sqrt{3}$ is derived; if any fails, the amplitude is not. This is the sharply defined replacement for "why $b/\sqrt{3}$?": the proof debt is now exactly *"does the projected closure geometry act C_3-democratically on the three Killing-orthonormal branch supports?"*

4.3.2 The democracy test is pre-readout

A caution sharper than the support/visible split: the democracy test must be applied *before* the mass/readout projection, not after it. Since $P\_W = P\_Y P\_R P\_C$, the final readout P_Y may itself select the visible 2↔3 branch. Testing the fully read-out block and finding only a 2↔3 entry would then *not* disprove C_3 democracy — the democratic support could have been present before readout, with P_Y choosing the Cabibbo-protected branch. The test therefore splits into two stages with two distinct objects:

Pre-readout (democracy): $K\_even^q = \tfrac{1}{2}\,\mathrm{Tr\_role}(P\_q\, P\_R\, P\_C\, H\_cl\, P\_C\, P\_R\, P\_q)$ — P_Y omitted. Post-readout (visible curvature): $H\_even^q = \tfrac{1}{2}\,\mathrm{Tr\_role}(P\_q\, P\_Y\, P\_R\, P\_C\, H\_cl\, P\_C\, P\_R\, P\_Y\, P\_q)$ — the §4.5 entry object.

On the three-generation space, use the Hermitian branch quadratures

$$S\_ij = (E\_ij + E\_ji)/\sqrt{2},\quad A\_ij = i(E\_ij - E\_ji)/\sqrt{2},\quad B\_ij = \mathrm{span}\{S\_ij, A\_ij\},$$

which are Hermitian and orthonormal under $\langle X, Y\rangle\_K \propto \mathrm{Tr}(X\,Y)$, with the three supports B_{12}, B_{23}, B_{31} mutually orthogonal and cycled by the C_3 rotation (Appendix C.5). The pre-readout branch weights are

$$w\_ij = \|\Pi\_\{B\_ij\}\, K\_even^q\|^2\_K \;\left(\propto |(K\_even^q)\_ij|^2 \text{ in a fixed branch normalisation } \right),$$

where all weights are taken in the *same* Killing normalisation as the inherited norm b. In a raw matrix-entry convention an overall constant appears (for the Hermitian quadratures above, $\|h\, E\_ij + h^*\, E\_ji\|^2\_K = 2|h|^2$); that constant cancels in the democracy score D_C_3 but must be fixed once and used consistently in the absolute norm-inheritance test $\Sigma\, w = b^2$.

and the democracy and norm-inheritance tests are

$$D\_C_3 = \left[(w_{12} - \bar{w})^2 + (w_{23} - \bar{w})^2 + (w_{31} - \bar{w})^2 \right] / \bar{w}^2 \to 0 \quad (\bar{w} = \tfrac{1}{3}\Sigma\, w),\quad w_{12} + w_{23} + w_{31} = b^2.$$

If both pass, each branch carries $w\_ij = b^2/3$ and the inherited curvature amplitude is $|\zeta| = b/\sqrt{3}$. The score discriminates: a democratic block gives $D\_C_3 = 0$, while a 2↔3-lopsided block of the *same* total norm gives $D\_C_3 = 6$ (Appendix C.6) — so finding the inherited norm is not enough; it must be shared. The post-readout test is then separate: H_even^q must show the 2↔3 branch as the first visible common-mode curvature, with $|(H\_even^q)_{23}| = b/(\lambda\, \varepsilon\_W \sqrt{3})$ and $\arg = \pi/3$.

The required democratic pass-values are explicit, giving the next calculation a concrete target. In convention-invariant curvature units the test object is normalised so that

$w_{12} = w_{23} = w_{31} = b^2/3 = 0.00054675$, branch amplitude $\sqrt{w} = b/\sqrt{3} = 0.0233827$ ($b = 81/2000$).

In Hamiltonian-entry units the weights carry the convention factor, $w_{ij} = b^2/(3\,\lambda^2\,\varepsilon_W^2)$: with $\varepsilon_W = 1$ this is 0.00054675 at $\lambda = 1$ and 0.002187 at the $\lambda = ½$ of this paper — the explicit reason the convention factor must stay visible. Crucially, these weights must be read off the *pre-readout* K_{even}^q. Evaluating them instead on the read-out entry $(H_{even}^q)_{23} = (b/(\lambda\,\varepsilon_W\,\sqrt{3}))\,e^{i\pi/3}$ — which carries only the 2↔3 channel — would return $(w_{12}, w_{23}, w_{31}) = (0, b^2/(3\,\lambda^2\,\varepsilon_W^2), 0)$, a maximal D_{C_3} that *falsely* fails democracy. The democratic support and the visible single channel are not in conflict; they live on opposite sides of P_Y.

The safe statement of the amplitude claim is therefore: **the C_3-democratic norm-sharing is a pre-readout support property** — the generation/role projection $P_R\,P_C$ must carry a three-branch Killing-orthonormal C_3 orbit with equal support weights — **and the mass/readout projection P_Y then selects the Cabibbo-protected 2↔3 branch as the first visible curvature.** With the explicit H_{cl} and projectors still owed, the current status is precise:

- **Test A — support geometry:** the Killing-orthonormal C_3 orbit exists. *Proven* (§4.3.1, C.5).
- **Test B — democratic population:** $w_{12} = w_{23} = w_{31} = b^2/3$ for the pre-readout K_{even}^q. *Owed*, but now reduced to a single commutator, $[P_C\,H_{cl}\,P_C, R_3] = 0$ (§4.3.4).
- **Test C — readout selection:** P_Y selects the visible 2↔3 branch. *Conditionally argued* (Cabibbo protection, §4.1), not yet computed.

4.3.3 A minimal substrate cell for the C_3 democracy

The democratic population (Test B) can be pushed from "owed" to "derived in a minimal model" for one sector, without inventing any CKM-specific structure. The construction restricts the *inherited* substrate free-energy functional $F[\rho] = \int V_{sub}\,d\mu + A_C + A_T + A_R$ (with $V_{sub} = (\lambda_{sub}/4)(\|\rho\|^2 - \rho_c^2)^2$) to the minimal weak-doublet cell

$$\mathcal{H}_{min} = \ell^2(\mathcal{G}_3) \otimes \mathbb{C}^2_{role} \otimes \mathbb{C}^8_{closure},$$

where $\mathcal{G}_3 = \{1, 2, 3\}$ carries the three generation supports and the generation-cycle operator R_3 cycles the branch supports $B_{12} \to B_{23} \to B_{31}$. No CKM matrix, z, or empirical phase enters the functional.

Closure-sector engine (derived). Write small role-even branch transports as $U_{ij} = \exp(i\varepsilon\,X_{ij})$ with $X_{ij} \in B_{ij}$. The closure-holonomy penalty $A_C \propto \|U_{12}\,U_{23}\,U_{31} - \mathbb{1}\|^2_K$ expands as

$$U_{12}\,U_{23}\,U_{31} - \mathbb{1} = i\varepsilon\,(X_{12} + X_{23} + X_{31}) + O(\varepsilon^2),$$

so at quadratic order $A_C^{(2)} \propto \|X_{12} + X_{23} + X_{31}\|^2_K$. By Killing-orthogonality of the branches the cross terms vanish and this collapses to $\|X_{12}\|^2_K + \|X_{23}\|^2_K + \|X_{31}\|^2_K$ — i.e. the closure

Hessian on the branch orbit is the scalar $\mathbf{2\cdot\mathbb{1}_3}$ (Appendix C.7). The inherited closure penalty therefore acts *democratically* on the three Killing-orthonormal branches, with no CKM input.

Honest scope (the Schur step corrected). This does not yet make the *full* Hessian scalar. The transport and record penalties A_T, A_R are not automatically democratic: C_3 is abelian, so the three-branch representation is reducible (a symmetric isotype ⊕ a 2-dimensional doublet), and a C_3-invariant quadratic form may carry *two distinct* isotype eigenvalues rather than one scalar (Appendix C.7 exhibits eigenvalues [0.9, 0.9, 2.1]). An SU(8)-Schur argument does not repair this, because SU(8) acts on the closure fibre $\mathbb{C}^8$ and is a singlet on the generation branches — it is the wrong group for branch democracy. The correct sufficient condition is supplied below: not a scalar Hessian, but C_3-covariance of the pre-readout block.

4.3.4 Conditional reduction theorem

Theorem (C_3 cell reduction). The actual weak-doublet projection reduces to the minimal C_3 Killing-orthonormal substrate cell at leading flavour order if:

1. **Generation closure** — P_C closes the substrate Hessian onto the three-generation module, P_C H_cl P_C ⊂ End($\mathcal{G}_3$).
2. **C_3 covariance** — the pre-readout generation-completion block commutes with the generation cycle, [P_C H_cl P_C, R_3] = 0.
3. **Role separation** — P_R preserves the weak-role grading and Ω_mix is gapped (§3, Rule 3).
4. **Charge-blind readout** — P_Y alters no gauge charge and acts only after the C_3 support has formed (§3, Rule 2).
5. **No branch-bias term** — before P_Y, H_cl contains no term distinguishing B_{12}, B_{23}, B_{31}.

Under (1)–(5) the pre-readout role-even block K_even^q = ½ Tr_role(P_q P_R P_C H_cl P_C P_R P_q) is C_3-covariant, R_3 K_even^q $R_3^\dagger$ = K_even^q. The cyclic permutation then forces equal branch magnitudes, $|(\text{K\_even}^q)_{12}| = |(\text{K\_even}^q)_{23}| = |(\text{K\_even}^q)_{31}|$ (Appendix C.7) — democracy follows from covariance of the *configuration*, not from a scalar Hessian. With Killing-orthonormality and inherited norm b, $w_{12} = w_{23} = w_{31} = b^2/3$, hence $|\zeta| = b/\sqrt{3}$. The readout P_Y may then select the Cabibbo-protected 2↔3 branch as the first visible curvature without disturbing the support-level democracy.

To be precise about what covariance does and does not give: C_3 covariance forces equal branch *magnitudes* in the off-diagonal generation cycle, but it does not by itself force every internal isotype coefficient of the Hessian to be equal (§4.3.3). The amplitude proof requires covariance on the relevant role-even branch support — equivalently, equal branch weights of the configuration K_even^q — not the stronger blanket claim that the full Hessian is scalar.

The bottleneck. Conditions (1), (3), (4) are inherited or structural; condition (5) is the no-isotype-splitting requirement of §4.3.3. The entire amplitude derivation therefore reduces to one commutator on the role-even committed-quark support:

[P_C H_cl P_C, R_3] = 0.

If it vanishes, the minimal C_3 cell is not an extra assumption — it is the leading projected substrate cell, and $|\zeta| = b/\sqrt{3}$ is derived from inherited structure with no CKM input. If it does not vanish, the amplitude remains a conditional theorem. This is a *conditional reduction theorem*, not a completed derivation: the commutator has not been evaluated from an explicit H_cl, because the substrate functional and projectors are not yet specified numerically (§4.6). But the owed quantity is now a single, sharply defined object — the next calculation is to compute or prove $[P\_C\ H\_cl\ P\_C, R_3] = 0$.

C_3 Equivariance Lemma. That commutator need not be computed entry-by-entry; it follows from a symmetry premise. Let R_3 act on the generation-completion cell by cyclic permutation of the three branch supports, and let $G = \delta^2F|\_\{\mathcal{M}\_adm\}$ be the Hessian of the substrate functional at the admissible background x*. If

- (a) F is C_3-invariant on the pre-readout generation-completion sector, $F[R_3\ x] = F[x]$, and
- (b) the admissible background is C_3-symmetric, $R_3\ x^* = x^*$,

then differentiating $F[R_3\ x] = F[x]$ twice at x* gives $R_3^\dagger\ G\ R_3 = G$, i.e. $\mathbf{[G, R_3] = 0}$. If in addition (c) the su(8) closure restriction and the generation projector P_C are R_3-equivariant ($P\_C\ R_3 = R_3\ P\_C$, and R_3 preserves the closure sector — it only permutes generation labels and adds no u(1) trace), then $[P\_C\ H\_cl\ P\_C, R_3] = 0$ on the role-even committed-quark support, and the amplitude democracy $w_{12} = w_{23} = w_{31} = b^2/3$ follows. The premise (b) is essential and not cosmetic: at a background that spontaneously breaks C_3, the Hessian does *not* commute with R_3 (Appendix C.7), so the lemma also names "no spontaneous C_3 breaking before readout" as a condition.

What this relocates, honestly. The lemma is a genuine advance — it replaces an opaque commutator computation with the standard fact that the Hessian of an invariant functional at a symmetric critical point commutes with the symmetry, using no CKM input. But it relocates the debt rather than discharging it: the owed items become (a) C_3-invariance of F on the pre-readout cell, (b) a C_3-symmetric admissible background, and (c) C_3-equivariance of P_C. Of these, (a) is the substantive one — it asserts that the substrate functional treats the three generations democratically *before* readout, with the mass hierarchy entering only through P_Y. This is physically natural (the hierarchy is a readout effect, not a substrate asymmetry), but it is an assumption about the explicit F, and a careful reader will note that functional-level C_3-invariance is close to the weight-democracy it implies — the two are the same democracy stated at different depths. The value of the lemma is that it pushes the assumption to the most fundamental and most checkable level: a single symmetry of the substrate functional, which the explicit F will either possess or not once it is supplied.

4.4 Phase audit: $\sigma\_\varphi\^q$

Complement-half-transport gives the quark phase equation

$$h\_q^2 = C \cdot O\_q,\ C = e\^(i\pi),\ O\_q = e\^(2\pi i/3) \Longrightarrow h\_q^2 = e\^(5\pi i/3).$$

The two antipodal square roots are

$h_q = e^{(5\pi i/6)}$ or $h_q = e^{(11\pi i/6)}$.

They are not interchangeable: they carry opposite imaginary signs and select opposite CP orientations. Define the branch sign **$\sigma_\varphi^q = \mathrm{sgn}(\mathrm{Im}\, z)$**. Then

$\sigma_\varphi^q > 0 \rightarrow \arg z = 5\pi/6$, $\sigma_\varphi^q < 0 \rightarrow \arg z = 11\pi/6$.

The CKM triangle encodes an *oriented* CP-violating area, not merely a magnitude. A two-branch curvature residue therefore cannot be the final prediction: it fixes the size of the triangle repair but leaves its handedness undetermined, exactly as the atmospheric octant is undetermined before σ_W is computed (§7). Until σ_φ^q is returned by $P_W H_{cl} P_W$, the framework has derived the triangle-residue mechanism and its magnitude target, but not the physical CKM orientation. The sign σ_φ^q is thus load-bearing: it is the invariant that turns a two-branch curvature mechanism into a single oriented CKM triangle.

"Principal branch" is not a derivation. But the branch is not a vague holonomy choice either: it reduces to the sign of one projected-Hamiltonian matrix element.

Lemma (Phase-Branch Extraction). Let $H_{even}^q = \frac{1}{2} \mathrm{Tr}_{role}(P_q P_W H_{cl} P_W P_q)$ be the role-even quark block (Hermitian), and $\Omega_{even}^q = i\, \varepsilon_W H_{even}^q$ its frame generator, with the fixed frame-orientation convention $\varepsilon_W > 0$. If the leading role-even curvature is $\Omega_0(z) = z E_{23} - z^* E_{32}$, then writing $h = (H_{even}^q)_{23}$ one has $z = i\, \varepsilon_W h$, hence $\mathrm{Re}\, z = -\varepsilon_W \mathrm{Im}\, h$ and $\mathrm{Im}\, z = \varepsilon_W \mathrm{Re}\, h$, so

$\sigma_\varphi^q = \mathrm{sgn}(\mathrm{Im}\, z) = \mathrm{sgn}\, \mathrm{Re}(H_{even}^q)_{23}$.

The proof is one line: $z = (\Omega_{even}^q)_{23} = i\, \varepsilon_W h$, and Hermiticity gives $(\Omega_{even}^q)_{32} = -z^*$ automatically, so the $\Omega_0(z)$ structure is not assumed but enforced. The lemma carries two riders. First, it presumes the fixed convention $\varepsilon_W > 0$; a silent sign flip in ε_W reverses the predicted CP orientation, so the convention must be stated once and held. Second, it extracts only the discrete *handedness*: it requires $(H_{even}^q)_{23}$ to be genuinely complex. If the projected role-even $2\leftrightarrow3$ block were real-symmetric ($\mathrm{Im}\, h = 0$), then $z = i\, \varepsilon_W \mathrm{Re}\, h$ is purely imaginary, forcing $\arg z = \pi/2$ rather than the $5\pi/6$ magnitude of the O_q structure above. The complex character of $(H_{even}^q)_{23}$ is therefore a separate owed fact (the source of the phase magnitude), distinct from the sign the lemma returns.

Quark phase-branch failure modes. The branch sign must be treated separately from the curvature magnitude, and four outcomes are distinct:

- **Orientation incompleteness** — σ_φ^q is not computed from the projection. The framework has then derived at most the *size and channel* of the CKM triangle repair; the oriented triangle is not yet derived. This is incompleteness, not falsification.
- **Branch success** — the projection returns $\sigma_\varphi^q = +1$, equivalently $\mathrm{Re}(H_{even}^q)_{23} > 0$ (with $\varepsilon_W > 0$). The curvature lies on the positive-imaginary branch, $\arg z = 5\pi/6$.
- **Branch mismatch** — the projection returns $\sigma_\varphi^q = -1$, selecting the antipode $\arg z = 11\pi/6$. This is not a harmless convention: it reverses the CKM CP orientation. Since the

triangle is an oriented CP object, this *falsifies* the proposed branch rule against the observed handedness.

- **Branch collapse** — $\sigma_\varphi^q = 0$. The projection supplies no first-order CP handedness; the residue may carry a magnitude but encodes no oriented area. This is a failure of the curvature mechanism *as a CP derivation*.

Only the second outcome completes the CKM CP structure; the third is a genuine falsification, and the first and fourth are, respectively, an incompleteness and a collapse.

4.5 The CKM curvature reduced to one Hamiltonian entry

Sections 4.1–4.4 collapse into a single target. Combining the amplitude (§4.3), the phase grammar (§4.4), the convention-invariant curvature (§4.2), and the extraction lemma, the entire C_3 curvature candidate $\zeta_{C3} = (b/\sqrt{3})\, e^{5\pi i/6}$ — the prior paper's z_{C3} in the $\lambda = 1$ normalisation — is equivalent to **one complex entry** of the projected role-even quark block:

$$(H_{even}^q)_{23} = (b/(\lambda\, \varepsilon_W \sqrt{3})) \cdot e^{i\pi/3}.$$

The equivalence is exact: $z = \zeta/\lambda = i\, \varepsilon_W\, (H_{even}^q)_{23}$, and the factor i rotates the Hamiltonian phase $\pi/3$ into the generator phase $5\pi/6$ ($i \cdot e^{i\pi/3} = e^{5\pi i/6}$). The convention-robust split is

$\arg(H_{even}^q)_{23} = \pi/3$ (phase; λ- and ε_W-independent), $|(H_{even}^q)_{23}| = b/(\lambda\, \varepsilon_W \sqrt{3})$ (modulus; carries the $\lambda\, \varepsilon_W$ normalisation).

So the whole CKM curvature claim reduces to whether the projected Hamiltonian returns one matrix element with a fixed phase and a fixed modulus.

Explicitly, in real and imaginary parts, the target entry and its Hermitian partner are

$$(H_{even}^q)_{23} = b/(2\lambda\sqrt{3}\, \varepsilon_W) + i \cdot b/(2\lambda\, \varepsilon_W),\quad (H_{even}^q)_{32} = b/(2\lambda\sqrt{3}\, \varepsilon_W) - i \cdot b/(2\lambda\, \varepsilon_W),$$

so the required 2↔3 block is

$$H_{even,C3}^q = [b/(2\lambda\sqrt{3}\, \varepsilon_W)]\, (E_{23} + E_{32}) + [i\, b/(2\lambda\, \varepsilon_W)]\, (E_{23} - E_{32}).$$

In this paper's convention ($\lambda = ½$) and frame units ($\varepsilon_W = 1$), with the inherited generator value $b = 81/2000$, the real part is $b/\sqrt{3}$ and the imaginary part is exactly **b**, giving $(H_{even}^q)_{23} = 0.023383 + 0.040500\, i$ and $z = i(H_{even}^q)_{23} = -0.040500 + 0.023383\, i = 2\, \zeta_{C3} = (2b/\sqrt{3})\, e^{5\pi i/6}$ — the half-split image of the invariant ζ_{C3} (Appendix C.3). In the full-Ω_0 convention ($\lambda = 1$) the same invariant gives the simpler $(H_{even}^q)_{23} = b/(2\sqrt{3}) + i{\cdot}b/2 = 0.011691 + 0.020250\, i$ and $z = \zeta_{C3}$.

One caution on reading this entry. Its positive real part is *consistent with* $\sigma_\varphi^q = +1$, but it is not a *derivation* of the sign: the entry was reconstructed by imposing $\zeta = \zeta_{C3}$, which already carries the +1 branch, so "Re > 0" here restates the chosen branch rather than establishing it. The arrow runs the other way — the owed task is to show that H_cl returns $\mathrm{Re}(H\_even^q)_{23} > 0$ *without* using ζ_C3 (or CKM) as input. This is the reconstruction-versus-derivation distinction made explicit in §4.6.

Conditional Derivation Theorem. The candidate $\zeta\_C3 = (b/\sqrt{3})\, e^{(5\pi i/6)}$ — equivalently the entry target above — is derived from H_W **if and only if** the projected role-even quark block satisfies all five conditions:

1. **Channel.** The leading shared committed-quark curvature is 2↔3, with no leading 1↔2 and no direct 1↔3 shared curvature. *(Selects the channel; §4.1.)*
2. **C_3 orbit.** The 2↔3 curvature is a democratic component of a three-branch Killing-orthonormal C_3 orbit — democratic at the support level, with the visible 2↔3 channel selected by Cabibbo protection (§4.3.1). *(With 3, fixes the modulus; §4.3.)*
3. **Norm inheritance.** The total inherited 2↔3 committed transport norm is b, giving $|\zeta| = b/\sqrt{3}$. *(Fixes the modulus; §4.3.)*
4. **Phase grammar.** The phase obeys $h\_q^2 = C \cdot O\_q$ with $C = e^{(i\pi)}$, $O\_q = e^{(2\pi i/3)}$, fixing arg ζ to $\{5\pi/6, 11\pi/6\}$ (the phase is λ-independent, since z and ζ differ only by real positive λ). *(Fixes the phase magnitude; §4.4.)*
5. **Positive branch.** $\mathrm{Re}(H\_even^q)_{23} > 0$, equivalently $\sigma\_\varphi^q = +1$, selecting $5\pi/6$ over $11\pi/6$. *(Fixes the sign; §4.4 lemma.)*

Conditions 1–3 deliver the modulus, condition 4 the phase magnitude, and condition 5 the sign. **Conditions 4 and 5 are independent.** Condition 5 alone fixes only the half-plane $\mathrm{Im}\,\zeta > 0$, not the angle $5\pi/6$: an entry with Re > 0 but arg ≠ π/3 still gives the correct branch while missing the phase. Dropping condition 4 therefore yields $|\zeta| = b/\sqrt{3}$ with the right branch but the wrong angle — the half-plane, not the point. This is exactly the §11 split: condition 5 failing is orientation incompleteness (8a), or CP-even failure if the entry is real (8b); condition 4 failing is a separate phase-magnitude failure.

None of the five conditions uses observed CKM data: the entry target is what the projection must return *before* any comparison with measurement (the data-side constraint $\sigma\_\varphi^q = +1$ is quarantined in §9). If H_cl returns the single entry $(b/(\lambda\, \varepsilon\_W \sqrt{3}))\, e^{(i\pi/3)}$, the CKM curvature is derived; if it returns anything else, the failed factor — channel, modulus, phase magnitude, or sign — localises the revision.

4.5.1 Hamiltonian-side minimal-lift lemma (conditional)

The branch sign (condition 5) can be reduced from an owed data-fit to a CKM-free admissibility statement. The move is to work on the Hermitian side. With $h = (H\_even^q)_{23}$ and $z = i\, \varepsilon\_W\, h$ ($\varepsilon\_W > 0$), note that the i-rotation carries $(i)^2 = e^{(i\pi)} = C$, the complement reversal. The generator-side phase grammar $z^2 \propto C \cdot O\_q$ (§4.4) therefore becomes, on the Hamiltonian side,

$$\mathbf{h^2 \propto O\_q = e^{(2\pi i/3)},}$$

the C_3 generation orientation alone — the i-rotation has already supplied C. The two square-root roots are

$h / |h| \in \{ e^{i\pi/3}, e^{i4\pi/3} \}$.

Lemma. If admissibility realises the *minimal continuous lift* of the C_3 orientation from the identity *on the Hermitian generator H_cl* — the root of smaller principal-phase magnitude, π/3 rather than 2π/3 — then

$h / |h| = e^{i\pi/3}$, $\mathrm{Re}(H_{even}^q)_{23} > 0$, $\sigma_\varphi^q = +1$, $\arg z = \pi/2 + \pi/3 = 5\pi/6$,

selected with no CKM input.

Load-bearing assumption (the lemma's entire content). The minimal-lift rule is not branch-neutral, so *where* it is applied is the whole claim. Applied to the anti-Hermitian z it selects the opposite branch: the roots of z are at arg 5π/6 and 11π/6 (i.e. 150° and −30°), whose minimal-magnitude lift is −30°, giving $\sigma_\varphi^q = -1$. The lemma asserts that the admissibility principle acts on the Hermitian H_cl, *before* the i-rotation — equivalently, that h, not z, is the minimal-reconfiguration generator. This is plausible (the admissibility programme casts H_cl as the minimal generator of admissible state reconfiguration) but it is owed, not proven: the admissibility results establish that a generator of this type is forced, yet do not compute its entries or fix this selection. A skeptic is entitled to observe that the h-side is also the side that yields the observed branch, so the justification must come from admissibility, not from the outcome.

Grade. The lemma upgrades $\sigma_\varphi^q = +1$ from an owed sign (returned by holonomy, or otherwise checked against data in §9) to a *conditional derivation* resting on two CKM-free conditions: **(A)** $h^2 \propto O_q$ — the role-even 2↔3 entry is the square-root transport of the C_3 orientation; and **(B)** minimal Hermitian lift — admissibility selects the shortest continuous square-root lift, acting on H_cl. Given A and B, the branch is derived without data. Neither A nor B is established here, so this is a branch-orientation theorem, not closure; and it does not touch the amplitude, which remains owed through the C_3 norm-sharing of §4.3.

4.6 Three levels: reconstruction, conditional derivation, first-principles calculation

The entry target of §4.5 must not be over-read, and the paper states plainly where it sits. H_cl is not itself a measured object, and in this paper it is not yet a computed one: the substrate admissibility/free-energy functional is not specified explicitly here, so the projected block cannot be evaluated numerically from first principles. The CKM curvature target therefore occupies one of three distinct levels, and this paper claims only the first two.

Level 1 — target reconstruction (achieved). From the measured CKM structure one can infer what the projected Hamiltonian entry *must* be if VERSF is correct,

$(H_{even}^q)_{23} = (b/(\lambda\, \varepsilon_W \sqrt{3}))\, e^{i\pi/3}$ (in the $\lambda = 1$ convention, $b/(\varepsilon_W\sqrt{3})$).

This is a reconstruction, not a first-principles result. It states a necessary condition — "if the framework holds, the projection must contain this entry" — and is, by construction, consistent with the data it was read from. It is not evidence that H_cl actually returns the entry.

Level 2 — conditional derivation (achieved). Granting the five structural conditions of §4.5 — the 2↔3 channel, the democratic Killing-orthonormal C_3 orbit, the inherited norm b, the complement-reversal half-transport grammar, and the positive branch — the same entry (equivalently ζ_{C3}) follows *independently of the data*. This is stronger than reconstruction: it shows the target is the natural output of a small set of structural rules rather than a free fit. But it remains conditional on the projected Hamiltonian actually possessing that structure.

Level 3 — first-principles calculation (owed). A genuine derivation requires the full chain, with no CKM input at any step,

$$F[x] \rightarrow G = \partial^2 F/\partial x_i \partial x_j \rightarrow H_{cl} = G|su(8) \rightarrow P_W \rightarrow H_W = P_W\, H_{cl}\, P_W,$$

and only then the evaluation of $\mathrm{Re}(H_{even}^q)_{23}$ and its sign. The admissibility/Hamiltonian results justify that a reversible generator of this type should exist, but they do not fix its entries, spectrum, or interaction content; those require the explicit functional, the symmetry data, and the boundary conditions. Until this chain is carried out, $\sigma_\varphi^q = +1$ is reconstructed from CKM (§9), not derived from F.

Stated plainly, to pre-empt the obvious objection: **H_cl is not measured — CKM is measured, and this paper reconstructs the H_cl projection that would explain it.** The owed step is to derive that same projection from the substrate functional without using CKM as input. The distinction is kept explicit precisely so the Level-1 reconstruction is never mistaken for the Level-3 calculation; a critic is entitled to say "you have not calculated H_cl, you have inferred the entry from the observed triangle," and the correct reply is that the paper claims exactly that and no more.

---

5. Neutrino-Sector Audit

In the lepton doublet the charged lepton is closure-committed while the neutrino is only weakly committed:

$$\gamma_e \simeq 1,\ \gamma_\nu = \varepsilon \ll 1.$$

The audit target is the weak-commitment scale-shape form

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

This form is powerful because T_ν and M_ν share eigenvectors for **every** $\varepsilon > 0$: if $M_\nu v = m v$, then $T_\nu v = (D_0 + \varepsilon m) v$. The small neutrino scale ε therefore suppresses *eigenvalues*, not the *mixing frame*. This is the structural reason tiny neutrino masses can coexist with large PMNS mixing. The first neutrino audit condition is

$$\mathbf{P\_\nu\ H\_W\ P\_\nu \to D_0\ I + \varepsilon\ M\_\nu.}$$

If the residual eigenframe collapses as $\varepsilon \to 0$, the weak-commitment explanation fails.

5.1 Target residual kernel

The benchmark weak-commitment residual kernel has the form (or its $\mu\leftrightarrow\tau$ reflected branch)

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

with symmetric core

$$M_0 = \begin{pmatrix} r\_e & A & A \\ A & r\_s & B \\ A & B & r\_s \end{pmatrix}.$$

Rotating to $\boldsymbol{\nu_+ = (\nu\_\mu + \nu\_\tau)/\sqrt{2}}$ and $\boldsymbol{\nu_- = (\nu\_\mu - \nu\_\tau)/\sqrt{2}}$ produces the solar block

$$\begin{pmatrix} r\_e & \sqrt{2}\cdot A \\ \sqrt{2}\cdot A & r\_s + B \end{pmatrix},$$

so that

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

Defining $\boldsymbol{\rho\_\odot = (r\_s + B - r\_e)/A}$ gives the compact relation

$$\mathbf{\tan 2\theta_{12} = 2\sqrt{2} / \rho\_\odot,}$$

with audit target $\boldsymbol{\rho\_\odot = \sqrt{(3/2)}}$.

5.2 Solar-ratio audit

The solar ratio is a root-support claim: the projected residual cloud compares a threefold completion support against a twofold coherent non-electron pair support, giving

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

The audit must check whether the projected weak-commitment support actually carries this 3:2 root comparison. If it does not, the solar-ratio theorem is not realised.

5.3 Pair-ratio audit

The residual support space is expected to split as

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

with target pair ratio **$B/A = 6/5$**. The denominator 5 is dim R. The numerator 6 requires exactly **one** additional return-continuation channel for the internal μ–τ pair relative to the electron-to-pair link:

$$\mathbf{B/A = (5 + 1)/5 = 6/5.}$$

This is a sharp audit: zero extra channels, two extra channels, or a root-normalised alternative would each change the result.

5.4 Weak-neutrality audit

The target diagonal μ–τ split is **$\delta = 0 + O(\beta^2)$**. This is not a convenience. It says the *first* atmospheric breaking should not be a diagonal μ–τ stiffness difference; it should appear as off-diagonal electron-attachment leakage. The audit condition is therefore that first-order diagonal μ–τ splitting is **suppressed**. If an unsuppressed first-order δ appears, the PMNS kernel becomes easier to fit but less predictive, and the architecture loses one of its sharpest constraints.

5.5 Neutrino phase audit: σ_φ^ν

The neutrino phase is again a branch question. Complement-half-transport gives

$$\mathbf{h_\nu^2 = C \cdot O_\nu,\ C = e^{i\pi},\ O_\nu = e^{i\pi/2} \Longrightarrow h_\nu^2 = e^{3\pi i/2},}$$

with antipodal roots

$$\mathbf{h_\nu = e^{3\pi i/4} \text{ or } h_\nu = e^{7\pi i/4}.}$$

Define **$\sigma_\varphi^\nu = \mathrm{sgn}(\mathrm{Im}\, h_\nu)$**. Then

$$\mathbf{\sigma_\varphi^\nu > 0 \to \psi = 3\pi/4,\ \sigma_\varphi^\nu < 0 \to \psi = 7\pi/4.}$$

As in the quark case, the branch sign is owed: it must come from the projected weak-commitment holonomy, not from a naming convention.

---

6. Leakage Support Trace

The weakest historical target was the reactor leakage amplitude **$\beta = \sqrt{3}/20$**. This paper strengthens its status by moving the denominator 20 into a support-trace audit. The numerator $\sqrt{3}$ is the root-normalised threefold leakage source; the denominator must be the trace of the projected leakage support:

$\beta = \sqrt{3}$ / Tr_support(Π_leak), with target Tr_support(Π_leak) = (K − 2) × 4.

With **K = 7** this gives **5 × 4 = 20**. The interpretation is precise:

- The K = 7 closure carrier supplies seven admissibility directions.
- Two are fixed by boundary/readout anchoring: one by the committed charged-lepton frame, one by the weak-doublet complement/readout basis. The live residual support is therefore **dim R = K − 2 = 5**.
- The leakage channel carries a fourfold weak-attachment orientation fibre, **dim Q = 4 = 2 × 2**, where the first factor is the weak-neutral quadrature and the second is an **owed independent closure-orientation binary** (e.g. a closure two-cycle / holonomy binary). The second factor must *not* be read as a Hermitian forward/return orientation: Hermiticity fixes the return entry once the forward entry is set, so it adds no new support slot, and counting it would double-count. Identifying the genuine second binary from the closure structure is part of the owed support-trace calculation.

The audit target is the factorisation

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

For the argument to be non-circular, numerator and denominator must come from the *same* projector Π_leak, not from two separately-asserted supports. They do: the numerator $\sqrt{3}$ is the root of the threefold (C_3) multiplicity of the leakage source *within* Π_leak — the same $\sqrt{3}$ of the C_3 branch sharing in §4.3 — while the denominator is Tr_support(Π_leak) = dim R · dim Q. So $\beta = \sqrt{3}$ / Tr_support(Π_leak) is one projector's source-multiplicity over its own support trace, not a ratio of two independent guesses.

If the factorisation is verified, then **$\beta = \sqrt{3}/20$**. If it is not, β remains a benchmark rather than a derived output. This is the central support-trace test of the paper — and it is the single most load-bearing un-pinned factorisation in the manuscript, since the entire reactor sector hangs on the owed second binary of dim Q. **Identifying that closure-orientation binary from the closure structure (and confirming it is not the Hermitian return) is therefore the explicit top priority of the next calculation** — ahead of the quark-sector commutators, because every other leakage quantity inherits from it.

---

7. Atmospheric Octant Audit

The leakage magnitude β controls the *size* of the atmospheric departure from maximality; it does not select the *branch*. The branch is selected by which electron-attachment channel receives the leading leakage. Define

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

and the octant sign

$$\sigma_W = \mathrm{sgn}(A_{e\tau} - A_{e\mu}) = \mathrm{sgn}(\|P_e H_W P_\tau\|^2 - \|P_e H_W P_\mu\|^2).$$

The branch rule is

$\sigma_W > 0$ → electron–tau leakage → upper-octant branch, $\sigma_W < 0$ → electron–muon leakage → lower-octant branch, $\sigma_W = 0$ → no first-order octant selection.

The atmospheric octant is thus no longer a loose ambiguity; it is the sign of one projected-Hamiltonian quantity.

---

8. First Projection-Audit Theorem

Theorem (Minimal weak-doublet projection audit). Assume $H_W = P_W H_{cl} P_W$ is well-defined, charge-blind in its mass/readout component, and role-gapped at leading flavour-frame order. Then the first calculational audit of the VERSF Standard Model flavour programme reduces to the finite output list

$$\Omega_q,\ \Omega_0^{(23)},\ z,\ M_\nu,\ \Pi_{leak},\ \mathrm{Tr\_support}(\Pi_{leak}),\ \sigma_W,\ \sigma_\varphi^q,\ \sigma_\varphi^\nu,$$

with success conditions

$\Omega_{odd} \to \Omega_q$ with $a \gg b \gg c$, Ω_{even} contains the leading $2\leftrightarrow3$ shared curvature $\Omega_0^{(23)}(z)$, $[\Omega_0(z), \Omega_q]_{13} = -a\,z$ (exact), $\Delta_{13} = \tfrac{1}{2}\, a\, \zeta$ (= $(\lambda/2)\, a\, z$; leading order, coefficient exact), $|\zeta| = b/\sqrt{3}$, $T_\nu = D_0 I + \varepsilon M_\nu$, $\rho_\odot = \sqrt{3/2}$, $B/A = 6/5$, $\delta = 0 + O(\beta^2)$, $\mathrm{Tr\_support}(\Pi_{leak}) = 20$, $\beta = \sqrt{3}/20$,

and the three signs **σ_W, σ_φ^q, σ_φ^ν** nonzero and computed from the projection.

Proof sketch. The weak-role decomposition is exact from the Pauli partial-trace coefficients of §2. The CKM residue is exact from the single matrix-element sum of §4.2, and the BCH coefficient ($\lambda/2$, i.e. ¼ in the half-split convention) is exact from the BCH expansion, giving the convention-independent $\Delta_{13} = \tfrac{1}{2}\, a\, \zeta$ (verified in Appendix C; the numerics confirm implementation, not the algebra). The neutrino eigenframe result is exact from the shared-eigenvector property of $D_0 I + \varepsilon M_\nu$. The amplitude $|\zeta| = b/\sqrt{3}$ holds only if the projected role-even curvature is a democratic, Killing-orthonormal C_3 component. The leakage amplitude $\beta = \sqrt{3}/20$ holds only if Π_{leak} factors as $R \otimes Q$ with dimensions 5 and 4. The octant and CP phases require the discrete signs σ_W, σ_φ^q, σ_φ^ν. The listed outputs are therefore the minimal finite audit of the weak-doublet closure programme. ■

---

9. Post-Audit Consistency Check

The audit rules fix the required projected-Hamiltonian outputs **without allowing new flavour matrices to be inserted by hand** (Rule 1); some numerical targets are inherited from earlier sector audits and remain first-principles outputs still owed by the substrate calculation, not fresh fits to the data below. As an after-the-fact check only — one that does not feed back into the construction — it is worth recording where the support targets land when compared with measurement.

The solar target $\rho_\odot = \sqrt{3/2}$ gives

$$\tan 2\theta_{12} = 2\sqrt{2} / \sqrt{3/2} \Longrightarrow \sin^2\theta_{12} \approx 0.30.$$

As an after-the-fact comparison with current global fits, $\sin^2\theta_{12}$ lies near the measured solar value; the specific comparison should be updated against the latest NuFIT / PDG release at the time of submission. The leakage target $\beta = \sqrt{3}/20 \approx 0.087$ is of the right order for the reactor-angle suppression.

The CKM phase branch admits the same after-the-fact constraint. The earlier *CKM Curvature Residue* audit shows that the J + |V_td| repair selects a curvature phase near arg $z \approx 145.5°$ (a band roughly [133°, 154°]), with the C_3 value $5\pi/6 = 150°$ at its upper edge. Because the measured triangle apex lies in the upper half-plane ($\bar{\eta} > 0$, $J > 0$), only the positive-imaginary branch repairs the triangle; the antipode $11\pi/6 = 330°$ reverses the CP orientation. A successful projection must therefore return **σ_φ^q = +1**. This is recorded strictly as a consistency requirement on the eventual output, computed *from* observed |V_td|, J, and $\bar{\eta}$: it does not derive σ_φ^q, it does not enter the construction of H_cl, and under Rule 1 it may not. The derivation obligation remains exactly as stated in §4.4 — to compute sgn Re(H_even^q)$_{23}$ from the projected Hamiltonian and check that it comes out positive.

These agreements are reported as **consistency**, not as inputs: the audit forbids tuning toward them, and a clean numerical landing near the data is evidence that the named targets are non-vacuous, not that they have been fitted.

---

10. Calculation Protocol

The next microscopic calculation must **not** start from CKM or PMNS data. It must start from H_cl and the projectors P_C, P_R, P_Y, in this sequence:

1. Build the projected block **$H_W = P_W H_{cl} P_W$**.
2. Form the anti-Hermitian generator **$\Omega_W = i \varepsilon_W H_W$**.
3. Take the role partial traces: **$\Omega_{even} = ½ Tr_{role}(\Omega_W)$, $\Omega_{odd} = ½ Tr_{role}(\Omega_W \tau_3)$, $\Omega_{1,2} = ½ Tr_{role}(\Omega_W \tau_{1,2})$**.
4. Check role-gap admissibility: **$\|\Omega_{mix}\| \ll \|\Omega_{even}\| + \|\Omega_{odd}\|$**.
5. Test the role-odd hierarchy: **$\Omega_{odd} \to \Omega_q$, $a \gg b \gg c$**.
6. Test the first role-even shared curvature: *Ω_{even} contains $\Omega_0^{(23)} = z E_{23} - z E_{32}$**.
7. Compute the projected orbit norm and phase branch: **$\zeta = \lambda z$?= $b/\sqrt{3}$ ($|\zeta|$), $\sigma_\varphi^q = sgn(Im\ z)$**.
8. Extract the neutrino weak-commitment block: **T_ν ?= $D_0 I + \varepsilon M_\nu$**.
9. Read the residual kernel ratios: **$\rho_\odot$, B/A, δ, $\beta \cdot e^{(i\psi)}$**.
10. Compute the leakage support projector and trace: **Π_{leak}, $Tr_{support}(\Pi_{leak})$**.
11. Compute the attachment and phase signs: **σ_W, σ_φ^q, σ_φ^ν**.

Only after this sequence may the resulting CKM and PMNS predictions be compared with measured data.

---

11. Falsification Conditions

The audit fails **locally** if any of the following occurs.

1. **Projection failure** — H_W cannot be defined without using observed flavour matrices.
2. **Charge failure** — the mass/readout block violates charge-blindness.
3. **Role failure** — Ω_{mix} is not gapped or subleading.
4. **Quark-hierarchy failure** — Ω_{odd} does not return a Cabibbo-dominated hierarchy.
5. **Curvature-channel failure** — the leading role-even shared curvature is not 2↔3.
6. **Triangle failure** — the CKM correction is not $\Delta_{13} = ½ a \zeta$ (equivalently $(\lambda/2) a z$ in the chosen convention).
7. **Amplitude failure** — the projected role-even orbit does not give $|\zeta| = b/\sqrt{3}$ or a controlled deformation of it.
8. **Quark-branch outcomes.** The branch sign is separate from the magnitude, and four outcomes are distinct:
    - **8a — orientation incompleteness:** σ_φ^q is not computed from the projection. The triangle is sized but not oriented — incompleteness, not falsification (the octant-before-σ_W status).
    - **8b — branch mismatch:** the projection returns $\sigma_\varphi^q = -1$ (arg $\zeta = 11\pi/6$). This reverses the CKM CP orientation and *falsifies* the branch rule against the observed handedness.
    - **8c — branch collapse:** $\sigma_\varphi^q = 0$ (or $(H_{even}^q)_{23}$ real). No first-order CP handedness is supplied; the residue may have a magnitude but encodes no oriented area — failure of the mechanism as a CP derivation.
    - ($\sigma_\varphi^q = +1$ is the success case: positive branch, arg $\zeta = 5\pi/6$.) Outcomes 8a–8c collapse to this success case *if* the §4.5.1 minimal-Hermitian-lift admissibility axiom holds; absent that axiom, the branch remains owed and 8a applies.

9. **Weak-commitment failure** — the neutrino block is not $D_0 I + \varepsilon M_\nu$.
10. **Solar-support failure** — the residual support does not give $\rho_\odot = \sqrt{3/2}$.
11. **Pair-continuation failure** — the continuation count does not give $B/A = 6/5$.
12. **Weak-neutrality failure** — first-order diagonal μ–τ splitting appears unsuppressed.
13. **Neutrino-branch failure** — σ_φ^ν is undefined or not returned by the projection.
14. **Leakage-factorisation failure** — Π_{leak} does not factor as $R \otimes Q$ with dim R = 5, dim Q = 4.
15. **Reactor-amplitude failure** — $\mathrm{Tr\_support}(\Pi_{leak}) \neq 20$.
16. **Octant failure** — $\sigma_W = 0$ or selects the branch opposite to stable data.
17. **Methodological failure** — any target block is inserted because it improves agreement with data.

Condition 17 is essential: a numerically successful fit is not a derivation if the target structures are supplied by hand.

---

12. What This Paper Adds

This paper does not add a new flavour mechanism — that is the point. It adds a stricter accounting system for the mechanism already proposed.

The previous theorem said the weak-doublet block must take the admissible shape $\mathbf{\Omega_W = \Omega_{even} \otimes I + \Omega_{odd} \otimes \tau_3 + \Omega_{mix}}$. This paper says what the next calculation must extract from that shape:

- the committed quark role-odd block Ω_q;
- the 2↔3 role-even shared curvature $\Omega_0(z)$;
- the C_3-normalised amplitude $|\zeta| = b/\sqrt{3}$;
- the quark CP branch σ_φ^q;
- the weak-commitment neutrino residual kernel M_ν;
- the solar and pair support ratios;
- the weak-neutral suppression of first-order diagonal μ–τ splitting;
- the leakage support trace $\mathrm{Tr\_support}(\Pi_{leak}) = 20$;
- the neutrino CP branch σ_φ^ν;
- the atmospheric octant sign σ_W.

The programme is now accountable at the level of named projected-Hamiltonian outputs.

---

13. Conclusion

The VERSF Standard Model flavour programme has reached a useful but demanding stage. The earlier papers identified the closure Hamiltonian H_{cl}, the weak-doublet projection P_W, and the admissible block structure of $H\_W = P\_W H\_cl P\_W$. This paper strengthens the next step by refusing to treat the remaining targets as adjustable parameters; it converts them into a finite projection audit.

In the quark sector, the decisive question is whether $P\_W H\_cl P\_W$ returns a role-even 2↔3 shared curvature. If it does, the missing CKM triangle contribution follows from exact non-commutation with the Cabibbo doorway,

$$[\Omega_0(z), \Omega\_q]_{13} = -a z, \quad \Delta_{13} = \tfrac{1}{2} a \zeta,$$

with amplitude and CP orientation then requiring the projected orbit norm $|\zeta| = b/\sqrt{3}$ and the branch sign $\sigma\_\varphi\hat{}q$.

In the neutrino sector, the decisive question is whether weak commitment gives $\mathbf{T\_\nu = D_0 I + \varepsilon M\_\nu}$. If it does, the small neutrino scale suppresses masses but not the PMNS eigenframe, and the residual kernel must return $\rho\_\odot = \sqrt{(3/2)}$, $B/A = 6/5$, and $\delta = 0 + O(\beta^2)$.

The reactor-angle leakage is now tied to a support trace rather than a loose denominator,

$$\mathrm{Tr\_support}(\Pi\_leak) = (K - 2) \times 4 = 20 \Longrightarrow \beta = \sqrt{3/20},$$

and the atmospheric octant and two CP branches are not conventions but the projected signs σ_W, $\sigma\_\varphi\hat{}q$, $\sigma\_\varphi\hat{}\nu$.

This paper therefore marks a change in the standard of proof. The VERSF flavour programme can no longer rest on architectural plausibility; it must compute the projected-Hamiltonian outputs it has now named. If those outputs are returned, the programme advances toward a genuine Standard Model flavour derivation. If not, the failure is clean: the failed block, sign, norm, or support trace identifies exactly where the closure programme must be revised.

The amplitude and the phase have not, however, reached the same epistemic status, and it is worth stating the split plainly. The amplitude $|\zeta| = b/\sqrt{3}$ is now a conditional theorem: once the C_3-equivariant pre-readout support is established — reducing to the single symmetry condition that the substrate functional is C_3-invariant on the generation cell with a C_3-symmetric background — the root-normalised value follows by norm algebra. One honesty note belongs at this claim, not buried in §4.3: that premise is close to its own conclusion. A substrate C_3-invariant before readout is precisely a substrate that treats the three generations democratically, which is what equal branch weights mean. The reduction is therefore a reformulation of the amplitude claim into its most fundamental and checkable form — a single symmetry property of the explicit F — not a decrease in logical content. Its value is that the property is now sharply true-or-false of F, not that the assumption has been eliminated.

The phase is weaker still. The unsigned phase magnitude (the $|2\pi/3|$ half-transport structure giving $\arg \zeta$ near $5\pi/6$ in modulus) is supplied by the grammar of §4.4. But the *signed* content — which of $5\pi/6$ versus $11\pi/6$, equivalently $\sigma\_\varphi\hat{}q = \pm 1$ — is fixed by the sign of the C_3 orientation ($e^{+2\pi i/3}$ vs $e^{-2\pi i/3}$), a single owed object that sets the magnitude branch and the CP handedness together. The minimal-Hermitian-lift rule (§4.5.1) is the proposed route to that one signed object, and it remains an admissibility condition still to be derived from the substrate dynamics; a successful projection must independently return it. The frontier is therefore two named theorems to prove — **C_3 covariance** for the amplitude and **minimal Hermitian lift** for the signed orientation — with only the unsigned magnitude coming free from the grammar. Neither theorem invokes CKM data, which is the property the whole audit exists to protect.

Appendix A — Minimal Matrix Outputs

The next technical calculation should evaluate or bound the projected blocks

P_even P_W H_cl P_W P_even, P_odd P_W H_cl P_W P_odd, P_e P_W H_cl P_W P_μ, P_e P_W H_cl P_W P_τ,

together with the leakage projector and its support trace **Π_leak, Tr_support(Π_leak)**. The decisive outputs are

Ω_0^(23), ζ_C3, M_ν, Π_leak, Tr_support(Π_leak), β, σ_W, σ_φ^q, σ_φ^ν.

The audit succeeds only if these objects are *returned by* the projected Hamiltonian rather than inserted as target structures.

Appendix B — Compact Audit Scorecard

Output	Pass condition	Failure meaning
Ω_mix	gapped or subleading	no clean weak-role frame
Ω_q	Cabibbo-dominated, a ≫ b ≫ c	quark role-odd split not derived
Ω_0^(23)	leading role-even 2↔3 curvature	CKM curvature channel fails
$(H\_even^q)_{23}$	= (b/(λ ε_W √3)) e^(iπ/3)	CKM curvature not returned by H_cl
$[\Omega_0, \Omega\_q]_{13}$	= −a z	algebraic mechanism absent if channel absent
ζ (= λz)		ζ
σ_φ^q	nonzero branch sign (= +1 if §4.5.1 minimal-lift axiom holds)	CKM CP branch not derived
T_ν	= D_0 I + ε M_ν	weak-commitment PMNS frame fails
ρ_⊙	= √(3/2)	solar support ratio fails
B/A	= 6/5	return-continuation rule fails

Output	Pass condition	Failure meaning
δ	$= 0 + O(\beta^2)$	weak-neutrality suppression fails
Π_leak	= R ⊗ Q (dim 5 ⊗ 4)	leakage support factorisation fails
Tr_support(Π_leak)	= 20	β denominator not derived
β	$= \sqrt{3}/20$	reactor leakage remains benchmarked
σ_φ^ν	nonzero branch sign	PMNS CP branch not derived
σ_W	nonzero attachment sign	octant not selected

Appendix C — Numerical Checks

The two **exact** claims of §4.2 are independent of the substrate evaluation and can be checked directly on generic anti-Hermitian generators. Both are verified for generic parameters — for example $a = 0.225$, $b = 0.041$, $c = 0.0035$, $\varphi = 1.2$, $z = 0.013 + 0.021\,i$ — chosen with no relation to observed flavour data.

C.1 Triangle residue. With $\Omega_0(z) = z\,E_{23} - z^*\,E_{32}$ and Ω_q in the form of §4, the (1,3) entry of $[\Omega_0, \Omega\_q]$ reproduces $-a\,z$ to machine precision ($|\text{difference}| \approx 10^{-17}$). The full commutator is *not* a single-entry matrix: with $c \neq 0$ and complex z it also carries (1,2), (2,1), (3,1) and diagonal (2,2), (3,3) entries. Only the (1,3) triangle element equals $-a\,z$ exactly, which is precisely the element that the BCH correction promotes into the CKM triangle. Because the (1,3) entry reads only $(\Omega\_q)_{12} = a$ and the vanishing first row of Ω_0, it is unchanged when $b = c = 0$ (and for any φ): the residue $-a\,z$ is reproduced identically in the $b = c = 0$ case and at generic large b or c, confirming its independence of the rest of the role-odd block.

C.2 BCH coefficient ¼. Forming $U\_u = \exp((\Omega_0 - \Omega\_q)/2)$, $U\_d = \exp((\Omega_0 + \Omega\_q)/2)$, $V = U\_u^\dagger\,U\_d$, and reading $\Delta_{13} = (\log V)_{13} - (\Omega\_q)_{13}$, a scaled-generator test ($\Omega_0, \Omega\_q \to s\cdot\Omega_0, s\cdot\Omega\_q$) isolates the leading term:

scale s	Δ_{13} (numeric)	¼ a z at scale s	ratio Δ_{13} / (¼ s^2 a z)
1.00	$+7.320\times10^{-4} + 1.183\times10^{-3}\,i$	$+7.313\times10^{-4} + 1.181\times10^{-3}\,i$	1.0012
0.30	$+6.582\times10^{-5} + 1.063\times10^{-4}\,i$	$+6.581\times10^{-5} + 1.063\times10^{-4}\,i$	1.0001
0.10	$+7.313\times10^{-6} + 1.181\times10^{-5}\,i$	$+7.313\times10^{-6} + 1.181\times10^{-5}\,i$	1.0000
0.03	$+6.581\times10^{-7} + 1.063\times10^{-6}\,i$	$+6.581\times10^{-7} + 1.063\times10^{-6}\,i$	1.0000

The ratio approaches 1 as $s \to 0$, confirming $\Delta_{13} = \tfrac{1}{4}\,a\,z$ as the leading contribution, with the residual O(s) departure accounted for by the higher nested commutators. (These checks validate only the data-independent algebra of §4.2; they say nothing about whether the projected H_cl actually returns Ω_0^(23) or the C_3 amplitude, which remain conditional targets and owed outputs.)

C.3 Entry target (Level-1 reconstruction). With $\varepsilon\_W = 1$ and the inherited $b = 81/2000$, the §4.5 target entry $(H\_even^q)_{23} = b/(2\lambda\sqrt{3}) + i\cdot b/(2\lambda)$ reproduces the invariant $\zeta\_C3 = (b/\sqrt{3})\, e^{(5\pi i/6)}$ to machine precision in either convention. At $\lambda = 1$: $(H\_even^q)_{23} = 0.011691 + 0.020250\, i$, $z = i(H\_even^q)_{23} = -0.020250 + 0.011691\, i = \zeta\_C3$. At $\lambda = ½$ (this paper): $(H\_even^q)_{23} = 0.023383 + 0.040500\, i$, $z = -0.040500 + 0.023383\, i = 2\, \zeta\_C3$. The block $H\_even,C3^q$ is Hermitian by construction. This is arithmetic on the reconstructed target (§4.6, Level 1); it does not evaluate H_cl and does not derive the sign of the entry.

C.4 Convention invariance of Δ_{13}. Building $V\_CKM = \exp(-\lambda\Omega_0 + ½\Omega\_q)\cdot\exp(\lambda\Omega_0 + ½\Omega\_q)$ with $z = \zeta\_C3/\lambda$ (so the *physical* curvature $\zeta = \lambda z$ is held fixed) and reading the full residue $\Delta_{13} = (\log V\_CKM)_{13} - (\Omega\_q)_{13}$, the result is numerically identical for $\lambda = 1$ and $\lambda = ½$: full $\Delta_{13} \approx -2.28\times10^{-3} + 1.32\times10^{-3}\, i$ (three significant figures). Beyond that precision the full residue's value is set by the higher nested commutators of the BCH tail, which read the *full* inherited role-odd block Ω_q — its (b, c, φ) entries, not the 1↔3 entry alone; the fourth–fifth digit therefore shifts with those inherited inputs and is not independently pinned here. The leading term, by contrast, is exact and input-light: $½\, a\, \zeta\_C3 = -2.2781\times10^{-3} + 1.3153\times10^{-3}\, i$, depending only on a and ζ. Leading and full differ by ≈ 0.1% (the tail). The essential point is convention invariance: for any fixed inherited (b, c, φ), the full residue is identical in both λ — every order of the expansion is convention-invariant, structurally, because $\lambda\Omega_0(z) = \Omega_0(\zeta)$ (§4.2) makes V_CKM a function of ζ alone. The bare amplitude $z = \zeta/\lambda$ and the entry modulus $b/(\lambda\varepsilon\_W\sqrt{3})$ carry the λ normalisation; the entry phase $\pi/3$ does not.

C.5 C_3 branch orbit (existence and orthonormality). The branch generators $X_{12} = E_{12} - E_{21}$, $X_{23} = E_{23} - E_{32}$, $X_{31} = E_{31} - E_{13}$ have Killing/trace Gram matrix $\langle X\_i, X\_j\rangle\_K = -\mathrm{Tr}(X\_i\, X\_j) = 2\, \delta\_ij$ — mutually orthogonal with equal norm 2. The cyclic generation permutation (1→2→3→1) acts as $P\, X_{12}\, P^T = X_{23}$, $P\, X_{23}\, P^T = X_{31}$, $P\, X_{31}\, P^T = X_{12}$, so the triple is a single C_3 orbit. This establishes claims (A) and (B) of §4.3.1 — the orbit exists and is Killing-orthonormal — at the level of the generation algebra; it does not establish (C), that the projected role-even block populates the orbit democratically, which remains the owed equivariance fact.

C.6 Pre-readout democracy score. Using the normalised Hermitian branch quadratures of C.5 and the democracy score $D\_C_3 = [\Sigma\_ij\, (w\_ij - \bar{w})^2]/\bar{w}^2$ with $\bar{w} = ⅓\, \Sigma\, w\_ij$: a democratic pre-readout block carrying amplitude $b/\sqrt{3}$ in each of the three branches gives $D\_C_3 = 0$, $\Sigma\, w = b^2$, and $\sqrt{(w_{23})} = b/\sqrt{3} = 0.02338$. A lopsided block placing the entire inherited norm b in the 2↔3 branch gives the *same* $\Sigma\, w = b^2$ but $D\_C_3 = 6 \neq 0$. The total norm alone therefore does not certify the amplitude; democratic sharing ($D\_C_3 = 0$) is the discriminating condition, and it is owed from the explicit pre-readout block $K\_even^q = ½\, \mathrm{Tr}\_role(P\_q\, P\_R\, P\_C\, H\_cl\, P\_C\, P\_R\, P\_q)$.

C.7 Substrate-cell democracy (closure Hessian and block covariance). Two checks support §4.3.3–§4.3.4. (a) *Closure Hessian.* With anti-Hermitian branch generators X_{12}, X_{23}, X_{31}, the product $\exp(\varepsilon X_{12})\, \exp(\varepsilon X_{23})\, \exp(\varepsilon X_{31}) - \mathbb{1} = \varepsilon(X_{12} + X_{23} + X_{31}) + O(\varepsilon^2)$ to numerical precision, and the quadratic-form Hessian $\langle X\_i, X\_j\rangle\_K = -\mathrm{Tr}(X\_i\, X\_j)$ is exactly $2\, \delta\_ij$ — the closure penalty is scalar (democratic) on the branch orbit. (b) *Covariance vs. Schur.* A generic C_3-invariant quadratic form on the three branches has eigenvalues [0.9, 0.9, 2.1] — invariant but not scalar, confirming that C_3-invariance of the Hessian does not force democracy. By contrast, a C_3-covariant Hermitian block ($R_3\, K\, R_3^\dagger = K$, R_3 the cyclic generation permutation) has exactly equal off-diagonal branch magnitudes, $|K_{12}| = |K_{23}| = |K_{31}|$. Democracy therefore follows from covariance of the configuration, which is the content of the $[P\_C\ H\_cl\ P\_C, R_3] = 0$ bottleneck — not from a scalar Hessian or an SU(8)-Schur argument (SU(8) acts on the closure fibre, not the generation branches). (c) *Hessian equivariance.* For a C_3-invariant test functional $F[R_3\ x] = F[x]$, the numerical Hessian at a C_3-symmetric background (equal components) commutes with R_3 to within finite-difference precision, while at a generic non-symmetric background it does not — confirming both the C_3 Equivariance Lemma of §4.3.4 and the necessity of its no-spontaneous-breaking premise (b).