The CKM Curvature Residue in VERSF

Computing the Weak-Doublet Commutator, the BCH Correction, and the Jarlskog / |V_td| Repair

Keith Taylor VERSF Theoretical Physics Programme — Standard Model Flavour Series

Working manuscript, v0.2

Successor to *The Electroweak Flavour-Frame Operator in VERSF*, *The Yukawa Operator from Completion-Channel Misalignment in VERSF*, and *Deriving Flavour Mixing from Closure Geometry*.

Summary for the General Reader

The previous electroweak flavour-frame paper made the next Standard Model flavour task very sharp. Quark mixing is not best understood as one family being diagonalised by itself: the CKM matrix is the *mismatch* between two families, the up-type frame and the down-type frame. If those two frames are equal-and-opposite tilts about one shared weak-doublet frame, the leading CKM pattern follows cleanly.

But the same paper named the remaining problem. The leading construction gets the large pattern right yet misses two things: the amount of matter–antimatter asymmetry, measured by the Jarlskog invariant J, and the long side of the CKM unitarity triangle, |V_td|. Ordinary tuning cannot fix this neatly — increasing the direct 1↔3 link changes too many entries at once, and moving the phase helps one quantity while harming another. The right place to look is a *structured* correction: a small curvature residue left over when the common weak-doublet frame does not commute perfectly with the role-odd up/down split.

This paper takes that residue seriously and reduces it to one object,

$\kappa = [\Omega_0, \Omega\_q]$.

Here Ω_q is the relative quark mixing generator already used in the Yukawa paper, and Ω_0 is the common-mode weak-doublet transport that the up and down sectors share. If these commute, the clean CKM exponential is exact. If they do not, a Baker–Campbell–Hausdorff (BCH) correction appears — not arbitrary noise, but a correction of precise shape, and the only natural place to repair the missing CP and triangle structure without destroying the leading CKM entries.

The central proposal is a *minimal* common-mode curvature. Rather than a fully free correction, we test a single complex 2↔3 common-mode rotation whose amplitude is tied to the inherited 2↔3 scale b by a threefold sharing rule, $|z| = b/\sqrt{3}$, and whose phase is the bisector between complement reversal and the C_3 generation-loop holonomy:

$z = (b/\sqrt{3}) \cdot \exp(5\pi i/6)$.

This small twist has a large effect because it is dragged through the wide Cabibbo doorway: the commutator of a 2↔3 common-mode rotation with the 1↔2 Cabibbo rotation produces precisely a 1↔3 curvature — exactly the missing long-triangle term.

Numerically the result is striking. Starting from the clean exponential at $\varphi = 2\pi/3$, the minimal C_3 curvature lifts $|V_{td}|$ from ≈ 0.0061 to ≈ 0.00869 and raises $|J|$ from $\approx 1.84\times10^{-5}$ to $\approx 3.00\times10^{-5}$ (observed $\approx 3.12\times10^{-5}$), while $|V_{us}|$ and $|V_{cb}|$ stay put. This is not a free lunch: the same curvature that builds the triangle and CP structure *spends* $|V_{ub}|$, which the clean exponential had already placed well (≈ 0.00393, near the exclusive value), dropping it to ≈ 0.00356. The trade is strongly favourable — a few-percent cost on $|V_{ub}|$ buys a ≈ 28% gain on $|V_{td}|$, a ≈ 37% gain on J, and an improved ratio — but it does convert $|V_{ub}|$ from the construction's one correct third-generation magnitude into its one liability, and the draft says so plainly.

The draft also makes the predictive content precise. The 2↔3 ansatz carries two real numbers (the amplitude and phase of z), and two observables — J and $|V_{td}|$ — pin them to within current experimental error. That leaves $|V_{ub}|$ and the ratio $|V_{ub}|/|V_{cb}|$ as *forced* outputs, not fitted ones, and they run into a sharp **$|V_{ub}|$–versus–ratio wall**: no admissible curvature reaches both the exclusive $|V_{ub}|$ and the direct ratio at once. So "it works" is not cheap — once the two parameters are set, the remaining predictions are nearly falsifying, and z_{C3} lands at the ratio-favouring end of the wall.

This is not yet a first-principles derivation: the curvature still has to be *forced* by the weak-doublet closure Hamiltonian rather than selected because it works. But the CKM problem has been narrowed from "find some correction" to "derive this one C_3-shaped common-mode rotation, or reject it."

Abstract

The electroweak flavour-frame paper reduced the remaining CKM problem in VERSF to a BCH curvature residue. The clean relative-exponential construction assumes the product-form weak-doublet split

$$U_u = U_0 \cdot \exp(-\Omega_q/2),\quad U_d = U_0 \cdot \exp(+\Omega_q/2),$$

giving $V_{CKM} = U_u^\dagger U_d = \exp(\Omega_q)$. The natural single-exponential doublet construction instead yields *additive* sector frames, $U_u = \exp(\Omega_0 - \Omega_q/2)$, $U_d = \exp(\Omega_0 + \Omega_q/2)$, and hence

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

The entire departure from the clean exponential is therefore controlled by the commutator $\kappa = [\Omega_0, \Omega_q]$. The previous paper named this object QF-2 and left its computation open.

This paper gives the first candidate computation of that residue. We retain the inherited relative generator ($a = 9/40$, $b = 81/2000$, $c = 243/100000$, $\varphi = 2\pi/3$) and test a minimal role-even common-mode curvature concentrated in the 2↔3 completion channel,

$$\Omega_0(z) = z\cdot E_{23} - z^*\cdot E_{32},\quad z = (b/\sqrt{3}) \cdot \exp(5\pi i/6),$$

where E_ij is the elementary matrix unit. The amplitude $b/\sqrt{3}$ distributes the inherited 2↔3 transport strength over the threefold closure loop; the phase $5\pi/6$ is the bisector of complement reversal and the C_3 holonomy. The commutator is then explicit — its dominant off-diagonal term is

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

so the first BCH correction contributes $+\frac{1}{2}\cdot a\cdot z$ to the 1↔3 sector. A small common-mode 2↔3 curvature is converted by the Cabibbo 1↔2 link into a 1↔3 triangle correction, without directly touching the leading 1↔2 entry.

The exact additive-frame matrix $V\_CKM = \exp(-\Omega_0 + \Omega\_q/2)\cdot\exp(\Omega_0 + \Omega\_q/2)$ with this z gives

$|V\_us| = 0.2231$, $|V\_cb| = 0.04054$, $|V\_ub| = 0.00356$, $|V\_td| = 0.00869$, $|J| = 3.00\times10^{-5}$, $|V\_ub|/|V\_cb| = 0.0877$.

Relative to the clean exponential this lifts $|V\_td|$ ($0.0061 \rightarrow 0.00869$) and $|J|$ ($1.84\times10^{-5} \rightarrow 3.00\times10^{-5}$) while leaving $|V\_us|$ and $|V\_cb|$ stable, at the cost of dropping $|V\_ub|$ from an already-good 0.00393 to 0.00356. We further show that (i) the ansatz carries two real parameters (the amplitude and phase of z), which the two observables J and $|V\_td|$ pin to within current experimental error — so the apparent ($|z|$, phase) "band" is the data error bar mapped into parameter space, not model slack, and collapses to a point as the tolerance shrinks; and (ii) with the two parameters thereby fixed, $|V\_ub|$ and the ratio $|V\_ub|/|V\_cb|$ are *forced* outputs that run into a structural **|V_ub|–ratio wall**: hitting the exclusive $|V\_ub| \approx 0.00382$ forces the ratio up to ≈ 0.098, hitting the direct ratio 0.083 forces $|V\_ub|$ down to ≈ 0.00353, and no admissible curvature — not even the data-preferred best fit — reaches both. The wall survives loosening the J and $|V\_td|$ windows to ±18%/±9%. A pure 2↔3 curvature cannot satisfy both; z_C3 sits at the ratio-favouring end.

The result is therefore not a completed CKM derivation but a sharply constrained curvature-residue candidate with no continuously fitted amplitude once the C_3 ansatz for z is accepted. The next first-principles task is exact: derive or reject $z = (b/\sqrt{3})\cdot\exp(5\pi i/6)$ from the weak-doublet closure Hamiltonian.

Table of Contents

---

0. Predictive-Content Ledger

This paper addresses one target.

QF-2: compute the CKM curvature residue created by the failure of the common weak-doublet frame to commute with the role-odd quark split, and test whether that residue repairs the Jarlskog invariant and $|V_{td}|$ without spoiling the leading CKM magnitudes.

Object	Construction	Status in this paper
Relative CKM generator	Ω_q with entries a, b, c and phase φ	inherited from the Yukawa / electroweak-frame papers
Baseline CKM matrix	$V_0 = \exp(\Omega_q)$	inherited leading result
Common-mode generator	Ω_0	the object QF-2 must compute
Curvature commutator	$\kappa = [\Omega_0, \Omega_q]$	sole BCH control object (exact)
BCH residue	$\Delta = -\tfrac{1}{2}\kappa + \tfrac{1}{6}[\Omega_0, \kappa] + \cdots$	derived from additive sector frames
Minimal curvature ansatz	$\Omega_0(z) = z\cdot E_{23} - z^*\cdot E_{32}$	proposed candidate
C_3 value of z	$z = (b/\sqrt{3})\cdot\exp(5\pi i/6)$	proposed, geometry-motivated, **not derived**
Dominant residue term	$\kappa_{13} = -a\cdot z$	exact in the minimal ansatz

Object	Construction	Status in this paper
Curved CKM matrix	$V_{curv} = \exp(-\Omega_0 + \Omega_q/2)\cdot\exp(\Omega_0 + \Omega_q/2)$	exact audit
Main repair	$\lvert V_{td}\rvert$ and $\lvert J\rvert$ lifted to observed scale	successful at candidate level
Parameter count	2 params ($\lvert z\rvert$, phase) vs 2 observables (J, $\lvert V_{td}\rvert$)	**pinned to a point, not free, not derived**
Forced outputs	$\lvert V_{ub}\rvert$, $\lvert V_{ub}\rvert/\lvert V_{cb}\rvert$	**predictions; hit the $\lvert V_{ub}\rvert$–ratio wall**
Cost	$\lvert V_{ub}\rvert$ 0.00393 → 0.00356	redistribution: spends the one good 3rd-gen magnitude
Proof obligation	derive z (esp. the phase) from closure geometry	open

The paper separates three layers.

1. **Exact algebra.** If Ω_0 exists, the BCH residue is controlled by $\kappa = [\Omega_0, \Omega_q]$; for the minimal $2\leftrightarrow3$ ansatz, $\kappa_{13} = -a\cdot z$ exactly.
2. **Candidate construction.** The C_3 curvature $z = (b/\sqrt{3})\cdot\exp(5\pi i/6)$ repairs the CP and $|V_{td}|$ deficits at the exact-matrix level; its two parameters are pinned (not freed) by the two observables, and the resulting $|V_{ub}|$ / ratio are forced predictions that hit a structural wall.
3. **Physical derivation, still owed.** VERSF closure geometry must derive that z rather than letting the numerical repair select it.

The discipline is unchanged: **an ansatz that works is not yet a derivation.** This paper is valuable only insofar as it turns QF-2 into a concrete, checkable closure-geometry target.

1. The Inherited QF-2 Problem

The electroweak flavour-frame paper showed that the clean CKM construction rests on product-form QF-1,

$$U_u = U_0 \cdot \exp(-\Omega_q/2),\quad U_d = U_0 \cdot \exp(+\Omega_q/2),$$

so the common frame cancels and $V_{CKM} = U_u^\dagger U_d = \exp(\Omega_q)$. With $\varphi = 2\pi/3$ and inherited entries $a = 9/40$, $b = 81/2000$, $c = 243/100000$, this gives leading magnitudes near $|V_{us}| = 0.223$, $|V_{cb}| = 0.040$, $|V_{ub}| = 0.0039$.

But the same paper made two things clear. First, some of those successes are *inherited*: $|V_{us}|$ is essentially the Cabibbo input 9/40 read back out, and $|V_{cb}|$ is a consistency check because b is assembled from the transport survival factor but is still close to an input scale. The genuinely hard outputs are the loop-sensitive quantities |V_ub|, |V_td|, |V_ub|/|V_cb|, and J. Second, the clean exponential underpowers the triangle and CP geometry: at φ = 2π/3 it gives

|V_td| ≈ 0.0061, |J| ≈ 1.84×10^{-5},

against observed targets near 0.0086 and 3.12×10^{-5}. The leading hierarchy is right; the triangle is incomplete.

The missing correction was named QF-2: not an arbitrary patch to CKM, but the curvature residue created if the common doublet frame fails to commute with the role-odd split. This paper attempts its first concrete computation.

---

2. The Clean Relative Exponential and Its Missing Outputs

The inherited quark relative generator is anti-Hermitian, Ω_q† = −Ω_q:

$$\Omega\_q = \begin{bmatrix} 0 & a & c\cdot\exp(+i\varphi) \\ -a & 0 & b \\ -c\cdot\exp(-i\varphi) & -b & 0 \end{bmatrix}$$

with a = 9/40 = 0.225, b = 81/2000 = 0.0405, c = 243/100000 = 0.00243, φ = 2π/3. The clean relative exponential is V_0 = exp(Ω_q), with exact magnitudes:

quantity	clean exponential	comment
|V_us|	0.22307	essentially inherited from a = 9/40
|V_cb|	0.04028	consistency check near b
|V_ub|	0.00393	close to exclusive central value
|V_td|	0.00611	**too low**
|V_ub|/|V_cb|	0.0976	high vs direct ratio 0.083 ± 0.004
|J|	1.84×10^{-5}	**too low**

The goal is not to adjust a, b, c by hand, but to compute the structured residual that appears when the common weak-doublet frame is allowed to curve against Ω_q.

---

3. Product Split, Additive Split, and the BCH Residue

There are two ways to describe the sector frames. The **product split**,

U_u = U_0 · exp(−Ω_q/2), U_d = U_0 · exp(+Ω_q/2), gives exact cancellation, $U_u^\dagger U_d = \exp(+\Omega_q/2)\cdot U_0^\dagger U_0\cdot\exp(+\Omega_q/2) = \exp(\Omega_q)$. This is QF-1. The **additive split**,

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

is the natural form from a single weak-doublet generator with role-even part Ω_0 and role-odd part $\Omega_q/2$. In general it does *not* reduce to $\exp(\Omega_q)$, because exponentials of non-commuting generators do not factor.

Writing $A = -\Omega_0 + \Omega_q/2$ and $B = \Omega_0 + \Omega_q/2$, so $U_u^\dagger U_d = \exp(A)\cdot\exp(B)$, the BCH expansion

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

with $A + B = \Omega_q$ and $[A, B] = -[\Omega_0, \Omega_q]$ gives, defining $\kappa = [\Omega_0, \Omega_q]$,

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

So the residue is

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

the exact formal target of QF-2. (Both displayed coefficients are standard BCH and were confirmed by direct matrix computation: subtracting $-\tfrac{1}{2}\kappa$ leaves an $O(t^3)$ residual, and subtracting $\tfrac{1}{6}[\Omega_0, \kappa]$ as well leaves $O(t^4)$, where t scales the generators.)

---

4. The Commutator Algebra

The CKM deficits are concentrated in the 1↔3 / triangle sector. A correction to the direct c entry is too crude — it inflates $|V_{ub}|$ and $|V_{td}|$ together and can worsen the ratio. A commutator correction is more structured, because it creates new entries through *paths*.

The most economical common-mode curvature is a role-even 2↔3 generator,

$$\Omega_0(z) = \begin{bmatrix} 0 & 0 & 0 \\ 0 & 0 & z \\ 0 & -z^* & 0 \end{bmatrix}, \qquad \Omega_0^\dagger = -\Omega_0.$$

This common curvature does not directly touch the Cabibbo 1↔2 entry, but its commutator with Ω_q does. A direct calculation gives

$$[\Omega_0, \Omega_q]_{13} = -a\cdot z,\quad [\Omega_0, \Omega_q]_{12} = c\cdot\exp(i\varphi)\cdot z^*,\quad [\Omega_0, \Omega_q]_{23} = 0,$$

with small diagonal phase terms $\propto b(z^* - z)$. The dominant term is the 1↔3 entry, $\kappa_{13} = -a\cdot z$, so the first BCH correction contributes

$\Delta_{13} = -\tfrac{1}{2}\kappa_{13} = \tfrac{1}{2}\cdot a\cdot z$.

The mechanism in one line:

1↔2 Cabibbo transport × 2↔3 common curvature → 1↔3 triangle correction.

This is exactly what the clean exponential was missing. It is not a new direct 1↔3 knob; it is curvature generated by composing the existing Cabibbo doorway with a common 2↔3 weak-doublet twist. (All three off-diagonal commutator entries and the diagonal $\propto b(z^* - z)$ were verified to machine precision.)

---

5. A Minimal Common-Mode Curvature Ansatz

What should z be? A fully free complex z carries two parameters and risks becoming a fit. The aim here is a geometrically constrained candidate.

The inherited 2↔3 relative scale is $b = 81/2000$. The common-mode curvature should be *smaller* than b, since it is not the relative 2↔3 mixing itself but the residual curvature of the shared doublet frame. A threefold sharing rule distributes b over the C_3 generation loop:

$|z| = b/\sqrt{3}$.

The phase should not be the CKM holonomy $\varphi = 2\pi/3$, because Ω_0 is common-mode, not role-odd. It sits halfway between the C_3 loop orientation ($2\pi/3$) and the complement reversal (π) associated with moving from one doublet reading to the other. The bisector phase is

$\theta_0 = (\pi + 2\pi/3)/2 = 5\pi/6$.

The minimal C_3 common-mode curvature ansatz is therefore

$z_{C3} = (b/\sqrt{3}) \cdot \exp(5\pi i/6) = -0.02025 + 0.01169\, i$, $|z_{C3}| = 0.02338$, $\|\Omega_0\|_F = \sqrt{2}\cdot|z_{C3}| = 0.03307$.

This is small compared with the Cabibbo scale $a = 0.225$ and comparable to the 2↔3 scale $b = 0.0405$ — small enough not to spoil the leading entries, large enough that $\tfrac{1}{2}\cdot a\cdot z$ is comparable to c. Indeed

$|\tfrac{1}{2}\cdot a\cdot z_{C3}| = 0.00263$,

the same size as the inherited direct 1↔3 entry $c = 0.00243$. That is why the correction strongly affects $|V_{td}|$ and J while leaving $|V_{us}|$ and $|V_{cb}|$ nearly stable.

Status. z_C3 is not derived here. It is a no-continuous-parameter candidate once b and the C_3/complement phases are accepted. But the two inputs are not equally secure, and they should not be bundled as one object "z_C3" when the derivation is written. The **amplitude** $|z| = b/\sqrt{3}$ has a genuine story — the 2↔3 transport scale distributed over the threefold loop — and is the kind of quantity a closure-Hamiltonian calculation tends to deliver. The **phase** $\theta_0 = 5\pi/6$ is thinner: it is the arithmetic mean of π and $2\pi/3$, and one can fairly ask why those two angles, why their mean, and why not bisect against φ itself. The phase is therefore the weaker of the two assumptions and the real target of the derivation. §7 sharpens this empirically: the value the data prefers is ≈ 145°, so $5\pi/6 = 150°$ sits about 5° high, at the upper edge of the allowed window. (The proximity of the data-preferred phase to $4\pi/5 = 144°$ is a coincidence to resist, not a rationale to reach for: $4\pi/5$ is a C_5 angle and this programme is C_3-based.) The proof obligation is to derive the amplitude and — especially — the phase from the weak-doublet closure Hamiltonian.

---

6. Exact CKM Audit with the C_3 Curvature Residue

Using the additive-frame matrix $V\_curv = \exp(-\Omega_0 + \Omega\_q/2)\cdot\exp(\Omega_0 + \Omega\_q/2)$ with $\Omega_0 = \Omega_0(z\_C3)$, the exact matrix exponential gives:

quantity	**clean exp(Ω_q)**	**C_3 curvature**	**target / comment**
|V_us|	0.22307	0.22307	stable; inherited Cabibbo scale
|V_cb|	0.04028	0.04054	stable; near target
|V_ub|	0.00393	0.00356	lowered; below exclusive central value
|V_td|	0.00611	0.00869	**repaired**
|V_ts|	0.04001	0.03975	stable
|V_ub|/|V_cb|	0.0976	0.0877	improved; still above 0.083
|J|	1.84×10^{-5}	3.00×10^{-5}	**near target**

The full magnitude matrix is

$$|V\_curv| = \begin{bmatrix} 0.974796 & 0.223071 & 0.003557 \\ 0.222930 & 0.973991 & 0.040538 \\ 0.008690 & 0.039755 & 0.999172 \end{bmatrix} \quad \text{(unitary to machine precision)}$$

The repair is exactly the one QF-2 wanted: the residue raises |V_td| and J while preserving the large Cabibbo and 2↔3 entries, and pulls |V_ub|/|V_cb| from 0.0976 toward 0.0877, closer to the direct 0.083 ± 0.004. But it should be read as a *redistribution*, not a free gain. At the clean exponential |V_ub| = 0.00393 was already close to the exclusive central value 0.00382 (+2.9%); the curvature spends that accuracy, dropping |V_ub| to 0.00356 (−6.8%) — low by ≈ 7% against exclusive and ≈ 17% against the inclusive ≈ 0.0043. So the construction converts its one already-correct third-generation magnitude into its one liability. The trade is strongly net-positive —

three of the four hard outputs (|V_td|, J, ratio) improve and only |V_ub| degrades, and the total triangle/CP error falls sharply — but it is a real and specific cost, and §7 shows it is structural to the 2↔3 channel rather than an artefact of the particular z.

---

7. Parameter Counting, the Repair "Band", and the |V_ub|–Ratio Wall

Two questions decide whether the §6 result is a coincidence or a mechanism: how tightly is z_C3 determined, and is the low |V_ub| an unavoidable cost? Both turn on a parameter count.

The construction carries exactly two real parameters. A 2↔3 curvature $\Omega_0(z)$ is one complex number — the amplitude |z| and phase of z. Against it stand two third-generation observables that the clean exponential misses, J and |V_td|. Two observables, two parameters: generically a *point*, not a region. Scanning the (|z|, phase) plane and demanding both |V_td| ≈ 0.0086 and J ≈ 3.12×10^{-5} does select a finite blob, but the blob is the experimental error bar mapped into parameter space — it shrinks toward a point as the tolerance tightens:

acceptance tolerance	|z| window	phase window
±10%	0.0153–0.0304	120°–161°
±5%	0.0191–0.0266	134°–154°
±2%	0.0214–0.0243	141°–149°
±1%	0.0221–0.0236	143°–147°

The area falls roughly quadratically with the tolerance, confirming a genuine point solution (non-degenerate Jacobian), not a flat direction. The exact-target best fit sits at **|z| = 0.0228, phase = 145.5°**. So the apparent "robustness to ±15% in |z| and ±10° in phase" is not model slack — it is a statement about how precisely J and |V_td| are currently measured. This cuts both ways and the draft states it symmetrically: the two parameters are *not* free (the data pins them), but neither are they *derived* (the data, not closure geometry, pins them). The bisector value 5π/6 = 150° lies ≈ 5° above the data-preferred 145.5°, at the upper edge of the window — consistent, but displaced, which is why §5 flags the phase as the input the closure derivation most needs to justify.

With the two parameters fixed, |V_ub| and the ratio are forced — and they hit a wall. This is what makes "it works" non-trivial rather than cheap. Once J and |V_td| set (|z|, phase), nothing remains free: |V_ub| and |V_ub|/|V_cb| are predictions. Moving along the J + |V_td| solution, those two predictions trade against each other:

point	|z|	phase	|V_ub|	|V_ub|/|V_cb|
best |V_ub|	0.0216	139°	0.00397	0.098
data best fit	0.0228	145.5°	0.00374	0.092
z_C3	0.0234	150°	0.00356	0.0877
best ratio	0.0243	151°	0.00353	0.0870

A pure 2↔3 curvature can reach the exclusive |V_ub| ≈ 0.00382 (up to ≈ 0.00397, near the inclusive branch) — but only by driving the ratio to ≈ 0.098; and it can reach the direct ratio ≈ 0.087 — but only by forcing |V_ub| down to ≈ 0.0035. It **cannot** hit both (0.00382, 0.083) at once. Even the data-preferred best fit leaves |V_ub| at −2.1% and the ratio at 0.092, so the wall is not a consequence of the C_3 phase choice; it stands at the optimum. It also survives loosening the J and |V_td| windows to ±18%/±9% — still no admissible point reaches both, with the closest joint approach leaving |V_ub| ≈ 7% low while the ratio sits at 0.088. The wall is therefore structural, not a tolerance artefact, and it is a near-falsifying prediction rather than a fitted nuisance:

Once J and |V_td| fix the two parameters of the 2↔3 ansatz, |V_ub| and |V_ub|/|V_cb| are forced onto a single curve whose ends are the exclusive |V_ub| and the direct ratio; no point on it reaches both.

Whether this is resolved by the experimental |V_ub| / ratio determinations settling (the direct Λ_b ratio already runs low relative to the individual exclusive extractions), or by a small *second* curvature component outside the 2↔3 channel, is the precise question §8 and §11 leave open. One thing it does *not* permit is an independent dial on J: raising |z| to push J from 3.00 toward 3.12×10^{-5} also pushes |V_td| above target and the ratio further from 0.083 (while |V_ub| rises only marginally), so J cannot be tuned alone — it moves you along the wall, not off it.

---

8. What the Residue Fixes, and What It Does Not

The C_3 common-mode curvature improves the right things for the right reason. It fixes what the clean exponential could not: |V_td| is lifted from too small to the observed scale; |J| from ≈ 40% low to within a few percent; |V_us| and |V_cb| stay stable; and |V_ub|/|V_cb| moves in the right direction.

It does not close CKM. The residual, made precise in §7, is the |V_ub|–ratio trade-off: the curvature builds the correct long side and CP area but ties the short-side magnitude to the ratio, and the pure 2↔3 channel cannot place both on their measured values. There are three readings, in increasing order of demand on the construction:

1. **The C_3 residue is essentially right, and the experimental |V_ub| / ratio will settle toward the ratio-favouring corner.** Then the construction becomes strong with no further structure.
2. **The C_3 residue is the leading term, but a small second curvature component (a 1↔2 or 1↔3 common-mode piece) is needed to slide |V_ub| up the band without undoing |V_td| and J.** Then QF-2 is not closed but strongly guided — and §7 already bounds where that component must act.
3. **The C_3 ansatz is not the true common-mode curvature.** Then the repair is an existence proof, not a derivation.

This paper takes the third reading as the honest default until the weak-doublet closure Hamiltonian is computed. What §7 adds is that readings 1 and 2 are now *quantitative*: the trade-off curve says exactly how much room exists and where a second component would have to enter.

9. Why the 2↔3 Common-Mode Channel Is the Right First Candidate

The 2↔3 choice is not arbitrary.

First, the leading 1↔2 entry is already inherited as the Cabibbo scale; a common-mode correction in the 1↔2 channel would move the one entry the construction should protect. Second, the deficits are third-generation quantities — |V_td|, J, |V_ub|/|V_cb| — so a 2↔3 curvature lives where the problem is. Third, the commutator of a 2↔3 common-mode curvature with the 1↔2 Cabibbo link produces a 1↔3 residue, [2↔3, 1↔2] → 1↔3, the standard closure of a triangular transport loop; the missing CKM quantities are triangle quantities, so the first correction should be a triangle commutator, not an edge rescaling. Fourth, the amplitude b/√3 is the simplest no-new-scale value once b is inherited (the 2↔3 scale distributed over the threefold loop), and 5π/6 is the simplest bisector between complement reversal and C_3 orientation — though, as §5 and §7 stress, the phase is the softer of the two motivations and the one the data would nudge downward by ≈ 5°.

So the ansatz is the most economical C_3-compatible common-mode curvature that can do the job. But economy is not derivation — the field calculation must still force it.

10. Relation to QF-1 and the Electroweak Flavour-Frame Paper

The electroweak frame paper separated two questions. **QF-1:** why does the leading split have product form, so common-mode transport cancels? **QF-2:** if QF-1 is not exact, what residue remains, and can it repair CP and |V_td|?

This paper addresses QF-2 first. It does not derive QF-1; it shows what a *controlled* failure of QF-1 looks like and identifies a minimal C_3-shaped common-mode curvature that repairs the missing outputs. This does not weaken QF-1 — it clarifies it. The clean exponential remains the leading approximation; the curvature residue is the next order:

leading CKM hierarchy: Ω_q • CP / triangle residue: [Ω_0, Ω_q] • first-principles demand: derive both from the weak-doublet closure operator.

The CKM sector now has a two-level architecture: the **relative split** produces Cabibbo and the leading hierarchy, and the **common-mode curvature** produces the CP area and the long unitarity-triangle side. That is more satisfying than forcing all CKM information into Ω_q alone.

11. What Is Proved, Conditional, and Open

11.1 Proven here, given definitions.

- The BCH residue for additive up/down frames is controlled by $\kappa = [\Omega_0, \Omega_q]$ (coefficients $-\frac{1}{2}$ and $\frac{1}{6}$ verified to third order).
- For a pure 2↔3 common-mode curvature $\Omega_0(z)$, the dominant residue term is $\kappa_{13} = -a \cdot z$, with $\kappa_{12} = c \cdot \exp(i\varphi) \cdot z^*$, $\kappa_{23} = 0$, diagonal $\propto b(z^* - z)$ (all verified exactly).
- The first BCH correction contributes $\Delta_{13} = \frac{1}{2} \cdot a \cdot z$.
- With $z_{C3} = (b/\sqrt{3}) \cdot \exp(5\pi i/6)$, the exact additive-frame matrix lifts $|V_{td}|$ to 0.00869 and $|J|$ to 3.00×10^{-5} while keeping $|V_{us}|$ and $|V_{cb}|$ stable.
- The ansatz carries two real parameters (amplitude and phase of z); the two observables J and $|V_{td}|$ pin them to a point ($|z| = 0.0228$, phase = 145.5° at best fit), to within current experimental error — the apparent ($|z|$, phase) band is the data error bar, and shrinks ~quadratically with the tolerance.
- With those two parameters fixed, $|V_{ub}|$ and $|V_{ub}|/|V_{cb}|$ are forced outputs lying on one curve whose ends are ($|V_{ub}| \approx 0.00397$, ratio ≈ 0.098) and ($|V_{ub}| \approx 0.00353$, ratio ≈ 0.087); no point — including the data best fit — reaches both the exclusive $|V_{ub}|$ and the direct ratio, and the wall survives loosening the windows to ±18%/±9%.

11.2 New structural content. The triangular commutator mechanism — Cabibbo 1↔2 transport composed with common 2↔3 curvature produces the missing 1↔3 triangle residue — turns QF-2 from a vague correction into a concrete closure loop: the missing CKM area is not patched by hand but produced by a commutator of two already-meaningful structures. The matched parameter count (two observables pin the two parameters) and the $|V_{ub}|$–ratio wall make the candidate predictive and falsifiable in a sharp, quantitative way rather than a free fit.

11.3 Conditional premises.

- The baseline Ω_q entries a, b, c and phase φ are inherited.
- The common-mode curvature is dominated by the 2↔3 channel.
- The C_3 sharing rule fixes $|z| = b/\sqrt{3}$.
- The common-mode phase is $\theta_0 = 5\pi/6$.
- Higher common-mode channels are absent or subleading.

11.4 Still open.

- Derive z_{C3} from the weak-doublet closure Hamiltonian — and in particular the **phase** $\theta_0 = 5\pi/6$, which is the weaker of the two inputs (the amplitude $b/\sqrt{3}$ has the threefold-loop story; the phase is an arithmetic bisector sitting $\approx 5°$ off the data-preferred 145.5°). The phase is the priority target.
- Decide whether a small second curvature component is required to slide $|V_{ub}|$ up the repair band while holding $|V_{td}|$ and J — §7 bounds where it must act.
- Derive the C_3 holonomy $\varphi = 2\pi/3$ at source rather than selecting it from the generation-loop picture.

- Connect the curvature residue to the full Yukawa eigenvalue programme, including charm and bottom.
- Confirm that the same weak-doublet closure operator is compatible with the PMNS weak-commitment regime.

Honest status: a candidate CKM curvature-residue reconstruction, not a completed first-principles derivation. The calculation is strong because one C_3-shaped curvature repairs exactly the quantities QF-2 was meant to repair, with its two parameters pinned by the two observables and the leftover |V_ub| / ratio standing as forced, near-falsifying predictions; it remains conditional because the curvature itself — above all its phase — is not yet forced by closure geometry.

---

12. Falsification Conditions

12.1 Wrong common-mode channel. If the weak-doublet closure Hamiltonian gives no dominant 2↔3 common-mode curvature, the minimal ansatz fails.

12.2 Wrong amplitude. If the common-mode amplitude is not near b/√3 — outside the band $|z| \in [0.019, 0.026]$ — the repair either underpowers J and |V_td| or spoils the leading entries.

12.3 Wrong phase. If the common-mode phase is outside ≈ [133°, 154°], the commutator lengthens the wrong side of the triangle or fails to raise J.

12.4 Too much common-mode curvature. If $\|\Omega_0\|$ is much larger than ≈ 0.03–0.04, Cabibbo and 2↔3 stability are endangered.

12.5 No |V_ub| / ratio resolution. If neither the experimental determinations settle toward the band nor an admissible second curvature component can move |V_ub| up the trade-off curve while holding |V_td|, J, and the ratio, the minimal residue is incomplete. (This is the sharpened §7 form of the old "no |V_ub| recovery" falsifier.)

12.6 Generic fitting. If the repair requires a fully free six-parameter Ω_0, the result collapses into a fit and stops being a VERSF derivation. The value of the present ansatz is precisely that it reduces Ω_0 to one C_3-determined complex number.

---

13. Conclusion

The electroweak flavour-frame paper reduced the unfinished CKM problem to one commutator, $\kappa = [\Omega_0, \Omega\_q]$. This paper gives the first concrete candidate for that commutator. A small role-even 2↔3 common-mode curvature, $\Omega_0(z\_C3)$ with $z\_C3 = (b/\sqrt{3})\cdot\exp(5\pi i/6)$, has the exact algebraic property the residue needs: its commutator with the Cabibbo part of Ω_q generates a 1↔3 correction,

$$\Delta_{13} = \tfrac{1}{2}\cdot a\cdot z_{C3},$$

naturally of the same size as the inherited direct 1↔3 entry. The exact matrix computation then lifts the two quantities the clean exponential missed,

$$|V_{td}|: 0.0061 \rightarrow 0.00869, \quad |J|: 1.84\times10^{-5} \rightarrow 3.00\times10^{-5},$$

while leaving $|V_{us}|$ and $|V_{cb}|$ stable, and dropping $|V_{ub}|$ from an already-good 0.00393 to 0.00356 — a redistribution that buys the triangle and CP structure at a favourable but real cost. The ansatz's two parameters are pinned, not freed, by the two observables J and $|V_{td}|$, which leaves $|V_{ub}|$ and $|V_{ub}|/|V_{cb}|$ as forced predictions. Those run into a sharp, stress-tested wall: the exclusive $|V_{ub}|$ and the direct ratio sit at opposite ends of one curve, beyond the simultaneous reach of any pure 2↔3 curvature.

The result is conditional, but it is the right kind of conditional. It does not say "there exists some correction." It says: if VERSF weak-doublet closure supplies this specific C_3 common-mode curvature, the CKM sector gains the missing CP and triangle structure; if closure geometry does not supply it, the repair fails cleanly. The next calculation is therefore exact —

derive or reject $z_{C3} = (b/\sqrt{3})\cdot\exp(5\pi i/6)$ from the weak-doublet closure Hamiltonian.

If that derivation succeeds, the CKM architecture becomes much stronger: the leading hierarchy from the role-odd relative split Ω_q, the CP and long-triangle residue from the common-mode curvature commutator $[\Omega_0, \Omega_q]$. That would move the VERSF flavour programme another step toward an actual Standard Model Yukawa derivation.

References

VERSF flavour architecture

- Keith Taylor, *The Electroweak Flavour-Frame Operator in VERSF.* Weak-doublet frame splitting; product-vs-additive split; QF-1 and QF-2; weak-commitment PMNS frame theorem.
- 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 and confinement programme

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

VERSF transport foundations

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

Numerical note. All CKM matrices in this draft were computed by direct complex matrix exponentiation. The clean baseline is $V_0 = \exp(\Omega_q)$; the curved matrix is $V_{curv} = \exp(-\Omega_0 + \Omega_q/2)\cdot\exp(\Omega_0 + \Omega_q/2)$ with $\Omega_0 = \Omega_0(z_{C3})$, $\|\Omega_0(z_{C3})\|_F = 0.03307$. The commutator entries ($\kappa_{13} = -a\cdot z$, $\kappa_{12} = c\cdot\exp(i\varphi)\cdot z^*$, $\kappa_{23} = 0$, diagonal $\propto b(z^* - z)$) and the BCH coefficients ($-\tfrac{1}{2}$, $\tfrac{1}{6}$) were verified to machine precision. The ($|z|$, phase) acceptance region for $J + |V_{td}|$ was mapped by exact-matrix grid scan and shown to shrink ~quadratically with the tolerance toward the best-fit point $|z| = 0.0228$, phase = 145.5°, confirming a non-degenerate two-on-two pin; the $|V_{ub}|$–ratio wall (no point reaching both exclusive $|V_{ub}|$ and the direct ratio) was confirmed to persist under windows loosened to ±18%/±9%. Empirical comparison values follow PDG 2024 (Navas et al., Phys. Rev. D 110, 030001), including the exclusive/inclusive $|V_{ub}|$ split and $|V_{ub}|/|V_{cb}| = 0.083 \pm 0.004$. The Jarlskog magnitude is reported as $|J|$; its sign is fixed by the orientation of the $1\to2\to3$ transport loop and is not used as an independent observable.*