The Conditional Derivation of Charged Flavour and the Unique Lowest-Order Quadratic CP Axle in VERSF

One action-derived Yukawa tensor, one symmetry-forced curvature structure, and the remaining same-action admission test

Keith Taylor VERSF Theoretical Physics Programme — Standard Model Flavour Series
Designation: CKP-2.4 / RRHF-FD1 — main paper Companion: *Provenance, Audits and Retractions for CKP-2.4* (all chronology, ledgers and superseded gradings)

General reader summary

The question. Quarks come in three families and can change from one into another. The likelihoods of those changes form a table of nine numbers, measured very accurately over decades. In standard particle physics that table is not explained — it is put in by hand.

What this paper achieves. It calculates that table from the framework's own rules. One frozen calculation produces the relative pattern of charged-particle masses — strictly, the ratios from which the mass hierarchy follows once the physical scale and running are supplied — and, from how the two kinds of quark sit relative to each other, a single mixing table. No measured mass, mixing angle or matter–antimatter quantity was accessible to the frozen execution core while it ran. The wider framework was developed with knowledge of flavour phenomenology, so the claim is **execution blindness rather than historical blindness**; the rules the calculation uses were themselves worked out within the framework rather than chosen to fit. **This is a genuine conditional action-level derivation**, carried out twice by separately implemented calculations agreeing to fourteen decimal places. Its remaining conditionality is specific rather than vague: parts of the inherited rule-set have their own unfinished provenance, and the absolute physical scale has not been independently determined.

How well it does. The table gets the broad structure right — the large mixing between the first two families, and roughly the right scale for the second and third. It fails on two entries, both in the part governing the matter–antimatter difference, each around a third to a half too small. On the usual scale, where single digits mean good agreement, it scores about 291.

A wrong answer of this kind is worth more than it sounds. A framework loose enough to accommodate anything would have quietly produced the right numbers and taught us nothing. This one produces a **definite** answer that can be checked and found wanting, which is what makes the rest of the investigation possible.

A second, separate result. The framework contains a six-point ring with something circulating around it, and the three families attached to the same object. A direct link from the circulation to the families is impossible — opposite points cancel exactly — as is every alternate link after it. The first that survives is squaring, and at that level it is the only one, delivering three exactly equal shares. **This is a permanent mathematical result.** It stands whatever happens to the rest, and it identifies the unique shape the correction can take at the lowest admissible order.

Why the favourable correction is not presently derived. Running the unique correction at one particular but unreturned direction improves the score dramatically, from about 291 to about 4 on the same five-quantity comparison. No continuous coefficient is fitted — but the direction is not derived, and it was tested with the empirical shortfall already known. It is a powerful diagnostic rather than a prediction. Whether the framework picks that direction has been the central question, and the answer so far is that it does not pick any. The most recent work shows the local structure is exactly even-handed between all directions: it can carry a direction but cannot choose one. So the favourable value is not derived, and neither are the alternatives that were briefly thought to be.

What has been established along the way. Within the declared quark carrier, symmetry fixes one local connection shape, up to its overall strength — which closes off a whole class of ambiguity. The present recording machinery is nevertheless unable to make that shape physically visible. And the obstacle is now precisely located rather than vague: the particular recording machinery tested for this correction cannot see family structure at all, so the information it would need is absent from that machinery rather than hidden in an unfinished calculation — the wider framework does return a family-resolved mass pattern. That is a sharper and more useful thing to know than a suspicion that something, somewhere, was incomplete.

Where it stands. A conditional derivation of the quark mass structure and one mixing table: yes. An exact result on the form of the possible correction: yes. A physical derivation of that correction's strength and direction: not yet. A successful account of the measured mixing pattern: not yet. A complete derivation from the deepest principles: no.

The investigation of the correction has done something useful even though it has not succeeded. It has established, with unusual precision, exactly what further structure the framework would need — and that is a considerably better position than not knowing what was missing.

CENTRAL RESULT — One frozen representation-resolved history functional conditionally returns the complete charged Yukawa tensor and one baseline CKM frame at two regulators, without target access during execution. Independently, a parity theorem forbids every odd-power descent from the sixfold source to the three-family carrier, while D_6 symmetry fixes a unique quadratic descent at the lowest admissible order with exact threefold democracy. The baseline CKM matrix is empirically poor. The quadratic axle is not part of the retained conditional derivation, because its physical coefficient, full internal installation, branch readout and phase remain unreturned: the completion jet is

component-preserving, but the local metric is exactly phase-isotropic, so no physical phase is presently selected.

Contents

§Final — Conditional charged-flavour derivation theorem, status block and decisive standard

Appendices

- **A** Numerical ledger — A.0 controlling derivation; A.0.1 Stage I certificate; A.0.2 admissible-power tower
- **B** Source-space envelopes · **C** Six-branch phase census · **D** Phase sensitivity around the diagnostic direction
- **E** Verification checklist · **F** What would move the boundary quickest
- **References**

Companion — *Provenance, Audits and Retractions for CKP-2.4* carries the inherited-constraint ledger and grading rules (§C0–§C3), the extension's transport algebra (§C4–§C13), the boundary and installation contract (§C20–§C29), the complete premise, falsifier and debt ledgers (§C30–§C32), the curvature-extension numerical ledger, and **Appendix G** — the invariant parameter budget and the full chronology of executions, audits and retractions.

---

§A.0 Is this a derivation? — the direct answer

This is the first question a reader asks, and the paper should answer it rather than leave it to be assembled from the status block. **Yes, with the object stated precisely.**

Conditional charged-flavour derivation — YES. Under the frozen premise package the parent history functional is differentiated and returns $T\_ch = Y\_u \oplus Y\_d \oplus Y\_e$ and hence $V\_CKM = U\_u^\dagger U\_d$, at two regulators, through independently executed bosonic and CAR/Fock routes agreeing at 1.94×10^{-14}. No measured mass, CKM entry, mixing angle or CP quantity is accessible to the execution core. **This is not "given a hand-built matrix, here is its CKM output"** — the representation, history and reduction architecture was developed within the corpus rather than introduced as a fitted model.

Exact lowest-order quadratic-axle derivation — YES. $\mathrm{Hom}(E_1,E_2) = 0$, every odd-power descent forbidden, and $\dim \mathrm{Hom}\_D_6(\mathrm{Sym}^2E_1,E_2) = 1$, with exact threefold democracy of the normalised image.

Physical derivation of the curvature correction — NOT YET. The **local gauge-compatible Q-sector intertwiner form is derived** (AXD-2), but a nonzero c_ax, the **complete full-carrier lift**, P_vis^curv, ρ_mult^physical and θ_phys are not returned from one completed primitive action.

Successful derivation of the observed CKM pattern — NO. The controlling conditional return has $\chi^2_5 \approx 290.7$ and badly underproduces |V_td| and J. **That is an empirical failure of a definite returned matrix, not a failure to return one** — a distinction the paper insists on throughout.

Unconditional first-principles Standard Model derivation — NO. Primitive admission still requires unique selection of the completed commitment-maintenance action, primitive derivation of the charged readout, independent physical-scale selection, the neutral/PMNS and global strong-phase sectors, and the continuum scheme and pole bridge.

§A.0.1 What "conditional" specifically means here

The word should not be left general, because one component of it is concrete and checkable. The execution's inputs include the count package (a, D, η, χ) from which b, c and A_gen follow exactly (§0.2). **The primitive provenance of that package is itself open** — a is a conditional geometric count with the typed overlap operator absent, D rests on a capacity-balance postulate, η on a democratic weighting postulate, and χ's physical weight on an isotropy assumption. That is debt I1, and it is a specific part of what "conditional" carries rather than a vague gesture at depth.

The correct one-sentence statement of the paper's claim:

This paper derives the charged-flavour tensor and one baseline CKM matrix at conditional action level, and separately derives the exact lowest-order form and threefold democracy of a possible quadratic correction — but does not derive that correction's physical coefficient, readout or phase, and the matrix it does return is empirically poor.

§A. Level statement — what this paper derives

This paper gives the flavour-sector chapter of the conditional VERSF Standard Model derivation. Its controlling object is the representation-resolved history functional RRHF-1. It does not present two routes to quark mixing. It presents **one action-derived charged-flavour return plus one unadmitted candidate refinement**, along the chain

$\Gamma\_RRHF^{(n)} \longrightarrow T\_ch^{(n)}=Y\_u^{(n)} \oplus Y\_d^{(n)} \oplus Y\_e^{(n)} \longrightarrow (U\_u^{(n)},U\_d^{(n)}) \longrightarrow V\_CKM^{(n)}=U\_u^{(n)\dagger}U\_d^{(n)}$.

The D_6 wheel and quadratic axle are exact mathematical results identifying the unique symmetry-allowed form of a possible role-even curvature correction. They do not enter the derived CKM matrix unless that correction is returned by the two-stage test of §7 — Stage I on the completed H_CMR, Stage II reintegration into RRHF.

Theorem FD-1 — Conditional charged-flavour derivation

Assume the frozen K=7 representation ledger, the regulator-covariant RCMO maintenance operators, the independently frozen FRGM quark transport operator and the functorial exterior-algebra lift used by RRHF-1. Then one representation-resolved parent history functional, **frozen and executed in an environment in which the RRMK arrays, comparison files and observed Standard Model targets were unavailable to the construction code**, returns by blind differentiation:

1. a rank-nine charged completion tensor $T\_ch^{(n)}=Y\_u^{(n)} \oplus Y\_d^{(n)} \oplus Y\_e^{(n)}$;

2. the charged singular spectra at two genuine regulators;
3. agreement between the independently differentiated bosonic and CAR/Fock routes at $\varepsilon\_CY3'^{(max)}=1.937\times10^{-14}$;
4. the left quark frames and therefore $V\_CKM^{(n)}=U\_u^{(n)\dagger}U\_d^{(n)}$;
5. one regulator-stable conditional CKM matrix, quoted in §5.

No observed mass, coupling, mixing angle, CKM entry or CP phase is used to choose a branch or coefficient before the action and its outputs are sealed.

This is execution blindness, not historical blindness. The programme was developed with knowledge of flavour phenomenology, and the action was designed to test whether prior structural returns could reproduce the inherited RRMK boundary. The claim is that the execution core could not access the targets — not that no one involved knew what RRMK or the measured CKM matrix looked like. Stated this way it is both weaker and defensible.

The representation, history and reduction structure this theorem uses is itself derived within the corpus rather than posited, so the result is an action-level derivation and not a model assembled from chosen inputs. **The conditionality is narrower and should not be read as covering the calculation's internal rules:** the most primitive MASD–HMGBC transfer has not been proved to *force* the registered K=7 representation ledger uniquely, and the absolute scale has not been independently selected. It does not claim the physical neutral block, PMNS, the global measure, the physical strong phase, the finite scheme bridge or pole observables.

§B. One-matrix rule

Until the source-clean Stage I forcing test of §7.1 returns Ω_0 from the completed H_CMR, Ω_0 is not part of the derivation.

The programme-level conditional CKM return here is therefore the RRHF/RRMK frame of §5. The curvature-corrected matrix remains a sealed diagnostic consequence of an additional unadmitted block; it may be compared with the controlling matrix and with experiment, but may not be presented as a second prediction or as part of the derivation.

Outcome	Binding consequence
Stage I-A fails	No physical axle is installed; RRHF remains the controlling return
Stage I-A passes	The axle **actually returned by the action** proceeds to Stage II
Stage I-B comparison	Report separately whether $c\_ax = b$, $\rho\_mult^{physical} = 1$ and $\theta\_phys = 5\pi/6$; **disagreement does not prevent Stage II**

There is no outcome in which two physical CKM matrices remain simultaneously admitted **under one premise set**.

§B.0 The controlling return is indexed by premise set

A third execution, RRHF-AX1 (§5.2), installs the axle as a *declared additional premise* rather than deriving it. That does not breach the one-matrix rule, because it is a different premise set — but it obliges the paper to say which set the programme stands behind.

Premise set	Controlling matrix	χ^2_5	Status
RH1 (RRHF alone)		$V^{(RRHF)}$	,
RH1 + AX1 at 150° *(quarantined)*		$V^{(AX1)}$	,
RH1 + AX1 at the provisional P_{23} branch	30° or 270° by chirality	**539 / 1032**	what the component-preserving jet returns *if* $P_{vis}^{curv} = P_{23}$, which is not established — **worse than RH1 alone** (companion §G.6.10)
(Part II curvature, bare generator)		$V^{(curv)}$	,

RH1 remains the controlling conditional return after the extension test. Stage I was run and did not return the physical installation of the axle (§7.0), so the RH1 row holds because **Stage I-A was executed and did not return a complete physical-installation tuple** — a stronger position than holding it because no test had been attempted.

A third row must be shown with the other two: RH1 + AX1 under the provisional P_{23} readout gives $\chi^2_5 = 539$ or 1032 by chirality, worse than RH1 alone. **It is not the programme's derived branch, because $P_{vis}^{curv} = P_{23}$ has not been established.** The provisional P_{23} reading gives $\chi^2_5 = 539$ or 1032, but **no physical branch has been derived**. Quoting only the 150° row would present a favourable quarantined diagnostic without showing the unfavourable consequences of the provisional alternative. **One matrix per premise set; one currently controlling premise set.** Stage I's failure to return an admissible axle leaves RH1 controlling, but **does not independently validate RH1's primitive premises or its empirical accuracy** — it establishes only that the tested correction has not been physically installed. ("Admitted" elsewhere in this paper means primitive physical admission, which is a much stronger claim.) The empirically good matrix costs the dependency premises of Part II, and the paper does not hide that trade by quoting the better number under the weaker assumptions. **What it costs is now stated precisely**: not a fitted continuous coefficient — AX1 has none — but the **physical coefficient c_{ax}, which remains unreturned alongside the phase carrier and readout**. The carrier selects among $\{150^\circ, 30^\circ, 270^\circ\}$ and thereby decides essentially the whole improvement. **The component-preserving character of the jet has been derived, but the physical readout — and therefore the physical phase — has not** (companion §G.6.11). A Stage I-A pass followed by Stage II would replace the RRHF row with **whatever matrix the action returns**. The first two rows collapse **only if** the Stage I-B comparisons also recover the AX1 values.

§B.1 The stakes of the forcing test, and the hazard they create

This was stated before the test was run, and is retained unchanged now that it has been. §6.1 records that the controlling RRHF matrix misses $|V_{td}|$ by -13.5σ and J by -10.1σ, while the curvature-corrected matrix brings both inside 1.5σ. **The stakes must be stated correctly.** The quarantined 150° execution shows that the exact lowest-order axle form is *capable* of repairing the principal empirical deficit. Stage I asks whether the action installs that structure and what values it returns — **it does not ask the action to reproduce the favourable matrix.** The action could return a different coefficient or phase, pass the physical derivation test, and remain empirically poor.

Stage I has since been executed and did not return the physical installation (§7.0). The hazard below is therefore not spent: the hazard now applies to any future calculation attempting the outstanding **Stage I-A installation tuple**, and with the empirical stakes unchanged. That Stage I returned a *negative* on the items that would have closed the gap is itself weak evidence that the sealing discipline is functioning — a test run under this much motivation that still reports non-return is more credible than one that reports a pass.

Two consequences bind the protocol.

The motivation to return a pass is now extreme and publicly known. Every constructor of the lift, projector and phase-carrier convention will know that one quarantined branch of the candidate structure moves the diagnostic score from about 291 to about 5, creating a strong incentive to reproduce that branch. The §7 freezing rules are necessary but no longer sufficient on their own: a later pass must be reported together with this fact, so that reviewers can weigh it.

Knowledge of the gap is unavoidable, so the seal cannot rest on intellectual blindness. Nobody can now unsee §6.1. Validity therefore rests on **procedural independence**: the lift, projector, carrier metric, phase convention, thresholds and regulator rule must be generated from representation and action data alone, preregistered and hashed **before the completed-H_CMR projected closure Hessian is opened**. No component may be revised after seeing either its projection coefficient or its effect on CKM observables.

The strongest implementation uses three artefacts:

1. a **construction package** containing only source-side representation data;
2. a **sealed execution package** holding the Stage I projection;
3. a **final comparison script** that joins them only after both hashes are published.

That is materially harder to game than relying on the constructor not to think about the gap.

§C. The generator is provenance, not a second derivation

The package $a=9/40$, $D=\mathrm{diag}(1,2,4)$, $\eta=3/5$, $\chi=2/5$, $\varphi=2\pi/3$ is the provenance decomposition of the orientation and score structure inside the returned charged tensor, not an alternative route from which a second CKM matrix is constructed. The downstream entries

$b = a\eta^2/(D_3 - D_2) = 81/2000$, $c = \chi a[a\eta^2/(D_3 - D_1)] = 243/100000$, $A\_gen = b/a^2 = 4/5$

are exact algebraic consequences once the upstream conditional ledger is granted. The remaining primitive provenance of a,D,η,χ,φ is recorded in Part II and does not alter the fact that the augmented RRHF action returns the charged tensor blindly under that ledger.

§D. Conditional derivation versus physical admission

The result proved here is a **conditional charged-flavour derivation at action level**. It is **not** an unconditional Standard Model derivation. Physical admission still requires: derivation of the K=7 representation ledger from the primitive H_CMR; same-action installation of the IORC preparation current; the post-commitment orientation and copy-maintenance lock; a primitive derivation of the physical charged readout P_Y; the typed map from the generation character O_q to the role-odd seed holonomy φ; independent selection of the absolute scale ξ — **conditionally supplied** by HCMR-XI-1 at 84.137117μm on the post-copy branch, subject to CPR and to importing ħ; the physical neutral mass block and PMNS; the global measure, anomaly line and strong phase; the finite native-to-continuum scheme bridge; and whole-action regulator convergence after those returns.

These are physical-admission debts. They do not negate the narrower conditional action-level derivation proved here.

Two of them are one debt. The typed map from O_q to φ, and the same-action return of c_ax and θ_phys, are the flavour-sector face of the H_CMR non-uniqueness recorded at programme level (§7.0.1.1). They close together with H_CMR or not at all.

Part I — The controlling RRHF charged-flavour derivation

§1. Frozen question and non-insertion protocol

The derivation tests whether a representation-resolved extension can be written inside the finite-regulator parent history functional using only prior VERSF representation, completion, support, record and transport returns, such that blind differentiation returns the charged and gauge tensors previously assembled in RRMK-1.

The core construction is forbidden to load an RRMK charged tensor; use the former RK1 or RK2 coefficient tables; insert a target gauge stiffness or Yukawa matrix; use an observed mass or mixing matrix to choose a branch, coefficient, root or phase; or inspect a CKM comparison before the core output has been sealed.

The core manifest freezes and hashes the parent-action source, the test contract, the K=7 representation ledger, the level-3 and level-4 RCMO maintenance operators, the inherited one-

Bit history metric, the independently frozen FRGM quark transport operator, and the output directory state before post-freeze comparison. The target directory remains empty during construction and execution; the RRMK reference is installed only after the action output has been sealed.

§2. Representation-resolved parent action

Charged carrier and ordered history transfer. $H_{ch}=H_u \oplus H_d \oplus H_e$, $\dim H_{ch}=9$. At regulator level n,

$$T_{ch}^{(n)}=U_{role}^{(n)} D_{score}^{(n)},$$

$$Y_f^{(n)} = \mathcal{E}_{fL}^{\dagger}\, P_Y^{RRHF}\, T_{ch}^{(n)}\, P_Y^{RRHF}\, \mathcal{E}_{fR},\ f = u, d, e.$$

The diagonal score step is generated by the K=7 support, localisation, hinge and maintenance ledger; the orientation step by the representation-overlap orbit and the regulator-covariant record commutator. Write this readout $P_Y^{(RRHF)}$. It and the embeddings belong to the frozen ledger; their primitive uniqueness is a later admission question and they are not selected from the numerical Yukawa output.

Two distinct readouts, not one debt. $P_Y^{(RRHF)}$ is the frozen charged-sector support/readout of the controlling action; $P_{vis}^{(curv)}$ is the unresolved branch-selection readout of the curvature extension (Part II). **$P_Y^{(RRHF)}$ is conditionally frozen but not primitively derived; $P_{vis}^{(curv)}$ is not returned at all.** They are different objects and different debts; the apparent contradiction between Part I using a frozen P_Y and Part II calling P_Y unidentified is an artefact of the shared symbol.

Bosonic route. With $R(T_{ch}^{(n)})$ the real block representation,

$$\Gamma_B^{(n)} = (\kappa_B/2)(\|x_L\|^2 + \|x_R\|^2) + x_L^T R(T_{ch}^{(n)}) x_R,$$

$$\partial^2\Gamma_B^{(n)}/\partial x_L \partial x_R=R(T_{ch}^{(n)}).$$

CAR/Fock route. The fermionic route does not read the bosonic cross-Hessian; it lifts the ordered steps separately to the exterior algebra, $\Gamma_{Fock}(U_{role}^{(n)})\Gamma_{Fock}(D_{score}^{(n)})$, descends the one-particle block and differentiates the CAR bilinear. Functorial composition $\Gamma_{Fock}(U)\Gamma_{Fock}(D)=\Gamma_{Fock}(UD)$ holds at residual 1.557×10^{-17} (level 3) and 5.715×10^{-18} (level 4). **The two returns are therefore genuinely separate differentiations of one ordered history transfer.**

§3. Discrete K=7 score ledger

The parent does not store flavour coefficients as free decimals. It evaluates:

Object	Derivation	Return
primitive interface coupling	2K/[2^K(2K+1)]	7/960
localisation stiffness	(K+1)/3	8/3
colour-triality readout	3×3	9
balanced support	C(6,3)	20
residual support	K-2	5
Cabibbo doorway	9/[2C(6,3)]	9/40
survival fraction		3/5
continuation fraction		2/5
2↔3 entry	$a\eta^2/(4-2)$	81/2000
1↔3 entry	$\chi\, a[a\eta^2/(4-1)]$	243/100000

The quark matching operator is independently frozen and verified generation-scalar, $R\_QCD=0.580788332942I_3$, so it changes the common quark scale without changing the relative left frame at the declared matching order.

§4. Blind charged-tensor return

At level 3 blind differentiation returns

$\sigma(Y\_u)=(7.214053900982\times10^{-6},3.704884912807\times10^{-3},9.725186481825\times10^{-1})$,
$\sigma(Y\_d)=(1.563045011879\times10^{-5},3.127500251071\times10^{-4},1.620450741891\times10^{-2})$,
$\sigma(Y\_e)=(2.990274737560\times10^{-6},6.196467859710\times10^{-4},1.070191327928\times10^{-2})$.

Maximum full-tensor difference from the post-freeze reference: 7.068×10^{-14} (level 3), 7.070×10^{-14} (level 4).

Level	bosonic derivative residual	CAR derivative residual	ε_CY3'^(max)
3	4.816×10^{-16}	1.111×10^{-16}	1.934×10^{-14}
4	4.350×10^{-16}	3.332×10^{-16}	1.937×10^{-14}

Minimum bosonic parent-Hessian eigenvalue $\lambda\_min=1.027481351817>0$. **The frozen augmented parent returns a positive rank-nine charged tensor at two regulators.**

§5. CKM as a direct tensor consequence

With $U\_u^{(n)\dagger}Y\_u^{(n)}Y\_u^{(n)\dagger}U\_u^{(n)}=D\_u^{((n)2)}$ and likewise for d, the mixing matrix is not an additional input:

$V\_CKM^{(n)}=U\_u^{(n)\dagger}U\_d^{(n)}$.

At level 4 the controlling conditional modulus is

$$|V\_CKM^{(RRHF)}|=\begin{bmatrix} 0.974818 & 0.222964 & 0.004198 \\ 0.222926 & 0.973988 & 0.040639 \\ 0.005895 & 0.040427 & 0.999165 \end{bmatrix},$$

$|J\_RRHF^{((4))}|=1.86032934918\times10^{-5}$,

with unitarity residual 2.22×10^{-16} on the full complex matrix. The level-3 frame gives $|J\_RRHF^{((3))}|=1.910812315633\times10^{-5}$.

Test	Level-3 to level-4 change	Gate
up / down / charged-lepton singular spectra	4.131280×10^{-4}	pass
CKM modulus	1.246729×10^{-3}	pass
relative Jarlskog	2.641964×10^{-2}	pass
CY3-prime maximum	1.937×10^{-14}	pass
minimum Hessian eigenvalue	1.027481351817	positive

The sign of J is not claimed as a physical prediction; the exact conjugate branch reverses it while preserving the modulus matrix and |J|.

Independent checks on the quoted matrix. The quoted moduli are unitary to $\approx7\times10^{-7}$ row-wise, consistent with six-decimal rounding; the standard-parameterisation |J| implied by those moduli is 1.86114×10^{-5}, agreeing with the quoted |J| to 4.3×10^{-4} relative — again rounding. The implied CP phase is $\delta=30.15^\circ$.

Relation to the bare baseline. The shared trunk $e^{(\Omega\_q)}$ has $|J|=1.8435\times10^{-5}$. The level-4 RRHF frame sits **+0.91%** from it and the level-3 frame **+3.65%**, with maximum modulus difference 4.14×10^{-4} from the bare baseline. **The controlling return is therefore very close to the bare inherited generator**; the curvature correction, by contrast, moves the same baseline by **+62.56%**.

Theorem FD-2 — One controlling conditional CKM return

Under the RRHF premise set, one frozen parent action returns the charged Yukawa tensor and therefore the CKM matrix above. This is the sole controlling conditional CKM return here. The curvature-corrected matrix of Part II is not part of Theorem FD-2, and is not part of the controlling derivation.

§5.2 RRHF-AX1 — the axle as a declared additional premise

A separate execution, **RRHF-AX1**, installs the unique D_6 quadratic axle directly into the charged history transfer as one new discrete premise:

AX1. The unique D_6-quadratic $E_1\rightarrow E_2$ axle is installed as the common role-even weak-doublet history step, with $\Omega\_ax=\zeta E_{23}-\zeta^* E_{32}$.

Its coefficient is **not fitted**: $|\zeta|= b/\sqrt{3}$ comes from the prior C_3 norm theorem and $\arg\zeta=5\pi/6$ from the principal Hamiltonian half-lift. The augmented transfer is $U\_u^{(n)}=\exp(\Omega\_ax-1/2\Omega\_q^{(n)})$, $U\_d^{(n)}=\exp(\Omega\_ax+1/2\Omega\_q^{(n)})$, differentiated independently through the bosonic and CAR/Fock routes.

Level-4 return, independently re-executed and confirmed:

$$|V\_CKM^{(AX1)}|=\begin{bmatrix} 0.974818801518829 & 0.222968367274218 & 0.003662157897164 \\ 0.222837161278604 & 0.973997689543530 & 0.040891323250282 \\ 0.008479608404144 & 0.040169739343630 & 0.999156888722876 \end{bmatrix}, \quad |J|=3.023691651925415\times10^{-5}.$$

Row and column sums of squares are unity to machine zero; the standard-parameterisation $|J|$ implied by the moduli reproduces the quoted value to 1.2×10^{-13}; the angles $(\alpha,\beta,\gamma)=(88.582382^\circ,23.057830^\circ,68.359788^\circ)$ reproduce and sum to 180.000000°; $\delta=68.39^\circ$. Action-level gates: Fock functorial residual 1.129×10^{-17}, $\varepsilon\_CY3'^{(max)}=2.173\times10^{-14}$, minimum parent eigenvalue $1.027883126507>0$, continuations $\Delta\_CKM^{(3\to4)}=1.247249\times10^{-3}$ and $\Delta\_J^{(3\to4)}=2.600385\times10^{-2}$.

§5.2.1 Two structural observations

(a) The AX1 construction is algebraically the CKP construction. Since Ω_ax and Ω_q are anti-Hermitian,

$$V = U\_u^{(n)\dagger}\, U\_d^{(n)} = \exp(-\Omega\_ax + \tfrac{1}{2}\Omega\_q^{(n)}) \cdot \exp(+\Omega\_ax + \tfrac{1}{2}\Omega\_q^{(n)}),$$

which is exactly the Part II curvature form with $\Omega\_ax=\Omega_0$. **The only difference is that AX1 uses the regulator-covariant $\Omega\_q^{(n)}$ rather than the bare inherited generator.** RRHF-AX1 is therefore not a new mechanism; it is the curvature construction evaluated on the RRHF regulator ladder.

(b) The axle and the regulator step are close to independent. The regulator correction alone moves $|J|$ from the bare baseline 1.843506×10^{-5} to 1.860329×10^{-5}, +0.913%. Applied on top of the axle it moves 2.996768×10^{-5} to 3.023692×10^{-5}, +0.898%. **The same +0.90% in both cases** — a genuine internal consistency check that neither execution was reported here and that the two effects factorise at this order.

§5.2.2 The cumulative grade of the axle

Withdrawn: the HCMR-AF1 programme-level verdict that the action does not force the axle. Its author retracts it on three grounds, all of which stand. (i) It tested the wrong object, computing $1/2(\log U\_u+\log U\_d)$ from the already-constructed charged frames, whereas the object the prior papers define is the projected closure Hessian $H\_W=P\_W^\dagger H\_clP\_W$ with $P\_W = P\_W^{supp} P\_R P\_C$, keeping the pre-readout C_3-democratic block distinct from the post-readout survivor. A zero in the final reduced frames does not evaluate that pre-readout object, so by the controlling typed-comparison rule the test was not licensed. (ii) HCMR-MR1 is not the completed H_CMR; its own paper refuses admission on five obstructions including a P_Y that is effectively the identity. (iii) The "must be a mixed third derivative" inference was premature.

Withdrawn: the grading of AX1 as a free additional constitutive premise. It is not freely chosen.

Theorem FD-3 — Conditional Axle-Transport Theorem

Given (i) the action-selected oriented E_1 preparation mode; (ii) the canonical D_6-equivariant quadratic descent $Q:\mathrm{Sym}^2E_1\to E_2$; (iii) the weak-complement descent of the committed role-odd profile; (iv) **a nonzero action-returned coefficient c_ax**; (v) a **one-dimensional allowed physical intertwiner space** — note this is not a survival factor of one; (vi) a physical readout returning **the branch and the retention factor ρ_mult^physical**; and (vii) an action-derived phase carrier or comparison holonomy —

the role-even weak-doublet curvature is uniquely

$\Omega_0 = \zeta E\_r - \zeta^* E\_r^\dagger$, $r \in \{12, 23, 31\}$ returned by the action,

with

$|\zeta| = \rho\_\mathrm{mult}^\wedge\mathrm{physical} \cdot c\_ax/\sqrt{3}$,

and its phase determined by the returned carrier. **The special conditional consequence** is

$c\_ax = b$ and $\rho\_\mathrm{mult}^\wedge\mathrm{physical} = 1 \Longrightarrow |\zeta| = b/\sqrt{3}$,

neither of which has been derived. Stating the theorem this way makes it consistent with the four-normalisation ledger of §20 and with the repaired Stage I pass condition of §7.1.

Proof. Part II establishes $\mathrm{Hom}\_D_6(E_1,E_2)=0$ and $\dim \mathrm{Hom}\_D_6(\mathrm{Sym}^2E_1,E_2)=1$, so the descent exists, is quadratic and is unique up to scale and sign; the normalised image has exact threefold democracy; premises (iii)–(vi) fix the physical scale, the intertwiner and the returned branch; premise (vii) fixes the phase given that branch. ■

Once the coefficient, retention, readout and carrier are returned by the action, no continuously adjustable axle freedom remains. Before those returns the form is unique but its physical installation is not.

§5.2.2.0 Four levels of derivation, kept separate

Level	What is established	Status
Mathematical derivation	every odd-power map is forbidden; the quadratic axle is the	**PASS at lowest order**; the even tower above is unbounded and is not excused by universality (companion §G.6.4)

Level	What is established	Status
	unique allowed map **at that order**	
Conditional axle-transport theorem	given the named action/readout conditions, the curvature has the unique profile and 1/√3 sharing	**PASS** (Theorem FD-3)
Executed action-level return	RRHF-AX1 returns the improved charged tensor and CKM matrix at two regulators	**PASS**
Primitive admission	the completed H_CMR itself returns the coefficient, internal lift, readout and phase carrier	**TESTED AND NOT RETURNED** (§7.0) — Stage I closes the source-mode gate and returns neither a nonzero c_ax nor the **complete full-carrier lift**, physical readout or phase carrier — **the local gauge-compatible Q-sector intertwiner form is derived** (AXD-2); the obstruction is underdetermination, not unfinished computation

CONDITIONAL AXLE-TRANSPORT THEOREM: PASS

PRIMITIVE SAME-ACTION ADMISSION: OPEN

§5.2.2.1 The grade, itemised by component

The cumulative flavour programme fixes the unique quadratic D_6 axle **at lowest admissible order** (AX-1.5b), its threefold branch geometry and the representation-geometric factor 1/√3. Conditional on the completed action identifying the axle coefficient with the inherited common-mode norm b and retaining one branch without further rescaling, its physical magnitude is b/√3.

Component	Grade
Axle form (unique quadratic $Sym^2E_1 \to E_2$, lowest admissible order)	**exactly fixed at that order**; higher even channels not excluded, and shown to change CKM observables rather than flow away, so quadratic-only is a **minimality prescription** (AX-1.5b, companion §G.6.4)
Threefold democracy	**exactly fixed**
Relative magnitude 1/√3	**exactly fixed**
Physical magnitude b/√3	**conditional** on c_ax=b and readout retention (§5.2.2.2)

Component	Grade
Quadratic character descent and source-phase candidate	**derived**
Physical phase *value*	**not returned.** The jet is component-preserving, the physical readout is unreturned and the local metric is exactly phase-isotropic. Under the provisional and unestablished choice $P\_vis^curv = P_{23}$ the selected component would carry 30° or 270°; the 150° execution remains a quarantined diagnostic

§5.2.2.2 The magnitude chain, typed explicitly

Two statements are routinely compressed into one and are **not** identical. The first is an exact geometric normalisation; the second requires the physical master action to supply the axle with norm b:

$\|Q(u)\|\_geom=1/\sqrt{3}$,

c_ax=action-returned physical axle norm,

$|\zeta\_phys| = \rho\_mult^physical \cdot c\_ax/\sqrt{3}$,

whence

$c\_ax = b \wedge \rho\_mult^physical = 1 \Longrightarrow |\zeta\_phys| = b/\sqrt{3}$.

A derivation must calculate c_ax, not merely observe that b is the natural shared norm. That calculation is **Stage I-A**; whether it equals b is a separate **Stage I-B** comparison (§7.1).

AXLE FORM AND DEMOCRACY: EXACT PHYSICAL MAGNITUDE: CONDITIONAL ON SAME-ACTION NORM TRANSFER QUADRATIC CHARACTER DESCENT AND SOURCE-PHASE CANDIDATE: DERIVED PHYSICAL PHASE CARRIER AND COMPARISON HOLONOMY: OPEN PHYSICAL PHASE VALUE AND BRANCH: NOT RETURNED; LOCAL PHASE SELECTION CLOSED NEGATIVELY

This is stronger than "free assumption" and weaker than "the full axle is forced". **RRHF-AX1 is the execution of the 150^∘ branch under those conditions.** It has no continuously fitted coefficient, but it is not yet the output of a complete same-action microscopic projection.

§5.2.3 The phase: character descent derived; physical carrier, holonomy and value not returned

The quadratic descent $Q(z)=z^2$ carries the selected sixfold character into the C_3 generation character, and the Hamiltonian-to-frame map then produces the **source-phase candidate** $5\pi/6$, without establishing that the physical branch coefficient inherits it — which is why Part II records the axle as *deriving* the half-lift relation $\hat{h}^2$=O_q rather than assuming it. The reduced scalar mechanism's 90^∘/270^∘ result is restricted to that reduced construction and is **not** a programme-level competitor.

The genuine remaining problem is narrower, and it is a question about which physical object inherits the phase.

Reading	Consequence
The branch coefficient inherits the **source mode's rim-to-rim advance**	argζ=150^∘ **independently of the visible branch**
The branch coefficient inherits the phase of the **selected descended component of Q(u)**	argζ∈{150^∘,30^∘,270^∘} according to the unreturned readout branch

CONTROLLING STATEMENT (companion §G.6.11). AXD-3 establishes that the completion jet is component-preserving. **It does not establish which component is physically read.** Under the provisional and unestablished choice P_vis^curv = P_{23} the selected component would carry 30° or 270°. **HCFC-1 shows that this projector is not derived and that the local metric selects no phase.** The quadratic character descent and a source-phase candidate are derived; the physical carrier, comparison holonomy and phase value are not returned — **150°, 30° and 270° are all physically unreturned.**

§5.2.3.1 The one equation that would close it

The derivation strengthens decisively the moment the action returns an explicit operator identity

** P_vis^curv 𝕀_int Q̂(u) = ρ_mult^physical e^(iθ_phys) (1/√3) E_r, r ∈ {12, 23, 31} returned by the action **

with θ_phys supplied by **the action, not by comparison with CKM**. That single equation fixes five open objects simultaneously: the internal lift 𝕀_int; the visible branch; the multiplicity retention ρ_mult^physical; the phase carrier; and the physical phase. Until it is calculated the correct claim is exactly:

quadratic character descent and source-phase candidate derived; physical phase branch open.

The fork is no longer abstract. Stage I supplies both readings as concrete **candidate objects** that disagree: its minimal-cell reconciliation reports source isotype k=0 (trivial), branch vector (1,1,1)/√3 with all components in phase, while its axle return gives Q(u) with character e^(2π i/3), branch arguments (0^∘,-120^∘,+120^∘). Identical moduli, phases differing by exactly ±120^∘ (§7.0.2).

Why this is the load-bearing gap. The empirical gain of AX1 is carried by the phase, not the magnitude. Holding |ζ|= b/√3 and varying only argζ:

$\arg\zeta$	J	χ^2_5
30° *(a carrier-reading value)*	$+2.997\times10^{-5}$	**539.0** (five-entry set, consistent with this table's heading; the eight-entry figure is 3000.7)
150° *(the value executed)*	$+2.997\times10^{-5}$	**5.25**
270° *(a carrier-reading value)*	-4.656×10^{-6}	1031.6

against $\chi^2_5=290.7$ for RRHF with no axle at all. The three carrier-reading values are not close: one reproduces the measured pattern and the others do not. **Determining the carrier is not a matter of convention or declaration.** It requires either an action-derived physical readout P_{vis}^{curv} or a realised comparison-loop holonomy fixing the phase independently of the local branch component (companion §G.6.11).

§5.2.3.2 Required perturbation tests

Three tests must accompany any published AX1 certificate, and their evidential weights differ.

(a) Axle-removal. Setting $c_{ax}=0$ must recover the baseline RRHF matrix. Verified: $\|V(c_{ax}=0)-e^{(\Omega_q)}\|=1.4\times10^{-16}$. **This is structurally guaranteed** — at $\Omega_{ax} = 0$ the construction reduces identically to $\exp(\Omega_q^{(n)})$ — so it is a coding check, **not physical corroboration**, and must be reported as such.

(b) Global conjugation. The valid opposite global branch must preserve all moduli and $|J|$ while reversing only $\operatorname{sgn} J$ (Part II, verified at exact machine zero).

(c) Carrier alternatives — published together, not selectively. All action-permitted phase-carrier readings must be computed and shown *before* comparison with experiment:

carrier reading	$\arg\zeta$	J	$\lvert V_{ub}\rvert$	$\lvert V_{td}\rvert$	χ^2_5
source-mode advance / branch 0	150°	$+2.997\times10^{-5}$	0.003557	0.008690	**5.3**
descended component, branch -	30°	$+2.997\times10^{-5}$	0.006541	0.004859	539.0
descended component, branch +	270°	-4.656×10^{-6}	0.003372	0.005763	1031.6

Showing only the 150° branch would be selection by presentation. The table was published in advance of the carrier declaration, precisely so that it could not be curated afterwards. **Under the provisional P_{23} diagnostic**, the B_{23} row would carry 30° or 270°, giving χ^2_5 = 539 or 1032. **This is not a returned readout:** the jet preserves whichever component is selected; it does not select the component. The 150° row remains a quarantined diagnostic.

§5.2.4 RRHF-AX1 in the derivation chain

RRHF-AX1 is best presented not as a separate premise-set execution but as a **conditional completion execution** of Theorem FD-3:

exact axle theorem+named physical premises⟶Ω_0⟶Γ_RRHF+axle⟶ V_CKM^(AX1).

RRHF-AX1 is an execution of the conditional axle-transport theorem. It is not a free numerical deformation. Its promotion to a primitive derivation depends only on whether the completed action returns the named physical premises.

RRHF-AX1 is not a same-action forcing result. It executes the exact lowest-order axle form under additional conditional choices about its physical norm and the 150° direction. It contains no continuously fitted parameter, but the physical direction is unreturned and it was tested with the empirical shortfall already known. **It is a diagnostic execution, not the controlling derivation.**

Two things remain true at once. *In its favour:* AX1 has no fitted continuous parameter; its form and threefold democracy are exact, and its magnitude follows conditionally from the inherited norm. Moving χ^2_5 from 290.7 to 4.10 with nothing tuned is not curve-fitting — **but that execution is on the 150° branch, which companion §G.6.10 shows the derivation does not return; at the provisional P_{23} branch the same construction gives 539 or 1032, worse than the 290.7 baseline.** *Against:* the decision to execute the 150^∘ branch was taken with the empirical gap of §6.1 known, and the phase carrier that selects that branch is precisely what remains undeclared (§B.1).

The question the paper now asks has changed. It is no longer *"is AX1 an arbitrary extra assumption?"* — that has been answered: no, its form and relative geometry are forced. It is:

Does the completed H_CMR physically install the unique lowest-order structure, and what coefficient, branch readout, multiplicity retention and physical phase does it return?

The returned values may then be compared post-freeze with b, 150^∘ and the CKM data. **The desired values must not appear in the derivational question itself** (F4).

That is narrower and stronger, and it is exactly what Stage I of §7 tests.

§5.1 A closed debt

The archived figure recording a frame "approximately 37.92% away from the frozen curvature boundary" is now identified: (2.996771-1.860329)/2.996771=37.92% exactly, i.e. **the level-4 RRHF frame**. An earlier inference in this series, that the figure belonged to a distinct RRMK frame differing from RRHF by ≈2.7%, is withdrawn. Debt I6b is closed.

§6. Status of the curvature module

The module proposes $\Omega_0(z)=zE_{23}-z^{*}E_{32}$ and V_curv=e^(-Ω_0+Ω_q/2)e^(+Ω_0+Ω_q/2). The exact D_6 results are retained: Hom_$D_6(E_1,E_2)=0$ and dim Hom_$D_6(\mathrm{Sym}^2E_1,E_2)=1$, proving the first permitted axle is quadratic and unique up to scale and sign. They do **not** prove that the RRHF parent contains that axle with a nonzero physical coefficient. Therefore

$\Omega_0 \notin$ the retained conditional derivation in CKP-2.4.

This is an admission statement, not a negative theorem about the primitive action. **Stage I-A has now tested the physical installation directly** (§7.0). The source-mode gate closes and the unique quadratic descent is reconfirmed on the actual selected mode. Stage I-A returns no nonzero c_ax, complete full-carrier lift, physical readout, physical retention or phase carrier. The local gauge-compatible Q-sector intertwiner form is derived. The admission statement therefore stands, now supported by an executed test rather than the absence of one. **The obstruction is underdetermination of the complete full-carrier lift and readout, not unfinished computation.**

§6.1 Empirical status of the controlling return — stated plainly

The exclusion of Ω_0 has a large empirical cost, and the paper records it rather than leaving it to be found.

Quantity	RRHF (controlling)	pull	curvature-corrected (not admitted)	pull
	V_us		0.222964	-1.58σ
	V_cb		0.040639	-0.38σ
	V_ub		0.004198	+1.89σ
	V_td		0.005895	**-13.52σ** (-31.4%)
J	1.860×10^{-5}	**-10.08σ** (-40.4%)	2.997×10^{-5}	-0.99σ
(α,β,γ)	(129.5^∘,20.4^∘,30.1^∘)		(84.5^∘,22.5^∘,73.0^∘)	
χ^2 (five entries)	**290.7**		5.3	

The retained conditional derivation reproduces the Cabibbo sector and the 2↔3 magnitude, and fails badly on the CP measure and the short side of the triangle. That is the honest current state of the programme's conditional flavour return, and it follows directly from the controlling frame sitting within one percent of the bare inherited generator (§5): without a role-even correction the construction has too little structure to generate the observed CP area and short triangle side.

AXD-3 has since shown that the provisional P_{23} branch does not close this gap — it widens it. With the component-preserving carrier under the provisional P_{23} choice the axle **would** carry 30^∘ or 270^∘; because that readout is unestablished, no physical phase is returned, the **provisional diagnostic** gives χ^2_5 = 539 or 1032 against **290.7 for no axle at all** (companion §G.6.10). **RRHF-AX1 closes this gap under an added premise, not by derivation** (§5.2): χ^2_5=4.10, with |V_td| at -0.60σ and J at -0.77σ. That is the correct comparison to keep in view — the gap is closable by exactly the block the geometry identifies, with no fitted coefficient, and the open question is whether the controlling action returns that block or merely tolerates it.

Nothing in this section is an argument for admitting Ω_0. Selecting a block because it improves agreement is forbidden (F4). The section exists so that the gap is on the record **before** the forcing test is run, and so that any later pass is read in that light (§B.1).

§7. The forcing test, restated in two stages

The previous formulation asked the wrong action. I6a was defined as a search inside the RRHF parent: differentiate Γ_RRHF, obtain $H\_even^{(n)}$, project onto the axle operator. But §5.2.2 accepts that the primitive object defined by the prior flavour papers is the projected closure Hessian

H_W=P_W† H_clP_W, P_W = P_W^supp P_R P_C,

with the pre-readout and post-readout stages kept separate. **Asking the RRHF charged transfer to force the axle repeats the error that invalidated HCMR-AF1**: it tests a downstream constructed object rather than the primitive projection. RRHF is the **integration and reproduction** test, not the source of the forcing theorem. I6a is therefore replaced by a two-stage test.

§7.0 Stage I has been executed — result and standing

C3PI-2 / AXD-1 was run on 2 August 2026 against the frozen prior packages. Its returns are re-executed and confirmed below. **One previously open gate closes; the already-proved quadratic descent is reconfirmed.**

Stage I item	Result	Independently reproduced
1. Source-mode identification E_1^(prep)=k=±1	**PASS**	overlap with the canonical k=-1 harmonic =1.0000000000000000; ‖Lv-2v‖=1.75×10^{-15}; hub amplitude 7.24×10^{-17}; rim spread 8.33×10^{-16}; phase step -1.047197551196597 against -π/3
2. Unique quadratic descent	**PASS EXACT**	linear pair-sum 5.82×10^{-16}; Q(u)=(1/3,-1/6-i√3/6,-1/6+i√3/6), norm 0.577350269189626=1/√3, moduli all 1/3, spread 6.1×10^{-16}, sum 1.3×10^{-15}; ‖LQ-4Q‖=5.21×10^{-15}, so the image lies in the same-wheel E_2 sector; C_3 character e^(2π i/3)
3. Nonzero action coefficient c_ax	**NOT RETURNED**	no same-action contraction maps the source axle into the weak-doublet Hessian
4. Weak-complement descent	**NOT RETURNED**	profile overlap unavailable

Stage I item	Result	Independently reproduced
5. Full-support multiplicity	**NOT RETURNED**	only the minimal regular copy is executable
6. **Stage I-B comparison c_ax = b**	**NOT REACHED** — Stage I-A returned no nonzero physical c_ax, so the comparison was never made	the $1/\sqrt{3}$ split is exact; the physical norm transfer is not
7. Physical readout P_vis^curv	**DECLARED-MODEL DIAGNOSTIC ONLY; FULL-ACTION NON-RETURN**	Π_23, t_12=t_31=0, retention 1 — but by construction inside the declared three-branch model
8. Phase-carrier declaration	**CHARACTER DESCENT PASS; PHYSICAL PHASE CARRIER AND VALUE NOT RETURNED** — AXD-3 derives a component-preserving jet, but P_vis^curv is unreturned and the local metric is phase-isotropic (companion §G.6.11)	character descent exact

Verdict: Stage I-A does not pass. By §7.4 the controlling matrix therefore remains |V_CKM^(RRHF)|, and the 150^∘ execution remains conditional. **A0 / VX-W6 is closed**, which is a genuine new cross-package result: the axle's source is now known to be the *actual* action-selected mode, not a separately chosen representative.

§7.0.1 The obstruction is underdetermination, not unfinished computation

This distinction decides how the remaining work should be planned, and the executed numbers support it.

The EX2F separable endpoint profile is a 7×50 array of **numerical rank one**, leading singular value $2.6457513110645907=\sqrt{7}$ exactly, the next at 1.13×10^{-15}. **A rank-one profile cannot distinguish internal directions**, which is precisely why no unique faithful lift is selected. Consistently, EX2F leaves local-gauge nullity 84 and **post-gauge physical flat nullity 836** of 920 total, and HM4-EX3B constructs an exact embedding family

$\varepsilon\_HM4(\theta)=\sin\theta$,

sweeping the whole interval [0,1] at residual 1.87×10^{-15} — from fail to pass continuously. **The supplied source metric does not select a faithful tangent embedding at all.** That is underdetermination, and no amount of further computation on the wheel or on the axle's mathematical form addresses it.

§7.0.1.1 The obstruction is not a flavour obstruction

This section relocates the blocker, and the relocation is the useful result.

The programme-level completion ledger of 2 August records, for the complete H_CMR, that the action still does not calculate

Δ I_ρ(η), ν_ρ(dp), η_realised, **z_realised**, Δ S, **relative phases**, γ_copy, R_physical,

and gives the reason: *the known invariant data do not select the needed noncentral coherences; multiple physically distinct action lifts have identical central traces and symmetry data; more diagonalisation cannot create the missing information.* It quantifies this with two inequivalent marked-score lifts reproducing the same supplied Hessian to 10^{-15} while differing by ≈1.7366 in relative norm.

Those two boxed items are this chapter's two open objects. The realised complex amplitude is c_ax; the relative phases are θ_phys. And the failure mode is identical to Stage I's: a rank-one endpoint profile with 836 post-gauge flat directions, and an exact embedding family ε(θ)=sinθ running from fail to pass, is the same statement as "multiple physically distinct lifts share all supplied invariants" — **the data are compatible with a continuum of inequivalent structures.**

Three consequences follow.

The flavour blocker sits inside the programme critical path, not beside it. It is not a peculiarity of the weak-doublet carrier that a faithful lift fails to be selected; it is the flavour-sector instance of the H_CMR non-uniqueness. Solving it in the flavour sector alone is not available.

No amount of further flavour calculation reaches it. The programme ledger's own phrasing — more diagonalisation cannot create the missing information — applies verbatim to §7.0.1.

The family can now be counted. the companion's Appendix G characterises it on the minimal carrier: 6 invariant quadratic couplings, 13 cubic, 3 in the preparation reference kernel, with the graph Laplacian sitting at a non-generic range-1 point whose two long-range coordinates are assumed rather than returned. That turns "the data are compatible with a continuum" into a dimension, a basis, and an address for c_ax.

And the two sectors now corroborate each other's diagnosis, having been reached independently: one by a rank census on the EX2F endpoint profile, the other by exhibiting two inequivalent score lifts with identical central data. Neither used the other's evidence.

§7.0.2 A concrete clash the execution does not flag

Two of the execution's own returns disagree about which branch vector is physical, and the disagreement is exactly the phase-carrier fork.

- The minimal-cell reconciliation reports **source isotype k=0, trivial**, whose branch vector is $Z_0=(1,1,1)/\sqrt{3}$ — all three components in phase.
- The axle return gives Q(u) with C_3 character e^(2π i/3), i.e. the ω isotype, whose normalised branch vector has component arguments (0^∘,-120^∘,+120^∘).

Both have **identical normalised component moduli**, so the geometric $1/\sqrt{3}$ sharing is unaffected. **This does not determine the physical amplitude**, which remains $\rho\_mult^{physical} \cdot c\_ax/\sqrt{3}$. They differ only in relative branch phase — by exactly $\pm 120^\circ$. **This is the phase-carrier fork appearing as a concrete inconsistency between two supplied results rather than as an abstract typing question**, and it must be resolved by naming which vector the action attaches to the physical branch. Until then, the **physical phase-carrier component of Stage I-A** remains open.

§7.0.3 What the branch-cost audit confirms

The execution reports that **the declared principal root is stable under a true coordinate pushforward** — independently confirming the base-point correction recorded in Part II — while the corpus does not derive the branch-specific cost carrier, metric or reference element from the action, so TYPE-ORD remains conditional.

§7.0.4 A naming note

The execution refers to the canonical wheel as W_7 (seven vertices); this paper writes W_6 ($C_6 \vee K_1$, six rim vertices plus hub). **They are the same graph.** All spectra quoted here — L=2 for E_1, L=4 for E_2, spec(L)={0,2,2,4,4,5,7} — agree under either convention.

§7.1 Stage I — full same-action forcing test

The completed H_CMR must return, without CKM input,

$K\_even,pre^q = \tfrac{1}{2} Tr\_role(P\_q\ P\_R\ P\_C\ H\_cl\ P\_C\ P\_R\ P\_q)$,

decomposed by role parity,

$H\_W^{(n)} = H\_even^{(n)} \otimes I + H\_odd^{(n)} \otimes \tau_3 + H\_mix^{(n)}$,

and return directly $c\_ax^{(n)}$, $\mathbb{I}\_int^{(n)}$, $P\_vis^{curv(n)}$ and $\theta\_phys^{(n)}$, delivering the following:

Stage I-A — physical installation. Require, all fully typed, regulator-stable and returned without target access:

$c\_ax \neq 0$, $\mathbb{I}\_int$, P_vis^{curv}, $r \in \{12,23,31\}$, $\rho\_mult^{physical}$, θ_phys.

Stage I-B — post-seal inherited-value comparison. Test separately: $c\_ax = b$, $\rho\_mult^{physical} = 1$, $\theta\_phys = 5\pi/6$.

Closed preliminary gates, not continuing requirements of every future Stage I-A execution: source-mode identification (CLOSED by Stage I, overlap 1.0000000000000000) and the unique quadratic descent.

This is the previously specified successor calculation strengthened by the D_6 axle, and it should be executed under its own designation:

C3PI-2 / AXD-1 — Full-Carrier Same-Action Derivation of the Quadratic Weak-Doublet Axle. *Does one frozen completed H_CMR return the quadratic $E_1 \to E_2$ channel, its physical norm, its faithful internal embedding, its branch readout and its phase carrier —* **and *does it return that channel alone, or also higher even channels?*** *(AX-1.5b)*

Executed 2 August 2026 (§7.0). The channel and its source are returned; the norm, embedding, readout and carrier are not. **The successor calculation is therefore narrower than C3PI-2 was**: it need not revisit the wheel, the descent or the democracy, and should address only the internal lift and readout. **But §7.0.1.1 shows that is an H_CMR calculation, not a flavour one**: the objects it must return are the realised complex amplitude and the relative phases, which the programme ledger lists among the complete action's unreturned outputs. The successor is therefore better placed in the H_CMR line than the flavour line, and the working designation **AXD-2 / I_int-1** is reserved for it on that understanding.

Its method has since been corrected, and the correction matters more than the designation. A first attempt under that name concluded the programme contains no weak-doublet flavour bridge and has been withdrawn: it searched the selected executable for a primitive-coordinate third derivative and read its absence there as absence from the programme. The bridge exists in descended coordinates as a second variation (companion §G.6.8). **The successor should therefore not begin by constructing a new third-derivative term.** It should first compile the existing projected-Hessian chain into the full modern carrier —

species and role addresses $\longrightarrow P\_C, P\_R \longrightarrow \mathbb{I}\_\mathrm{int} \longrightarrow H\_W = P\_W^\dagger H\_\mathrm{cl}\, P\_W \longrightarrow P\_\mathrm{vis}\string^\mathrm{curv}$ —

and test at regulator levels 3–6 that $[G\_C_3\string^\mathfrak{B}, R_3] = 0$, that $\dim \mathrm{Hom}\_C_3(\chi\_\mathrm{axle}, \mathfrak{B}\_\mathrm{full}) = 1$, that the role-even/role-odd overlap is unimodular, and that $P\_\mathrm{vis}\string^\mathrm{curv}\ \mathbb{I}\_\mathrm{int}\ \hat{Q}(u) = \rho\_\mathrm{mult}\string^\mathrm{physical}\ e\string^(i\theta\_\mathrm{phys})(1/\sqrt{3})E\_r$, with $\hat{Q}(u) = Q(u)/\|Q(u)\|$ the **normalised** image, **with $r \in \{12,23,31\}$, $\rho\_\mathrm{mult}\string^\mathrm{physical}$ and the retained norm all returned and tested against b rather than assumed** (F5). Only after that reconciliation is the fifth derivative the right instrument for the quartic question.

That question is sharper than asking whether RRHF "contains" Ω_0. The live obligations are the **incomplete Stage I-A installation tuple**: nonzero coefficient, complete full-carrier lift, physical readout, multiplicity retention and phase carrier. The inherited-value equalities belong only to the subsequent **Stage I-B** comparison.

§7.2 Stage II — RRHF integration test

Only after Stage I-A passes should the derived axle be inserted into the RRHF parent and the complete charged tensor differentiated again:

$$H\_CMR \longrightarrow \Omega_0 \longrightarrow \Gamma\_RRHF\string^(\mathrm{completed}) \longrightarrow Y\_u, Y\_d, Y\_e \longrightarrow V\_CKM.$$

Stage II tests that the primitive return survives the full charged-sector machinery and reproduces one regulator-stable matrix. It does not decide whether the primitive action forces the axle — that is Stage I's job alone, and no Stage II result may be quoted as though it were a forcing result.

§7.3 Sealing and anti-bias protocol

Knowledge of the empirical gap is unavoidable, so the seal cannot rest on intellectual blindness. Validity rests on **procedural independence**: the lift, projector, carrier metric, phase convention, thresholds and regulator rule must be generated from representation and action data alone, preregistered and hashed **before** the projected Hessian of Stage I is opened. No component may be revised after seeing either its projection coefficient or its effect on CKM observables. The implementation uses three artefacts: a **construction package** containing only source-side representation data; a **sealed execution package** holding the Stage I projection; and a **final comparison script** that joins them only after both hashes are published.

Stage I-A passes when the completed action returns a nonzero, regulator-stable and fully typed physical axle — c_{ax}, $\mathbb{I}_{int}$, P_{vis}^{curv}, r, $\rho_{mult}^{physical}$ and θ_{phys} — with a preregistered small residual. **No target CKM value, $b/\sqrt{3}$ or $150°$ phase may enter its construction.**

Stage I-B is a post-seal comparison only. It reports separately whether $c_{ax} = b$, $\rho_{mult}^{physical} = 1$ and $\theta_{phys} = 5\pi/6$. **Failure of any Stage I-B comparison does not invalidate a Stage I-A installation and does not prevent Stage II.**

§7.4 Matrices selected by the test

Stage I-A fails, or its physical-installation tuple is incomplete — ***this is the realised outcome*** *(§7.0):* the controlling matrix remains $|V_{CKM}^{(RRHF)}|$ with $|J|=1.86032934918\times10^{-5}$, and the 150° execution stays a conditional execution under a declared carrier choice.

Note (companion §G.6.11). AXD-3 derives a component-preserving jet. Under the provisional choice $P_{vis}^{curv} = P_{23}$ this would produce 30° or 270° rather than 150°. **However P_{23} has not been derived, so no physical phase and no Stage-II matrix is presently selected.** In particular the matrix a Stage-II pass would reissue is **not** the $|V^{(AX1)}|$ quoted elsewhere in this paper. The rule below was written to prevent exactly the substitution that would otherwise be tempting here.

Stage I-A pass, followed by Stage II: the complete action-returned axle proceeds whether or not its values agree with AX1. The Stage I-B comparisons with b, unit retention and $5\pi/6$ are reported separately. **The Stage I-A return — not the target value —** is inserted into the full charged parent and the complete charged tensor is re-differentiated.

The final matrix must be whatever that calculation returns. The attractive matrix cannot be preregistered as the required outcome.

§8. Orientation and the IORC preparation current

IORC-1 supplies a source-supported preparation-sector action whose minimiser gives $J\_rim^{(\sigma)}=\sigma 0.0602316978168341$ and $\Delta\Sigma\_Fold=\ln 2$, with antipodal quotient returning the conjugate generation characters $O\_q(\sigma)=e^{(i\sigma 2\pi/3)}$. This advances the origin of orientation but does not by itself derive the role-odd seed holonomy. The typed chain is

$j\_prep^{(\sigma)} \longrightarrow O\_q(\sigma) \xrightarrow{\mathbb{I}\_role\text{-}odd} \Omega\_q,13^{(\sigma)}$,

$\mathbb{I}\_role\text{-}odd[O\_q(\sigma)]=e^{(i\sigma 2\pi/3)}E_{13}-e^{(-i\sigma 2\pi/3)}E_{31}$,

with the physical seed $\Omega\_q,13^{(\sigma)}=c\mathbb{I}\_role\text{-}odd[O\_q(\sigma)]$ — the map carries the orientation, the amplitude c comes from the count package of §C. The role-odd holonomy belongs to the frozen RRHF ledger; IORC gives it a candidate action ancestry and a realised sign, but primitive closure requires same-action derivation of $\mathbb{I}$_role-odd, installation of the preparation block within the complete H_CMR, and a maintenance lock making the realised orientation persistent. **No statement about the observed sign of CP violation follows before those returns.**

§9. What this derivation proves and does not prove

Proved conditionally at action level: one positive representation-resolved charged parent; one rank-nine charged tensor; Y_u,Y_d,Y_e at two regulators; independent bosonic/CAR agreement; $V\_CKM=U\_u^\dagger U\_d$; one regulator-stable conditional CKM modulus matrix; no target loading during the core solve.

Exact but not dynamically installed: the W_6 graph and Laplacian structure; the odd-power no-go $\mathrm{Hom}(\mathrm{Sym}^k E_1,E_2)=0$ for every odd k, which requires only a $\mathbb{Z}_2$ parity and is therefore robust against any deformation preserving it; uniqueness of the quadratic $E_1\longrightarrow E_2$ axle **at that order**, which requires the full D_6; exact threefold democracy of its normalised image.

Tested and not returned by Stage I-A: a nonzero physical coefficient, the complete full-carrier lift, physical multiplicity retention, branch readout and phase carrier. **Stage I-B not reached:** $c\_ax = b$, unit physical retention and $\theta\_phys = 5\pi/6$. **Closed by Stage I:** source-mode identification. **Derived by AXD-2:** the local Q-sector intertwiner form.

Moot for the presently tested action: whether the quartic coefficient survives coarse-graining — a **fifth**-derivative question in primitive coordinates, $D^5S[u,u,u,u,\mathfrak{B}]$, with two channels available (companion §G.6.8). The emergent-time record-conditioned test of companion §G.6.6 was executed and correctly withdrew a stronger earlier claim, but its three findings are respectively an identity, a symmetry tautology and a self-conjugate correlator, and it ran on two dimensions rather than the effective one. The informative comparison — a Standard-Model-relevant quantity not conjugate to the coupling, at equal fact count, on the effective dimension — has not yet been made. **A cheaper alternative now exists** and may retire the debt without it: test whether the quartic channel moves $\|P\_A z\|^2$ for any commitment projector (companion §G.6.7).

Not proved, and shown not to be excusable: that the quadratic is the **only** even channel — the quartic enters the same $1\leftrightarrow 3$ entry with the conjugate C_3 character, so A 30% relative quartic source amplitude gives a 5% norm shift and roughly a factor-three variation in the diagnostic χ^2;

a 5% source amplitude gives only a 0.83% norm shift (companion §G.6.4) — parity gives a floor and no ceiling, and dim $\mathrm{Hom}(\mathrm{Sym}^4E_1,E_2)=2$ (AX-1.5b); primitive uniqueness of the representation ledger; primitive uniqueness of P_Y; I6a admission of Ω_0; physical sign of J; persistence of the IORC orientation; $O\_q\to\varphi$ as a same-action operator map; the neutral mass block and PMNS; the full strong phase; the absolute scale and scheme bridge; unconditional Standard Model admission.

And not achieved empirically: the controlling conditional return misses |V_td| and J by -13.5σ and -10.1σ (§6.1).

Part II — The D_6 axle as a candidate RRHF component

Position under the spine. This Part supplies the object whose membership in the RRHF parent the fork of §B decides. Its representation-theoretic content — $\mathrm{Hom}(E_1,E_2) = 0$ under a single $\mathbb{Z}_2$ parity, and dim $\mathrm{Hom}\_D_6(\mathrm{Sym}^2E_1,E_2) = 1$ under the full D_6 — is exact and holds whether or not the channel is physically installed; that is clause 4 of Theorem A. Its physical coefficient, embedding, readout and phase carrier are not returned here and may already appear in different coordinates inside RRHF.

§14. Source-locked inputs

The axle test uses five archived action- or substrate-side returns: the K=7 D36 history kernel; the isolated sixfold slow mode selected by PFDMI-EX2F; the preferred circulation returned by PFDMI-EX2G; the unique three-antipodal-axis decomposition of the D_6 wheel returned by PFDMI-EX2H; and the explicit canonical W_6 closure-transition graph and spectrum. **No CKM or PMNS datum, phase target, branch magnitude or readout coefficient enters.**

§14.1 Explicit canonical wheel certificate

Let $K=\{\kappa\_h,\kappa\_b_1,\ldots,\kappa\_b_6\}$. The canonical closure graph contains the six rim edges $\{\kappa\_b(i), \kappa\_b(i+1)\}$ and six spokes $\{\kappa\_h,\kappa\_b\_i\}$, hence $|V|=7$, $|E|=12$, $\beta\_1=|E|-|V|+1=6$. Its combinatorial Laplacian has the exact spectrum spec(L)={0,2,2,4,4,5,7}, **independently verified**: rim degree 3, hub degree 6, and the wheel closed form $3-2\cos(2\pi k/6)$ for k=1,…,5 together with 0 and |V|.

§14.1.1 The transition spectrum is not yet reproducible — $\hat{T}$ must be defined

The quoted transition spectrum {1,1/2,1/2,-1/4,-3/28,0,0} **is not returned by any standard transition operator on W_6**. Tested and excluded: the random walk $D^{-1}A$, the normalised adjacency $D^{(-1/2)}AD^{(-1/2)}$, $I-D^{-1}L$, $I-D^{(-1/2)}LD^{(-1/2)}$, and $I-L/\lambda$ for $\lambda\in\{4,6,7\}$. All return different spectra.

The single closest candidate is $\hat{T} = I - 1/4\, L$, spec($\hat{T}$)={1,1/2,1/2,0,0,-1/4,-3/4}, which reproduces **six of the seven quoted entries exactly** — including both the Perron value, the doubled 1/2, the doubled 0 and the -1/4. The sole discrepancy is the hub mode: -3/4 rather than -3/28, and -3/28=-3/4·1/7, suggesting a stray factor of |V| or of the hub Laplacian eigenvalue in transcription.

Nothing in the axle depends on the hub entry. Under the candidate operator $\hat{T} = I - \frac{1}{4}L$ the k=±1 harmonics have eigenvalue 1/2. The quoted spectrum contains a doubled 1/2, **but without an explicit operator its eigenspace cannot be identified** — a list of eigenvalues does not say which eigenvectors carry them. But $\hat{T}$ must be given an explicit definition before any of its eigenvalues is cited as a substrate return (new debt **A0b**, falsifier **F23**). The slowest non-trivial eigenspace is the two-dimensional pair of boundary Fourier harmonics $\psi_{\pm1}(\kappa_{b_j})=e^{(\pm 2\pi i j/6)}$, $\psi_{\pm1}(\kappa_h)=0$, with $L\psi_{\pm1}=2\psi_{\pm1}$, hence $\hat{T}\psi_{\pm1}=1/2\psi_{\pm1}$ under $\hat{T}=I-1/4L$. The half-turn acts exactly as $r^3\psi_{\pm1}=e^{(\pm i\pi)}\psi_{\pm1}=-\psi_{\pm1}$.

§14.1.2 The axle target is a second eigenspace of the same wheel

The rim harmonics split by antipodal parity, and the parity classes are not irreducible:

Rim sector	Full-wheel status	half-turn	D_6 content
k=0 rim constant	**trivial representation, but not by itself a full W_6 eigenvector** — it mixes with the hub to form the L=0 and L=7 modes	+I	trivial
k=±1	L=2	-I	$\mathbf{E_1}$
k=±2	L=4	+I	$\mathbf{E_2}$
k=3	L=5	-I	one-dimensional alternating

Note. the rim-uniform vector with zero hub component satisfies Lf=(1,…,1,-6), which is not proportional to f. The trivial D_6 sector is two-dimensional, spanned by the rim-uniform and hub directions, and splits into the full-wheel eigenvalues 0 and 7 — verified exactly. **The representation statement is unaffected:** the antipode-even rim sector is still trivial $\oplus$ E_2; only the eigenvalue assignment needed correcting.

Two consequences, both new and both verified.

(a) The three antipodal axes are trivial$\oplus$ E_2, realised on the same rim. The antipode-even rim functions are spanned by k=0,2,4, of dimension 3, decomposing exactly as trivial $\oplus$ E_2. **The unique quadratic axle $Q(z)=z^2$ therefore maps the k=±1 eigenspace of the wheel to the k=±2 eigenspace of the same wheel.** The generation target is not an external three-axis structure imported alongside the graph; it is a second harmonic sector of the graph itself. This is a strong internal consistency check that the earlier presentation did not make, and it removes a joint that previously had to be assumed.

(b) Antipodal oddness alone does not single out E_1. The antipode-odd rim sector is also three-dimensional, $E_1 \oplus$ the k=3 alternating mode at L=5. A mode is therefore *not* identified as E_1 by parity; the eigenvalue is needed too. This strengthened the case that VX-W6 was a genuine obligation rather than a formality — the EX2F selection had to be shown to land on k=±1 specifically, not merely on an odd sixfold vector. **Stage I has since done exactly that** (§7.0), matching on the eigenvalue and not on parity.

§14.1.3 The target lies in the kernel of the candidate operator

Under the candidate $\hat{T}=I-1/4L$ the k=±2 eigenvalue is exactly 0, so $E_2 \subset \ker\hat{T}$ — verified. **The unresolved quoted spectrum also contains a two-dimensional zero sector, but it cannot be identified with E_2 until $\hat{T}$ is explicitly defined and shown to commute with the D_6 action.** The claim therefore holds under the candidate operator only, not "under both readings".

This is not obviously a defect, but it is a constraint that must be discharged rather than passed over. The axle maps the *slowest* non-trivial mode into a sector the transition operator annihilates. Read favourably, the generation carrier is a non-propagating frozen record, which is consonant with the corpus's record language. Read unfavourably, an amplitude deposited in E_2 is not sustained by the transition dynamics, so the coefficient c_ax cannot arise from the transition channel and must come from a stationary or constitutive term. **Debt A1 is sharpened accordingly**: "nonzero and regulator-stable" must be demonstrated for a channel whose target is in $\ker\hat{T}$, and the mechanism supplying it must be named.

Mode-identification premise VX-W6 — now DISCHARGED. The canonical wheel and its **verified Laplacian spectrum** supply the k=±1, L=2 source candidate and the k=±2, L=4 target, using **no eigenvalue of the unresolved $\hat{T}$**. The premise was that the EX2F/EX2G-selected mode is the k=+1 or k=-1 member of that eigenspace. **Stage I closes it** (§7.0): direct cross-package matching gives overlap 1.0000000000000000 with the canonical k=-1 harmonic, $\|Lv-2v\|=1.75\times10^{-15}$, hub amplitude 7.2×10^{-17} and half-turn parity -I. The axle's source is the *actual* action-selected mode, not a chosen representative. The graph computation proves that such an eigenspace exists and fixes its spectrum; it does not by itself prove that the marked-history action selects it. That identification **was** an installation-level provenance obligation and is now **discharged by Stage I**.

Let $E_1=\mathrm{span}_{\mathbb{R}}\{\mathrm{Re}\,\psi_1, \mathrm{Im}\,\psi_1\}$ be the real two-dimensional D_6 irrep carried by the canonical sixfold slow eigenspace, with VX-W6 discharged by Stage I, and let E_2 be that carried by the three antipodal axes. For the normalised selected mode, $u_{n+3}=-u_n$ follows analytically from k=±1, in addition to the numerical residual 2.24×10^{-16}; the generation-axis readout is antipode-even.

§15. The linear no-go

Theorem AX-1 (linear void-to-generation no-go)

$\mathrm{Hom}_{D_6}(E_1, E_2) = 0$.

Proof. On the canonical $k=\pm1$ boundary harmonics, the half-turn acts as $e^{(\pm i\pi)}=-1$, so $r^3=-I$ on E_1. Antipodal axes are unchanged by the same half-turn, so $r^3=+I$ on E_2. For equivariant T, $T(-v)=T(r^3v)=r^3T(v)=+T(v)$, while linearity gives $-T(v)$; hence $T=0$. ■

Certificates. Characters $\chi\_E_1 = (2, 1, -1, -2, 0, 0)$ and $\chi\_E_2 = (2, -1, -1, 2, 0, 0)$ on the six D_6 classes give $\langle\chi\_E_1, \chi\_E_2\rangle = 1.48\times10^{-16}$, and the canonical antipodal pair-sum $L(u)\_a = (u\_a + u\_a+3)/\sqrt{2}$ vanishes at $\|L\| = 2.05\times10^{-16}$.

Consequence. The earlier direct current-to-generation descent failed for a structural reason: it sought a linear axle that the symmetry forbids.

Corollary AX-1.1 — the no-go needs only a $\mathbb{Z}_2$, not D_6, and forbids every odd power

The proof above uses exactly one group element: the half-turn. Nothing else about D_6 enters. The theorem is therefore stronger than stated, and holds in a setting with no hexagonal symmetry at all.

Let σ be any involution acting on a source V by $\sigma v=-v$ and on a target W by $\sigma w=+w$. Then every σ-equivariant linear $T:V\to W$ vanishes.

Proof. $T(v)=T(-\sigma v)=-T(\sigma v)=-\sigma T(v)=-T(v)$, so $2T(v)=0$. ■

The division of labour is therefore clean and worth stating explicitly:

Result	Symmetry actually required
every odd-power descent forbidden (AX-1, extended in AX-1.5b)	**a single $\mathbb{Z}_2$ parity**
quadratic descent **unique at that order** (AX-2)	the full D_6
no upper bound on the even tower	**not established** (AX-1.5b)

This matters for robustness. The linear no-go survives any deformation that preserves the half-turn parity, whether or not it preserves the hexagonal symmetry; only the *uniqueness* of the quadratic replacement is a D_6 statement.

Corollary AX-1.2 — the same parity governs the weak-doublet split

Scope: an observation about two constructions, verified for those two and claimed for no others.

The handshake structure of §2 is governed by the same mechanism under a different involution. Let σ_role exchange the up and down readings. Then Ω_q is role-**odd** and Ω_ax is role-**even**, which is exactly why

$U_u=\exp(\Omega_{ax}-1/2\Omega_q)$, $U_d=\exp(\Omega_{ax}+1/2\Omega_q)$

carries the odd generator with opposite signs and the even one with the same sign on both sides — verified directly. In both constructions **the odd object cancels in the symmetric combination and the even object survives**; on the wheel this forbids the linear term and makes the square the first admissible power, and in the doublet it produces the two-sided split.

Squaring is therefore not a motif but a consequence: where the source is parity-odd, the square is the first parity-even power, and the coupling has nowhere lower to sit. On the wheel this is checkable directly — $u_{n+3}+u_n$ vanishes at 3.4×10^{-16} while $u_{n+3}^2-u_n^2$ vanishes at 2.8×10^{-16}, so the mode is odd and its square is even.

Corollary AX-1.3 — what does *not* share the mechanism

The corpus contains several further quadratic relations. **They are not instances of this mechanism, and grouping them with it would repeat a failure mode this series has already suffered four times** — three distinct 120^∘ phases, two distinct ρ's, two distinct readout objects — P_Y^{RRHF} and P_{vis}^{curv} and two distinct five-dimensional carriers, each conflated because of shared notation.

Relation	Shares the parity mechanism?
half-lift $\hat{h}^2=O_q$	**the same object** — the axle's own character doubling, not a second instance
ε^2, $J^2=-I$	no — a representation-**type** invariant (Frobenius–Schur), with no parity obstruction
Born	ψ
$N_g L_g^2=\text{const}$	no — a capacity balance, and graded a **postulate** in the generator-provenance audit
η^2 attenuation	**unestablished** — the audit records that multiplicativity of the two filtrations is itself unproved
$\sigma_F=\Lambda_3^2$	no — dimensional

The test that would extend the mechanism is uniform and cheap: for each candidate, exhibit the group, show the lower power is forbidden by it, and **state the bilinear type of the survivor**. The last requirement is not pedantry — AX-1.4 shows two cases that pass the first two tests and still are not the same mechanism.

Corollary AX-1.4 — the Born case: parallel in shape, opposite in content

The corpus's *Double Square Rule* derives $P(A)=|\psi_A|^2$ by excluding selection on individual paths and admitting selection on **pairs**. That has the same outward shape as AX-1: linear forbidden, quadratic first admissible, uniqueness then argued. **The accurate comparison is more interesting than either a dismissal or a unification.**

	Axle (this paper)	Born / Double Square Rule
shape of the argument	linear forbidden, quadratic first-allowed	linear forbidden, quadratic first-allowed
group doing the forbidding	$\mathbb{Z}_2$ half-turn parity	U(1) phase invariance
status of the exclusion	**theorem** from the involution	**also a theorem** — the companion *Physical Necessity of Quantum Probability* derives A7 from gauge invariance, normalisation preservation and factorisation (its Thm 3.1), so both exclusions are now theorems
bilinear type of the survivor	**holomorphic**, z^2	**sesquilinear**, $\psi^*\psi$
effect on the source phase	**doubles** it	**cancels** it
what that achieves	carries the CP orientation	makes the probability real and positive

Verified directly: under $z\mapsto e^{i\theta}z$, $\arg z^2$ shifts by exactly 2θ while $|z|^2$ is unchanged.

The two are therefore dual rather than identical, and the duality is the point. The axle *needs* the phase to survive and double — that doubling **is** the generation character $e^{i\pi/3}\mapsto e^{2\pi i/3}$, and hence the CP content. The Born rule *needs* the phase to cancel — that cancellation is what makes a probability real. They are the two ways a quadratic can treat a phase, and each construction requires the one the other forbids.

What is genuinely common, and it is worth stating: in both cases the linear term is excluded by a symmetry acting on the source, and the quadratic is the first admissible power. That recurring *shape of argument* is real. What does **not** transfer is the object it produces — holomorphic against sesquilinear, phase doubled against phase cancelled. Recording the shape as a shared heuristic is legitimate; recording the objects as instances of one mechanism is the conflation F7a forbids.

Corollary AX-1.5 — the shared method is Schur, and the exclusions are not symmetric

Two refinements follow from reading the companion papers, and the second is a gap on this side.

(a) The uniqueness arguments are the same method. Both are one-dimensionality computations of an equivariant Hom space: $\dim \mathrm{Hom}_{D_6}(\mathrm{Sym}^2E_1,E_2)=1$ here, and "the only U(d)-equivariant map $\mathbb{C}^d\to\mathbb{C}^d$ is a scalar, so the invariants are generated by $c^\dagger c$" there. **That is a genuine shared structure, and it sits at the level of uniqueness rather than exclusion** — a more precise location for the recurrence than "squaring".

(b) The $\mathbb{Z}_2$ forbids *every odd power*, not only the linear one. Computing the full tower on the wheel:

(tower table in Appendix A.0.2: 0, 1, 0, 2, 0, 2, 0, 3 for k = 1…8)

The reason is weight bookkeeping: E_1 carries weight 1 and E_2 weight 2 in units of π/3, while Sym^kE_1 carries weights k,k-2,…,-k; the target weight is reachable only for even k. **So AX-1 generalises from "linear is forbidden" to "all odd powers are forbidden"** — a strictly stronger statement, proved by the same involution.

But the tower does not terminate, and that is an asymmetry with the Born case. There the exclusion is two-sided: U(d) invariance forbids **all** k≠2, so the quadratic is isolated. Here parity supplies only a **floor**. Nothing in the D_6 analysis excludes a quartic channel — indeed dim Hom($\mathrm{Sym}^4 E_1$,E_2)=2.

Open, and newly identified: is the axle the *only* even channel the action installs, or the lowest of several? Every statement in this paper concerns the quadratic term. If the master action also returns a quartic coupling, it would contribute to the same 1↔3 entry and the "unique axle" language would need qualifying to "unique at lowest admissible order". The Born side has an upper-bound theorem; this side has not attempted one.

§16. The quadratic axle, unique at lowest admissible order

Theorem AX-2 (uniqueness of the quadratic axle at lowest admissible order)

dim Hom_D_6($\mathrm{Sym}^2 E_1$, E_2) = 1, with representative $Q(x, y) = (x^2 - y^2, 2xy)$, i.e. $Q(z) = z^2$.

Scope (F7b). This is uniqueness *within the quadratic sector*, which is the lowest admissible order by AX-1.1. It is **not** a statement that the quadratic is the only even channel: dim Hom($\mathrm{Sym}^4 E_1$, E_2) = 2, and no upper bound has been established (Cor AX-1.5b). Nor may the gap be waved through on universality grounds: the quartic enters the same entry with the conjugate character and moves the observables directly (companion §G.6.4).

Proof. $\chi\_\mathrm{Sym}^2E_1(g) = ½(\chi(g)^2 + \chi(g^2))$ gives (3, 0, 0, 3, 1, 1) on the six classes, and $\langle\chi\_\mathrm{Sym}^2E_1, \chi\_E_2\rangle = 1$ exactly. Equivariance of $z \mapsto z^2$ is direct: under r acting by $e^{(i\pi/3)}$ on E_1, $z^2 \mapsto e^{(2i\pi/3)}z^2$, which is the E_2 action; under reflection $z \mapsto \bar{z}$, $z^2 \mapsto \bar{z}^2$. Verified for all six rotations. ■

The map is even under the half-turn, so opposite rim points contribute coherently rather than cancelling — which is exactly why second order succeeds where first order cannot.

Theorem AX-3 (the descent and its democracy)

For the normalised action-selected mode,

$Q(u) = ⅓(1, e^{(2\pi i/3)}, e^{(-2\pi i/3)})$, $\|Q(u)\| = 1/\sqrt{3} = 0.577350269190$,

with component moduli all exactly 1/3, **democracy spread 1.11×10^{-16}**, and $\Sigma_a Q(u)_a = 0$ confirming the image lies in E_2 rather than the trivial summand. After normalisation every branch component has modulus $1/\sqrt{3}$. ■

The threefold sharing is therefore not merely imposed inside a minimal flavour cell: it is the unique lowest-order image of the void's sixfold orientation. This is the axle's strongest and least ambiguous contribution.

§17. Character descent and the source-phase candidate

Corollary 6.1

Under the preferred wheel rotation the selected void half-lift carries character $\chi_{void} = e^{(i\pi/3)}$; the quadratic axle returns $\chi_{gen} = \chi_{void}^2 = e^{(2\pi i/3)}$, the forward C_3 generation character; and with $\Omega = i\varepsilon H$ the frame map gives

$\Phi_{src} = i\cdot e^{(i\pi/3)} = e^{(5i\pi/6)}$.

This is the derived **source-phase candidate**, not arg ζ_{phys}. The physical branch coefficient has not been shown to inherit it (companion §G.6.11).

What this genuinely achieves. The half-lift relation $\hat{h}^2 = O_q$, previously adopted as phase grammar and left owed after the complement-typing failure, is now *derived* as the statement that the unique equivariant descent is quadratic. That is a real answer to a long-standing debt, and it does not require the anti-linear complement at all.

§18. Qualifications established here

§18.1 The descent lands on a non-trivial isotype

$Q(u)$ is an exact R_3 eigenvector with eigenvalue $e^{(\mp 2\pi i/3)}$ — verified, eigenvalue ratio constant across all three components. **It occupies the ω or $\bar{\omega}$ isotype, not the trivial isotype $k_{src} = 0$ assumed by the earlier minimal-cell work.**

The amplitude survives this relocation only because Theorem 2 makes equidistribution isotype-independent. The **phase does not survive it automatically**: the six-branch census (Appendix C) shows the isotype label moves arg ζ by ±120° among {150°, 30°, 270°}. The earlier $k_{src} = 0$ statement must therefore be corrected to record that the void descent selects a non-trivial character, and any downstream argument that used trivial-isotype support must be re-examined.

§18.2 The phase-carrier typing — source candidate derived; physical mechanism and value unresolved

Corollary 6.1 does not state which object inherits the phase, and the two defensible readings disagree.

Reading	Hermitian branch coefficient h	arg ζ	Depends on the visible branch?
(i)	inherits the mode's rim-to-rim phase advance $u_1/u_0 = e^{(i\pi/3)}$	**150° for every branch**	No
(ii)	is the visible component of the descended vector Q(u)	**150°, 30° or 270°**	**Yes**

Under reading (ii) the three components of Q(u) carry arguments (0°, +120°, −120°), so with the overall convention phase fixed such that one branch returns 150°, the other two return 30° and 270° — **reproducing exactly the three principal entries of the frozen six-branch census**, with consequences:

Visible branch	arg ζ	J	‖V_ub‖	γ	χ²_all
the one giving 150°	150°	$+2.9968\times10^{-5}$	0.003557	73.02°	**11.3**
+120°	270°	-4.6562×10^{-6}	0.003372	−9.08°	9259.7
−120°	30°	$+2.9974\times10^{-5}$	0.006541	31.77°	3000.7

The entire numerical phase result turns on this typing choice, and under reading (ii) it is not independent of the unreturned readout P_vis^curv. By Rule R4 the paper must declare which object carries the phase and justify it from the action. This is the same class of question as the two prior corrections in this series, and it is now the load-bearing step.

Note the internal tension: §18.1 establishes that the descent carries a non-trivial character, which is naturally a statement about the *vector* Q(u) — favouring reading (ii). Reading (i) requires an independent argument that the branch coefficient tracks the mode rather than its image.

§18.3 The raw norm is not a second determination of the scale

The statement that 1/√3 is "supported in two independent ways" — C_3 equidistribution and the axle norm — is overstated. Both are the same combinatorial factor √3 arising from three branches; and since Q is quadratic, the raw norm ‖Q(u)‖ = 1/√3 is tied to the choice ‖u‖ = 1 and carries no physical scale. **c_ax remains the only scale**, exactly as §19 states. What the axle independently supplies is a second derivation of the *democracy*, not of the amplitude.

§18.4 Two selections remain inherited

The chirality of the mode (e^(+iπ/3) versus e^(−iπ/3)) and the half-lift root (e^(iπ/3) versus e^(−2πi/3)) are each one bit, and each moves the answer out of agreement. They are attributed to PFDMI-EX2G's preferred circulation and to the action-to-frame type order respectively. Both are inherited claims about other returns, and both carry their own provenance obligations; neither is established within this document.

§18.5 Graph-mode typing firewall

The canonical wheel has several numerically simple invariants, but they belong to different carriers:

Object	Dimension/count	Role here
Wheel vertices	7	closure-state catalogue
Wheel edges	12	admissible transition links
Cycle space	6	independent closed graph loops
Perron eigenspace of $\hat{T}$	1	stationary coherent sector
$k=\pm1$ eigenspace	2	oriented sixfold source E_1, $L=2$, parity $-I$
$k=3$ alternating mode	1	also antipode-odd — **not** E_1; why parity alone cannot identify the source
$k=\pm2$ eigenspace	2	generation target E_2, $L=4$, parity $+I$, **in ker(I-1/4L)**; membership in the kernel of the unresolved $\hat{T}$ is unproved
Antipodal-axis carrier	3 components	= trivial $\oplus$ E_2 exactly (§14.1.2)
Normalised branch image	$1/\sqrt{3}$ per component	representation-geometric democracy

The CKM axle uses the **two-dimensional k=±1 eigenspace**, not the one-dimensional Perron sector and not a quotient of the twelve edge links. Conversely, the charged-lepton proposal that relates twelve wheel edges to a terminal maintenance fraction is a different typed construction. Neither its 1/12 approximation nor its refined $(7/6)e^{(-8/3)}$ factor supplies c_ax, ρ_mult^physical, a CKM phase, or a CKM branch coefficient. Any such transfer requires a new action-level intertwiner and is prohibited by R5 until supplied.

§19. Norm transfer and the physical amplitude

The axle fixes the **relative geometry** of the curvature, not its overall physical coefficient. Let c_ax be the coefficient with which the master action inserts the unique quadratic channel into the role-even weak-doublet Hessian. Then

$\zeta_{phys} = \rho_{mult}^{physical} \cdot (c_{ax}/\sqrt{3}) \cdot e^{(i\theta_{phys})}$.

The source-character descent supplies the candidate $\Phi_{src} = i\cdot e^{(i\pi/3)} = e^{(5i\pi/6)}$, but the physical action has not shown that the visible branch coefficient inherits that source phase.

If the action proves c_ax = b **and** ρ_mult^physical = 1, then $|\zeta_{phys}| = b/\sqrt{3}$. **That identification is a physical norm-transfer statement and must not be attributed to symmetry alone.**

Proposition 7 (conditional physical amplitude)

If the master action places the unique quadratic channel in the physical role-even Hessian with an action-returned coefficient c_{ax} and retention $\rho_{mult}^{physical}$, and the readout preserves the surviving branch coefficient, then $|\zeta_{phys}| = \rho_{mult}^{physical} \cdot c_{ax}/\sqrt{3}$ — equal to $b/\sqrt{3}$ only under the Stage I-B comparison values $c_{ax} = b$ and unit retention. Exact conditional transport; the owed objects are the action-returned coefficient and its retention. ■

---

§F. Falsifiers governing the claims of this paper

The complete ledger is in the companion (§31). These four govern the claims made here.

ID	Fires if
F4	Any value is selected because it improves a measured quantity — including adopting the 150° branch on empirical grounds, or letting a desired value appear in a derivational question.
F5	A normalisation is quoted without naming which of the four of companion §20 it is; or a test condition is written with $1/\sqrt{3}$ in place of $\rho_{mult}^{physical} \cdot (1/\sqrt{3})$.
F7b	The axle is called the *unique* channel without "at lowest admissible order", or the even tower above it is dismissed as excused by universality.
F7g	A branch is selected by taking the orthogonal complement of an "inherited" support that omits $\mathfrak{B}_{23}$ — the (2,3) entry $b = 0.0405$ is nonzero, so that complement is empty.

---

§Final. Conditional charged-flavour derivation theorem

Under the frozen RRHF premise package, the representation-resolved parent action returns by blind differentiation a positive rank-nine charged tensor $T_{ch}=Y_u \oplus Y_d \oplus Y_e$ at two regulators, with maximum post-freeze tensor residual 7.070×10^{-14} and independent CY3-prime residual 1.937×10^{-14}. The left singular frames return $V_{CKM}=U_u^\dagger U_d$, with level-4 modulus

$$|V_{CKM}|=\begin{bmatrix} 0.974818 & 0.222964 & 0.004198 \\ 0.222926 & 0.973988 & 0.040639 \\ 0.005895 & 0.040427 & 0.999165 \end{bmatrix}, \quad |J|=1.86032934918\times10^{-5}.$$

No measured Standard Model target is accessible to the frozen execution core during the run, and none is used to choose a branch or coefficient before the output is sealed. The exact D_6 axle theorem identifies one possible curvature refinement but does not install it; the curvature-corrected matrix is therefore an unadmitted diagnostic rather than part of the derivation here.

CONDITIONAL CHARGED-FLAVOUR TENSOR DERIVATION: PASS QUADRATIC AXLE FORM AND DEMOCRACY: EXACT AT LOWEST ORDER QUADRATIC-

ONLY PHYSICAL TRUNCATION: MINIMALITY PRESCRIPTION LOCAL Q-SECTOR INTERTWINER FORM: DERIVED FULL PHYSICAL INTERNAL LIFT AND AXLE COEFFICIENT: OPEN PHASE: THE JET IS COMPONENT-PRESERVING, BUT THE PHYSICAL READOUT IS NOT DERIVED AND THE LOCAL METRIC IS EXACTLY PHASE-ISOTROPIC 150°, 30° AND 270° ARE ALL PRESENTLY UNRETURNED PHYSICAL BRANCHES CURVATURE-CORRECTED CKM AT 150°: A QUARANTINED DIAGNOSTIC ON A BRANCH THE DERIVATION DOES NOT RETURN CURVATURE-CORRECTED CKM UNDER THE PROVISIONAL P_{23} BRANCH: WORSE THAN NO AXLE (χ^2_5 = 539 OR 1032 AGAINST 291) NEXT BLOCKER: RETURN A NONZERO SAME-ACTION COEFFICIENT FOR THE ALREADY IDENTIFIED GENERATION-SENSITIVE CHARGED-COMPLETION MIXED-JET FORM; THEN ITS PHYSICAL READOUT, MULTIPLICITY AND COMPARISON-LOOP HOLONOMY UNCONDITIONAL STANDARD MODEL ADMISSION: NOT CLAIMED

What this status settles, and what it does not. The earlier conclusion that no weak-doublet bridge exists in the programme is withdrawn and wrong (companion §G.6.8). Equally, the remaining problem cannot be closed by further wheel or flavour algebra: **the source, the map Q, the target E_2, the democracy and now the Q-sector internal intertwiner are all settled** (companion §G.6.9).

The blockage has since been located exactly, and it is not a missing calculation. The terminal current-and-record action **tested for this correction** is **exactly blind to rotations in family space**: a Q-family rotation commutes with every gauge generator and with all seven terminal projectors at residual zero, so no committed weight moves and 42 post-gauge physical flat directions remain. No amount of further diagonalisation of the existing endpoint, current or selector arrays can return c_ax, P_vis^curv, ρ_mult^physical or θ_phys, because none of them is an observable of *that tested action* — a statement about the tested record action, not about the cumulative framework. **What is required is one new generation-sensitive same-action completion coupling** — a mixed jet contracting the unique Q-sector E_2 support plane with the charged-completion tensor. Until that exists, selecting c_ax = b or the 150° branch would remain an insertion.

Axle qualification. RRHF-AX1 evaluates the unique quadratic axle on the RRHF regulator ladder and returns $|J|=3.024\times10^{-5}$ with χ^2_5=4.10 — **on the 150^∘ branch, which companion §G.6.10 shows the derivation does not return; at the provisional P_{23} branch the same construction gives 539 or 1032.** Its **form and threefold democracy are exact**, and its **magnitude follows conditionally** from the inherited norm b. Its **quadratic character descent and source-phase candidate are derived**; the physical carrier, comparison holonomy and phase value are not returned. The **completion jet** is component-preserving, but the physical readout selecting the component is not derived (companion §G.6.10); the remaining unresolved step is therefore the coefficient and the physical branch readout, consequently the specific 150^∘ execution is not yet the programme's controlling conditional return. A full projected-H_CMR calculation must first return the axle coefficient, internal lift, readout and phase carrier (Stage I, §7.1), after which RRHF must reproduce the resulting charged tensor (Stage II, §7.2). An earlier grading of AX1 as a free additional constitutive premise is **withdrawn** (§5.2.2).

Empirical qualification, carried in the verdict rather than a footnote. The controlling conditional return matches the Cabibbo sector and the $2\leftrightarrow3$ magnitude but misses $|V_{td}|$ by -13.5σ and J by -10.1σ, with $\chi^2\approx291$ over five entries (§6.1). *Derivational* pass and *empirical* success are separate gradings, and only the first is claimed here. The unadmitted refinement would bring both quantities inside 1.5σ, which is precisely why the I6a protocol is sealed in advance and why any later pass must be reported alongside §B.1.

Stage I-A has been executed and does not pass (§7.0). It closes the source-mode gate — the action-selected mode *is* the canonical E_1 harmonic, verified at overlap 1.0 — and confirms on that actual mode that the unique quadratic descent reaches the same-wheel E_2 target with exact democracy. It returns neither a nonzero c_ax nor the complete full-carrier lift, readout or phase carrier; **the local Q-sector intertwiner form is derived**. **The obstruction is underdetermination rather than unfinished computation**: the EX2F endpoint profile is rank one with leading singular value exactly $\sqrt{7}$, 836 of 920 directions are post-gauge flat, and the EX3B embedding family $\varepsilon(\theta)=\sin\theta$ sweeps fail-to-pass continuously. **The critical path is therefore the full-carrier physical intertwiner and readout calculation — not further work on the wheel or on the axle's mathematical form.** Concretely, the next calculation must break the rank-one degeneracy of the endpoint profile and select one member of the $\varepsilon(\theta)=\sin\theta$ family; the axle's source, form and democracy require no further work at all.

And that calculation is not a flavour calculation (§7.0.1.1), **and its invariant parameter budget can now be counted**: the companion's Appendix G shows the invariant parameter budget on the minimal carrier is at least 22 real numbers — 6 quadratic, 13 cubic, 3 in the reference kernel — against which the certified corpus supplies one equation and one inequality. c_ax is one coordinate of the 13-dimensional cubic space, and its existence is generic **in the unrestricted family**, though the tested action is not a generic point of that family (companion §G.6.9); what is codimension-one is c_ax=b. The programme ledger records the same non-uniqueness for the complete H_CMR — two inequivalent score lifts with identical central traces and symmetry data — and lists among its unreturned objects precisely the realised complex amplitude and the relative phases, which are this chapter's c_ax and θ_phys. **The flavour blocker is the H_CMR blocker seen from the flavour sector.**

That claim now needs splitting, because part of H_CMR has closed and it is not the relevant part. A subsequent revision upgrades the two completion principles from adopted to derived — RSU from the single-source theorem, non-contextual weight and operational distinguishability; CPR from the record definition itself, since a record that fails to survive a causal split fails persistence, re-accessibility and composability simultaneously. **That upgrade is a real reduction of the premise debt and is accepted here**, with two qualifications recorded in §7.0.1.2. But it does not change the parameter count at all: a principle moving from *assumed* to *derived* alters its status, not the number of couplings the action leaves free. HCMR-XI-1 (2 August 2026) reports a conditional law-level completion under two adopted principles — Record-Sufficient Update and the Causal-Persistence Requirement — which turn the eleven free record-conditioned tilts into integer record multiplicities and fix the copy question by a minimal two-copy code. Independently re-executed: the completed path cost

$$M[X]=\Sigma\_n\ D\_KL(K\_n\|T_0)-\alpha\_\star\Sigma\_n\sigma\_n\langle\omega\_C\rangle\_K\_n-\Sigma\_e\in\rho\_n\ln\ p\_e(\eta\_n)+I\_hard$$

carries **no free coupling constants at all** — every relative weight is unity, ∞ or the one-Bit root — which is a substantial gain over eleven phenomenological tilts.

But the flavour blocker does not live in that sector. the companion's Appendix G locates it:

Ambiguity	Closed by HCMR-XI-1?
eleven record-conditioned tilts	**yes**, conditionally, via RSU
copy / maintenance law	**yes**, conditionally, via CPR
reference kernel T_0 (3 invariant parameters)	**no** — M is *defined relative to* T_0
D_6-invariant quadratic couplings (6)	**no** — M contains no wheel Hessian term
D_6-invariant cubic couplings (13), **where c_ax lives in primitive coordinates** — the same object is a *second* variation in descended branch coordinates (companion §G.6.8)	**no** — not addressed

Corrected statement. The record-conditioning sector of H_CMR has conditionally closed. The coupling sector has not. **The flavour blocker tracks the second**, so it is untouched by this completion, and c_ax and θ_phys remain exactly where Stage I left them.

§7.0.1.2 Two qualifications on the derivation upgrade

(a) Composition alone does not fix the likelihood exponent. The RSU argument excludes an arbitrary multiplier f_ρ(η) on four grounds, then concludes the likelihood is the chain product. But multiplicativity is preserved by every power: $(L_2 L_1)^\beta = L_2^\beta L_1^\beta$ identically, verified for β=1/2,1,2,3. So the stated grounds exclude arbitrary tilts and leave a **one-parameter family** L^β. What fixes β=1 is the single-record normalisation — that at N=1 the likelihood must *be* the Born probability. That step is available in the corpus's own universal Born result but is not cited in the derivation as written, and the argument is incomplete without it.

(b) The V-metric refinement may enlarge the undetermined set rather than reduce it. The proposal to replace the Dirac adjoint with $A^‡=V^{-1}A^† V$ rests on a correct quasi-Hermiticity result: for a PT-symmetric H there exists V with $VHV^{-1}=H^†$ and a conserved V-norm. Two problems with importing it here.

First, **VERSF's H_cl is the second variation of a real functional and is therefore symmetric by construction**; no reason is given why the physical metric should not be the Dirac one, so the refinement currently answers a question the corpus has not raised.

Second, and more seriously, **V is not unique.** For quasi-Hermitian H the metric operators form a family, and fixing one requires specifying additional observables. Verified explicitly: for a simple non-Hermitian H with real spectrum, a one-parameter family of valid V satisfies

$VHV^{-1}=H^T$ exactly. **Unless V is itself returned by the action, adopting it adds an undetermined object to a programme whose central difficulty is already underdetermination.**

Neither qualification touches the CPR derivation, which is sound: on a commutative algebra with a distinguished pointer basis, bilateral recoverability plus coassociativity force $\Delta|i\rangle=|i\rangle|i\rangle$ — verified — and no-cloning is not violated because the object copied is an already-committed classical record in an orthogonal algebra.

§7.0.1.2.1 The V-refinement is closed negatively — it does not exist

HCMR-XI-2 settles §7.0.1.2b more decisively than the objection raised there. The oriented D36 history kernel has four **nonreal** eigenvalues in conjugate pairs. If $T^\dagger V=VT$ with V positive definite, then $V^{(1/2)}TV^{(-1/2)}$ is Hermitian and T has real spectrum; contrapositively, **nonreal eigenvalues exclude any positive quasi-Hermitian metric on a single oriented branch.**

So the refinement is not merely non-unique, as §7.0.1.2b argued — **it is unavailable**. Probability preservation is instead carried by the completely positive marked instrument, whose Kraus completeness passes. Debt A5c is closed negatively.

§7.0.1.2.2 Microscopic execution: three passes and one decisive failure

HCMR-XI-2 moved RSU and CPR from formulae into the actual matrices. Independently checked where the arithmetic permits.

Object	Return	Check
marked instrument on $\mathbb{C}^7\otimes\mathbb{C}^{50}$, 319 marks	completeness 2.47×10^{-14}	
RSU chain identity p(α_2,α_1)=p(α_1)p(α_2	α_1)	1.39×10^{-17}
CPR copy isometry, 101761×319	‖C† C-I‖_F=0	$319^2=101761$ ✓
Wilson Fisher score	72×72, full rank, $\kappa=6.151487\times10^6$	λ_max/λ_min reproduces ✓
record-completed parent	rank 101, one normalisation null	drifts $7.51,1.85,0.46\times10^{-5}$, ratios 4.07 and 4.01 — **clean second-order convergence**; Richardson limit λ_min→1.81881×10^{-3}
same-action HM4 gauge reproduction	**FAIL**	see below

The Wilson score reaching full rank 72 is a genuine advance over the earlier primitive derivative that returned zero.

§7.0.1.2.3 The failure is misdiagnosed, and the correct diagnosis is worse

The report attributes the gauge mismatch to a **stiffness anisotropy**: the source instrument returns relative generator stiffnesses (1, 1.0758, 2.7336) where HM4 is isotropic (1,1,1). **That cannot be the main effect.**

If the two matrices differed *only* in those three stiffnesses within a shared eigenbasis, the Frobenius cosine would be $\cos=\Sigma\, d_i s_i/\sqrt{(\Sigma\, d_i s_i^2 \cdot \Sigma\, d_i)}$. Evaluated across every plausible split of the 72-dimensional quotient:

sector split	predicted cosine
(48,18,6)~8{:}3{:}1	0.926
(56,12,4)	0.942
(36,24,12)	0.900
(24,24,24)	0.895
(64,6,2)	0.965

The observed cosine is 0.115. The two reported diagnostics are internally consistent — $\sqrt{(1-0.993354^2)}=0.115103$ reproduces the quoted cosine exactly — so the number is right; it is the interpretation that is wrong.

The mismatch is not an anisotropy of eigenvalues. It is a near-orthogonality of eigenstructure. A pure stiffness disagreement would leave the matrices about 90% aligned; they are about 11% aligned.

This changes what the next calculation must repair. A stiffness mismatch would call for a term generating generator-dependent rigidity. **Near-orthogonality means the current endpoint/current instrument is not a rescaled version of the right object but a different one**, and the search should be over operator structure rather than over stiffness coefficients.

§7.0.1.2.4 The gauge-shape repair, and the one question it leaves

HCMR-XI-3 repairs the §7.0.1.2.3 failure, and does so by attacking structure rather than coefficients — which is what the corrected diagnosis required. Its construction is a minimum-relative-entropy oriented-mark law

$p_{a,\sigma}(X)=\exp(\sigma\sqrt{m}D_a(X))/2\Sigma_b\cosh(\sqrt{m}D_b(X))$,

whose Fisher information at the frozen background is exactly $I_X=J_D^TJ_D$. **Verified independently** by finite differences on a random residual derivative: relative residual 5.6×10^{-11}, and the mark probabilities sum to 1.000000000000000. The $\sqrt{m}$ normalisation is precisely what removes any leftover constant — a design feature, not a fitted coefficient.

With $D_g(A)=Z^{(1/2)}P_WA$ the gauge pullback is P_WZP_W, and the reported comparison against the stored HM4 array gives relative residual 4.485×10^{-16} with cosine 1.000000000000000.

Two observations, one supporting and one cautionary.

Supporting. The prior object sat at cosine 0.130631 after optimal rescaling, and the repair moves it to unity. That is a shape closure, not a normalisation — and it independently confirms the §7.0.1.2.3 diagnosis, which predicted that a stiffness fix could not work. The (48,18,6) split inferred there is also confirmed by Z's multiplicities, and 48+18+6=72 is the Wilson quotient exactly.

Cautionary, and it is the same question this paper has asked twice before. **A residual at machine zero admits two readings with opposite value:** an exact structural identity has been found, or the two objects are the same and the comparison is a reproduction. The corpus records Z=(1,5/2,298) as the reduced gauge response obtained as a Schur complement of the gauge block. **If that Schur complement was taken of the same quadratic form the HM4 array stores, then P_WZP_W reconstructs H_HM4 by construction**, and 4.485×10^{-16} is what one should expect with no physics involved. Relatedly, "the physical spectrum is exactly the source-response spectrum" is a tautology of P_WZP_W rather than a finding.

The stated firewall does not settle this. It certifies that the HM4 *array* was not opened during construction; it does not certify that Z, which **was** used, is independent of that array's contents.

What would settle it, in one line from the source chain: Z is independent if and only if it is returned by a calculation that never differentiates, inverts or Schur-reduces the HM4 gauge quadratic form. If it does descend from that form, the correct grade is **reproduction**, and XI-3's genuine new content is the marked-instrument construction rather than the gauge agreement.

What is new regardless of how that resolves: the exact Fisher identity from a minimum-relative-entropy law; one global Gibbs normaliser over 402 marks, so six sectors are typed outcomes of a single instrument rather than separately normalised models; C_comp annihilating both q_Y and q_B-L exactly, so the anomaly-admissible charge plane survives; the charged sector reproducing its own Hessian at 1.33×10^{-15} with bosonic/CAR agreement at 5×10^{-16}; and the neutral sector kept **branch-resolved** rather than selected, with Loewner incomparability still controlling. Shape checks pass: 402=2×201 and 402^2=161604.

§7.0.1.2.5 HCMR-XI-4: three gates close, one no-go retained

Neutral branch — Dirac selected. The claim is M_ν=S_DU_D† at 8.41×10^{-16} with S_D=√(M_ν M_ν†), the Majorana alternative failing at 1.57×10^{-1}. **A caution of the A5e kind applies and is the only one that matters here.** Every matrix admits a left polar decomposition, so the content is not that the factorisation exists but that the *independently returned* Dirac strain coincides with the polar modulus while the Majorana one does not. **That is a genuine discriminator if and only if S_D was not itself defined as that polar modulus upstream.** If it was, 8.41×10^{-16} is a tautology. Recorded as debt A5f.

Minor: the branch separation is quoted as "approximately 5.7%". The three eigenvalue pairs actually differ by **3.87%, 4.72% and 5.71%** — the quoted figure is the largest, not a common value, and should be stated as a range.

Charged alphabet — a real reduction of freedom, correctly sized. Verified: with $r_L=\sqrt{\kappa}\, x_L+\kappa^{(-1/2)}T_Rx_R$ and $r_R=(\kappa I-\kappa^{-1}T_R^TT_R)^{(1/2)}x_R$, the identity $\|r_L\|^2+\|r_R\|^2=x^TH_{ch}x$ holds exactly. This is the block LDL completion of the square, **unique given the L-before-R ordering**. So the alphabet is now fixed up to *block ordering* rather than up to a full O(36) rotation — a genuine reduction over XI-3's admitted residual freedom, and it should be claimed at exactly that strength.

An internal consistency check that does not quite close. The charged spectrum [1.027883126507, 2.972116873493] is symmetric about $\kappa=2$, confirming κ and giving $\sigma_{max}(T_R)=0.972116873493$. Then the smallest right-complement eigenvalue must be $\kappa-\sigma_{max}^2/\kappa=1.527494392135$. **XI-4 reports 1.527103739469** — agreement to three significant figures, a 0.026% discrepancy. Either the two numbers come from different regulator levels, which should be stated, or one needs rechecking (debt A5g).

Completion instrument. Completeness is exact by construction, since A is block diagonal with respect to the P_r so $\Sigma_r P_r A^\dagger A P_r=A^\dagger A$; the residual 7.62×10^{-15} confirms the implementation rather than the structure. The interval bound is real and honestly stated: $K_\emptyset$ requires $\sin^2 2\varphi_H\leq 1/\lambda_{max}(A^\dagger A)$.

A structural relation the report does not draw. With $\nu_F=(21/16\pi)(c/\xi)$, all four quoted numbers reproduce to twelve figures, and against the existing clock ledger

$\tau_F/\Delta t_0=2.000000000000$, $T_0/\tau_F=7.000000000000=K$.

The Fold period is exactly two elementary steps and the full cycle exactly K=7 Fold periods. That is a clean consistency with the independently frozen clock ledger and is worth recording.

The scale no-go, stated precisely. The arithmetic checks: 9×11 of rank 9 has nullity 2; 10×13 of rank 10 has nullity 3. But note what happened — **adding the continuum-current ratio introduced one equation and two unknowns, so the nullity rose from 2 to 3.** The added relation is therefore net-negative for fixing the scale rather than a strengthening of the no-go, and the report should say so. The no-go itself stands, consistent with ASAC-1.

§7.0.1.3 The budget is unchanged by the upgrade

M has no free coefficients *given its term structure*. The derivation upgrade strengthens the case that the term structure is forced, which is genuine progress on the **premise** debt. It leaves the **parameter** debt untouched: T_0 is still a three-parameter choice, M still contains no wheel Hessian term, and the six quadratic and thirteen cubic invariant couplings — including c_{ax} — are still unaddressed. If the V-refinement is adopted without deriving V, the budget rises rather than falls.

The decisive standard. To move beyond this paper, one sealed action must return the whole chain:

** $H_{CMR}\longrightarrow(E_1,Q,c_{ax},\mathbb{I}_{int},P_{vis}^{curv},\theta_{phys})\longrightarrow\Omega_0\longrightarrow Y_u,Y_d,Y_e\longrightarrow V_{CKM}$ **

with **no measured CKM quantity entering before the first arrow is frozen**. The strongest current grade is: charged-flavour derivation — conditional action-level pass; quadratic axle form and democracy — exact; curvature-completed CKM — conditional execution pass on a branch the derivation does not return; physical coefficient, embedding and readout — same-action return open; **quadratic character descent and source-phase candidate derived; physical carrier, holonomy and phase value open, physical value not** (companion §G.6.10).

The remaining primitive critical path was H_CMR‖ξ. **Both have moved, and unevenly.** ξ is conditionally selected at 84.137117μm, with the former 0.078% discrepancy resolved as a UV-ruler artefact rather than a second physical scale. H_CMR has conditionally closed **in its record-conditioning sector**, removing eleven free tilts. A later revision upgrades RSU and CPR from adopted principles to derived consequences, which is a real reduction of the **premise** debt. A microscopic execution then passed the RSU instrument, the CPR copy isometry and the record-completed parent at machine level while **failing** to reproduce the HM4 gauge Hessian from the same action — a failure whose diagnosis is corrected in §7.0.1.2.3 from stiffness anisotropy to near-orthogonality of eigenstructure. **None of it touches the coupling sector**, which is where c_ax and θ_phys live (§7.0.1.1, the companion's Appendix G).

The target has since been corrected, and the correction cuts both ways. A Standard Model derivation needs one physical universality class, not a literally unique microscopic action — actions differing by gauge choice, field redefinition or irrelevant operators are acceptable, as many lattice actions give one continuum theory. Applied to the two sectors this gives opposite answers: the gauge counterfamily is Symanzik-type improvement and plausibly flows to the same infrared theory, while **the flavour tower changes the observables directly**, since the quartic channel enters the same 1↔3 entry, carries the conjugate C_3 character and therefore moves the phase. A 30% relative quartic source amplitude gives a 5% norm shift and roughly a factor-three variation in the diagnostic χ^2; a 5% source amplitude gives only a 0.83% norm shift (companion §G.6.4). **The reframing improves the gauge grade and leaves the flavour grade where it was.** A later correction adds that an external-time equilibrium RG cannot settle a framework whose time is generated by the commitments being coarse-grained. That is accepted, and it bears on whether the quartic *survives* coarse-graining rather than on what it *does* once present — and the second is what was computed, by direct evaluation rather than by any RG. Its practical effect is to **name the test that would close the question** rather than to reopen the finding (companion §G.6.5).

That test has since been run, and At this historical stage, A6 remained open it — but the successor calculation supplies a sharper route. RUDM-1 closes the preparation measure and the commitment projectors and reduces the analogous K_4 question to a single algebraic condition: whether the Fold-to-tangent intertwiner commutes with the twelve commitment projectors, since ‖P_AUz‖ = ‖P_Az‖ whenever it does. Verified, along with the exchangeable two-fact law and the sampling audit. **The transferable form for this chapter** is to ask not whether a quartic channel survives coarse-graining but whether it moves $\|P\_Az\|^2$ for any commitment projector at all — if it does not, it cannot change a committed record law at any coefficient (companion §G.6.7).

On the earlier test itself, At this historical stage, A6 remained open. ETRG-1 correctly withdraws the earlier "universality fails" claim and correctly relocates the decisive question to the commitment map. But its pre-temporal verdict is an identity holding for every symmetric transfer matrix rather than a finding about this model; its single-cell universality is forced by $\mathbb{Z}_2$ symmetry for every value of the coupling; and its two-fact dependence is the correlator conjugate to the coupling being varied. The informative comparison — a quantity not conjugate to K_4 — remains outstanding, and the effective-dimension objection is untouched by a two-dimensional calculation with commitment appended (companion §G.6.6).

PSU-1 then changed the kind of the obstruction, not its size. Three explicit counterfamilies — a third Fold-edge coupling, an infinite Wilson-harmonic family and an infinite completion-amplitude family — reproduce every stored low-order return while remaining physically distinct, and none is removable by the physical quotient. **The admissible action family is therefore infinite-dimensional**, and the finite ≥22 count survives only as a statement about the D_6-invariant quadratic and cubic slice. Two consequences follow for this chapter: **using the quadratic axle alone is a minimality prescription rather than a derivation**, on exactly the standard PSU-1 applies to the gauge sector; and the remedy for the flavour tower is the same RG fixed-point theorem, not a further flavour calculation (companion §G.6.2, companion §G.6.3).

The critical path for *this chapter* is therefore now stated exactly:

the D_6-invariant coupling sector: T_0 (3) + quadratic (6) + cubic (13), with c_ax one coordinate of the last. The I6a curvature test, the primitive generator-provenance audit and the orientation-installation work are parallel grading and promotion tests. **They are not additional admitted CKM routes.**

Appendix A — Numerical ledger

A.0 Controlling derivation (Part I)

Object	Level 3	Level 4
ε_T^(max) (post-freeze tensor residual)	7.068×10^{-14}	7.070×10^{-14}
ε_CY3'^(max)	1.934×10^{-14}	1.937×10^{-14}
bosonic / CAR derivative residuals	4.816×10^{-16} / 1.111×10^{-16}	4.350×10^{-16} / 3.332×10^{-16}
Fock functoriality residual	1.557×10^{-17}	5.715×10^{-18}
	J_CKM	
CKM modulus continuation		1.246729×10^{-3}
relative Jarlskog continuation		2.641964×10^{-2}
singular-spectrum continuation (u, d, e)		4.131280×10^{-4}
λ_min (bosonic parent Hessian)	1.027481351817	positive

Object	Level 3	Level 4
R_QCD	$0.580788332942I_3$	frozen, generation-scalar
Final controlling matrix	continuation input	**controlling return**

Independent checks on the level-4 quoted moduli: row-wise unitarity $\approx 7\times10^{-7}$ (six-decimal rounding); standard-parameterisation $|J| = 1.86114\times10^{-5}$, agreeing with the quoted value to 4.3×10^{-4} relative; δ=30.15^∘; departure from the bare e^(Ω_q) baseline +0.91% in $|J|$ and 4.14×10^{-4} in maximum modulus.

A.0.1 Stage I forcing-test certificate (§7.0)

Object	Value
overlap of the EX2F mode with the canonical k=-1 harmonic	1.0000000000000000
‖Lv-2v‖ (E_1 sector membership)	1.75×10^{-15}
hub amplitude / rim-amplitude spread	7.24×10^{-17} / 8.33×10^{-16}
phase step vs -π/3	-1.047197551196597
half-turn odd residual	$\le 1.2\times10^{-15}$
linear antipodal pair-sum	5.82×10^{-16}
‖Q(u)‖	0.577350269189626=1/√3
component moduli / democracy spread	1/3 each / 6.1×10^{-16}
‖LQ-4Q‖ (E_2 sector membership)	5.21×10^{-15}
C_3 character of Q(u)	e^(2π i/3), residual 9.0×10^{-16}
EX2F endpoint profile	7×50, **rank 1**, leading singular value √7=2.6457513110645907
EX2F nullities	local gauge 84, **post-gauge physical flat 836**, total 920
EX3B embedding family	ε(θ)=sinθ on [0,1], residual 1.87×10^{-15}
Stage I-A installation tuple returned	**incomplete**

Manifest hashes for the controlling derivation are to be added in a dedicated reproducibility appendix when available — the parent-action source, test contract, K=7 ledger, RCMO level-3 and level-4 operators, one-Bit history metric, FRGM transport operator and pre-comparison output directory state (§1).

A.0.2 Admissible-power tower (Cor AX-1.5b)

k	1	2	3	4	5	6	7	8
dim Hom_D_6(Sym^kE_1,E_2)	0	**1**	0	**2**	0	**2**	0	**3**

This is an L2 result and is scheme-independent — a direct evaluation through the same exponentiation used throughout, with no renormalisation step, so no choice of ensemble or coarse-graining affects it (companion §G.6.5). Sensitivity of the observables to a quartic admixture at relative source strength ε, scanning its unknown phase (pure quadratic gives 5.25; no axle at all gives 290.7):

ε	shift in |ζ|	χ^2_5 range
0.05	0.83%	4.72 – 5.85
0.10	1.67%	4.25 – 6.51
0.30	5.00%	2.99 – 9.87

with $\|Q_4(u)\|/\|Q_2(u)\| = 0.166667$ and the quartic carrying the **conjugate** C_3 character e^(+2πi/3) against the quadratic's e^(−2πi/3).

Odd k vanish by the $\mathbb{Z}_2$ parity; the surviving even tower is unbounded, so the quadratic is a floor and not an isolated point. Contrast the Born sector, where U(d) invariance excludes all k≠2.

Appendix B — Source-space envelopes

Source subspace	Basis	dim	χ²_all	χ²_curv
$\mathbb{R}^6$ unrestricted	none	6	5.8158	2.9503
$g_{31} = 0$	depth adjacency	4	5.8200	2.9534
σ-symmetric families	reflection-even	2–3	7.6354–7.6640	5.22–5.26
only g_{23}	adjacency + occupancy (forbidden)	2	7.8178	5.4539
only g_{12}	σ image	2	11.8470	6.9855
supplied target	clean candidate	0	11.3268	8.9816

Appendix C — Six-branch phase census

Isotype	Root	arg ζ	J	‖V_ub‖	γ	χ²_all
trivial	principal	150°	$+2.9968\times10^{-5}$	0.003557	73.02°	11
trivial	antipodal	330°	$+6.8890\times10^{-6}$	0.005646	8.06°	2323
ω	principal	270°	-4.6562×10^{-6}	0.003372	−9.08°	9260
ω	antipodal	90°	$+4.1501\times10^{-5}$	0.005759	55.41°	1531
$\bar{\omega}$	principal	30°	$+2.9974\times10^{-5}$	0.006541	31.77°	3001
$\bar{\omega}$	antipodal	210°	$+6.8949\times10^{-6}$	0.001325	36.20°	2330

The three principal entries are exactly the three values reading (ii) of §18.2 produces from the three visible-branch choices. **J is isotype-blind** (two branches agree to 1.97×10^{-4}); the discriminators are |V_ub| and γ.

Appendix D — Phase sensitivity around the returned direction

Post-freeze sensitivity, **not** target-source selection:

g_A/g_S	arg ζ	J	γ	χ²_all
0	90°	$+4.1501\times10^{-5}$	55.41°	1530.83
1.0	135°	$+3.4744\times10^{-5}$	69.94°	107.35
√3	**150°**	$\mathbf{+2.9968\times10^{-5}}$	**73.02°**	**11.33**
2.0	153.43°	$+2.8750\times10^{-5}$	73.49°	24.44
3.0	161.57°	$+2.5729\times10^{-5}$	74.10°	109.96

Quarantine note (companion §G.6.10). The 150^∘ row of this table is the branch the derivation does *not* return; the provisional P_{23} reading gives 30^∘ or 270^∘. The scan is retained as a sensitivity diagnostic only. One-parameter diagnostic optimum 148.362°, separation 1.638° in a 360° range.

Appendix E — Verification checklist

W_6: |V|=7, |E|=12, β_1=6 · spec(L)={0,2,2,4,4,5,7} · **the quoted transition spectrum is NOT part of this checklist — it is unresolved (§14.1.1) and must not be verified against** · k=±1 harmonics have L=2 and r^3=-I; k=±2 have L=4 and r^3=+I · **VX-W6 identification with EX2F certified by Stage I: overlap 1.0000000000000000, ‖Lv−2v‖ = $\mathbf{1.75\times10^{-15}}$** · χ_E_1 = (2,1,−1,−2,0,0) and χ_E_2 = (2,−1,−1,2,0,0) with $\langle E_1,E_2\rangle = 0$ · the $\mathbb{Z}_2$ form of the no-go: T(v) = −T(v) from the involution alone · the admissible-power tower dim Hom(Sym^k E_1,E_2) = 0, 1, 0, 2, 0, 2, 0, 3 for k = 1…8, and the weight argument that only even k reaches the target · half-turn parity of the mode, u_n+3+u_n = 3.4×10^{-16} and u²_n+3−u²_n = 2.8×10^{-16} · $\chi\_\mathrm{Sym}^2E_1$ = (3,0,0,3,1,1) with $\langle \mathrm{Sym}^2E_1, E_2\rangle = 1$ · Q(r^k z) = e^(2πik/3)Q(z) for k = 0..5 and Q(z̄) = conj(Q(z)) · u_n+3 = −u_n at 2.24×10^{-16} · ‖L(u)‖ = 2.05×10^{-16} · Q(u) component moduli ⅓, norm 1/√3, sum 0, spread 1.11×10^{-16} · R_3Q(u) = e^(∓2πi/3)Q(u) · args of Q(u) = (0°, ±120°) · V unitary, det V = 1, ζ = λz bit-identical across conventions · $\Pi_{23}Z_0 = Z_{23}/\sqrt{3}$ · **AX1 diagnostic:** the anchoring transfer gives |ζ| = b, and under the declared AX1 normalisation the six phase-census branches share |ζ| = b/√3 — **neither is a primitive physical return** · IORC-1: π_hub = 7/D_α and π_rim = 2(1+cosh α)/D_α at ≤5.6×10^{-17}, rim current uniform at 6.9×10^{-17}, **spoke currents zero at** $\mathbf{2.1\times10^{-17}}$, entropy per event = ln 2 at 3.3×10^{-16}, α⋆ unique by monotonicity, C_3 characters exact eigenvectors at 7.2×10^{-17} **when paired with the correct conjugate eigenvalue**, and $K_- = \Pi^{-1}K_+^{T}\Pi$ at 1×10^{-15} **against the π-reversal, not the plain transpose**.

Appendix F — What would move the boundary quickest

VX-W6 no longer heads this list: Stage I closed it (§7.0). The source, the map Q, the target E_2 and the democracy are settled, and none of the remaining work is wheel or flavour algebra. The blockage is the master action's inability to select a faithful internal lift and a noncentral physical readout from a large family of equivalent ones.

1. **Return a nonzero same-action c_{ax} for the identified mixed-jet form**, including μ_F and H_F in physical units. The local form is fixed up to one real scalar; the missing result is its nonzero physical installation.
2. **Return the complete physical lift, P_{vis}^{curv} and $\rho_{mult}^{physical}$.**
3. **Realise the co-terminal comparison loop** L_{ax} as a pair of admissible histories.
4. **Evaluate its holonomy and return θ_{phys}** — the only candidate source of a branch-independent phase now that the local metric is known to be exactly phase-isotropic (companion §G.6.11).
5. **Perform the post-seal Stage I-B comparisons** with b, unit retention and $5\pi/6$.
6. **Perform sealed RRHF Stage II reintegration**, with the branch index and phase returned rather than fixed in advance.

Two further items are real but of a different kind and should not displace the six above: defining $\hat{T}$ (A0b, a graph-operator debt the axle does not depend on) and declaring the comparison scale (I3, which affects empirical confrontation rather than installation).

Provenance, Audits and Retractions for CKP-2.4

Companion to *The Conditional Derivation of Charged Flavour and the Unique Lowest-Order Quadratic CP Axle in VERSF*

Keith Taylor VERSF Theoretical Physics Programme — Standard Model Flavour Series Designation: CKP-2.4-C

This companion carries the material the main paper depends on but does not need to display: the inherited-constraint ledger and grading rules, the transport algebra of the curvature extension, the boundary and installation contract, the full premise/falsifier/debt ledgers, and the complete chronology of executions, audits and retractions. **The main paper's claims are stated there; the evidence and the record of how the gradings were reached are here.**

Contents

Part C1 — Inherited corpus constraints (main paper §A–§D depend on these)

Position under the spine. These constraints bind the extension before it begins. They are not the programme's derivation, which is RRHF's; they are the rules under which a proposed correction to that derivation may be graded. §0 predates the §A–§D front matter and is retained as the record of how that position was reached.

§C0. Level statement — how this position was reached

Added on reconciliation against the July 2026 Master Ledger and SMC-1U.

This paper audits **one proposed role-even correction** to the programme's controlling CKM return: the C_3 curvature construction, in which a role-even branch source is descended, read out, and exponentiated against the inherited role-odd generator. **The cumulative framework does not contain two independent routes to CKM mixing.** It contains one shared flavour construction, returned most completely by RRHF-1: the representation-resolved history functional RRHF-1 returns the charged Yukawa tensor at conditional action level, with singular values

$\sigma(Y\_u) = (7.22\times10^{-6}, 3.71\times10^{-3}, 0.971983)$, $\sigma(Y\_d) = (1.56\times10^{-5}, 3.13\times10^{-4}, 0.016196)$,

from which $V\_CKM = U\_u^\dagger U\_d$ follows directly without passing through a curvature-plus-readout construction at all.

Consequences for how this document must be read.

1. Every "not identified" and "not returned" grade below is a statement about **the curvature route**, not about the programme. The cumulative stack may already return the corresponding object by the RRHF route. Debts A0–A5 and R1–R3 are debts of this construction.
2. By the cumulative reading rule, a local calculation on a reduced carrier may test consistency but **may not re-open a quantity returned by the cumulative history-action stack unless it directly evaluates and contradicts that stack.** Part II does not independently re-execute RRHF-1 — Part I does. Its negative results concern the reduced curvature package and cannot downgrade the controlling RRHF return unless they evaluate the same operator on the same carrier and readout.
3. The scalar-readout no-go (§21) and the anti-linear typing result (Prop 9) are therefore statements about **the reduced scalar package**, not about the cumulative action. They are retained at that scope and must not be quoted at programme level.
4. What survives at programme level from this document is: the exact parity and D_6 representation theory (with the linear no-go holding under $\mathbb{Z}_2$ alone), the W_6 graph certificate, the intra-wheel character of the descent, the source–readout bijection, the freedom-budget bound, and the **phase-carrier typing question of §18.2** — which the Master Ledger independently records as open ("primitive orientation and admitted branch remain open"), so the withdrawal of this paper's programme-level conclusions does **not** extend to the orientation finding.

§0.1 One construction, not two routes

The resolution the paper states. RRHF-1 is the controlling programme-level return: it supplies the complete charged Yukawa tensor at conditional action level, from which V_CKM = U_u†U_d follows directly. The construction studied here is **not a second derivation of CKM**. It shares (a, b, c, φ), the symmetric weak-doublet split, the relative-frame architecture and the baseline generator, and then proposes an additional role-even correction $\Omega_0(z)$. The structure is

shared generator and relative-frame architecture ⟶ RRHF/RRMK charged-tensor return ⟶ { baseline CKM frame, **possible role-even curvature correction Ω_0** }.

Construction	Correct status
RRHF/RRMK	**Controlling conditional action-level CKM return**
Bare e^(Ω_q) frame	Shared baseline of the construction
The curvature calculation of this paper	**Candidate role-even correction to that baseline**
The D_6 quadratic axle	Independent mathematical mechanism that may generate the correction
Completed-H_CMR Stage I projection, then sealed RRHF Stage II reintegration	Test of whether the action installs the correction, then whether it survives integration

The unresolved question is therefore not which of two CKM paths is correct. It is: *does RRHF already generate Ω_0, or is Ω_0 additional constitutive information?* That is a correction question inside one construction, not a contest between routes. The status of the curvature block affects the refinement of the programme's return, not the existence of a competing second route.

§0.1.0 Why the two constructions are not independent

The previous version stated that the Master Ledger §6 modulus matrix and this paper's §11–§12 return agree to 5.6 significant figures, and treated that as a possible cross-route corroboration. **That comparison was invalid.** The Master Ledger's §6 explicitly labels its matrix "the full-curvature conditional reconstruction", and SMC-1U identifies the matrix it prints as the full-curvature result with the source-locked confrontation living in CKP-2 itself. The comparison was therefore CKP against a reproduction of CKP within one lineage, and the residual 2.81×10^{-6} is a version, transport or rounding difference, not evidence of anything.

The dedicated provenance audit returns: the routes are NOT disjoint. They share the relative-frame architecture $V_{CKM} = U_u^\dagger U_d$, the symmetric weak-doublet split $U_u = U_0 e^{(-\Omega_q/2)}$, $U_d = U_0 e^{(+\Omega_q/2)}$, and — decisively — the same quark orientation generator with the same coefficients. RRHF-1 recalculates $a = 9/40$, $b = 81/2000$, $c = 243/10^5$ in its discrete $K = 7$ ledger; the curvature construction inherits them from the Yukawa and electroweak-frame papers. **The agreement of the leading CKM hierarchy is therefore inherited, not corroborative.**

§0.1.1 Independent confirmation of the audit's structural picture

The audit's account is checkable from the baseline frame alone, and it holds. With the symmetric split, the shared trunk is exactly $V_{base} = e^{(\Omega_q)}$, which returns

$|V_{base}|$ = [[0.974795, 0.223069, 0.003931], [0.223020, 0.973981, 0.040285], [0.006108, 0.040013, 0.999181]], **$J_{base} = 1.8435063\times10^{-5}$.**

Against that baseline:

Frame	J	departure from the shared baseline
shared trunk $e^{(\Omega_q)}$	1.8435×10^{-5}	
RRHF level-3	1.9108×10^{-5}	**+3.65%**
curvature-corrected extraction (this paper)	2.9968×10^{-5}	**+62.56%**

RRHF sits within a few percent of the bare shared generator; the curvature correction moves the same baseline by nearly two-thirds. That is precisely the audit's picture — one trunk, one branch that adds Ω_0 and one that does not — and it is confirmed here without using any RRHF internal.

The modulus comparison agrees too: the maximum difference between the bare baseline and the curvature frame is **0.002582**, occurring at the (3,1) entry, against RRMK-1's reported maximum modulus difference of 0.00279523 from the frozen curvature boundary — the same quantity to within 8%, which is what one expects if RRMK is the baseline carrying its own small corrections.

§0.1.2 The 37.92% figure, resolved

The quoted 37.92% departure is the level-4 RRHF/RRMK Jarlskog magnitude relative to the frozen curvature value:

$(2.996771 - 1.860329)/2.996771 = 0.3792.$

The earlier attempt in this series to associate it with a separate RRMK frame was caused by comparing the **level-3** RRHF value 1.910812×10^{-5} with the **level-4** RRMK value 1.860329×10^{-5}; those differ by 2.71%, which is the regulator step, not a frame discrepancy. RRHF reproduces the corresponding RRMK frame at each regulator. **Debt I6b is closed.**

§0.1.3 What is and is not left independent

The curvature calculation is a **candidate correction to a shared baseline**, not a second derivation. What remains potentially independent is narrower than this paper previously implied:

- **Independent as mathematics:** $\mathrm{Hom}(E_1,E_2) = 0$ — a $\mathbb{Z}_2$ parity result (Cor AX-1.1) — and $\dim \mathrm{Hom}\_D_6(\mathrm{Sym}^2E_1,E_2) = 1$, the W_6 graph certificate, and the intra-wheel character of the descent. RRHF does not use these.
- **Potentially independent as physics:** the $\Omega_0(z)$ block itself, its coefficient, embedding, readout and phase carrier.
- **Not independent, as a diagnostic:** the **AX1 execution** sets the amplitude to $b/\sqrt{3}$, so that diagnostic value shares ancestry with the baseline generator. **The physical amplitude has not been returned; it is ρ_mult^physical·c_ax/√3, with c_ax unreturned** (FD-3). **The correct remaining test is not a matrix comparison. It is two-stage.** Stage I freezes the completed H_CMR and asks it to return the full Stage I-A tuple — c_ax, the complete full-carrier lift $\mathbb{I}$_int^full, P_vis^curv, r, ρ_mult^physical and θ_phys — without CKM input. Stage II inserts whatever Stage I returns into RRHF and re-differentiates the complete charged tensor. **RRHF tests survival and integration; it does not force the microscopic axle.**

§0.2 Provenance of the inherited generator

$a = 9/40,\ b = a\cdot\eta^2/(D_3 - D_2) = 81/2000,\ c = \chi\cdot a\cdot[a\cdot\eta^2/(D_3 - D_1)] = 243/100000,$

with $D = \mathrm{diag}(1, 2, 4)$, $\eta = 3/5$, $\chi = 2/5$. Verified in exact rational arithmetic: both b and c reproduce their target values identically, with no further number inserted.

Consequently A = 0.800 is not an ansatz. At leading Wolfenstein order,

$A\_gen = b/a^2 = \eta^2/(2a) = (9/25)/(2\cdot9/40) = \mathbf{4/5},$

an exact algebraic consequence once a, η and the stiffness gap $D_3 - D_2 = 2$ are granted. The question "where does the clean 0.800 come from?" is therefore closed. The generator has an exact **leading-order Wolfenstein-shaped form** with λ_gen = 9/40 and A_gen = 4/5; calling it "exactly the Wolfenstein parameterisation" is too strong, since the exact CKM entries receive exponential and curvature corrections. Against measurement, λ_gen is +0.31% and A_gen −3.15%.

The identity $A = 2\chi$ also holds exactly, but no source theorem identifies the adjacent attenuation coefficient with twice the non-adjacent continuation fraction. It is an arithmetic coincidence inside the frozen package and may not be cited as a derivation.

§0.2.1 Where the live debt actually sits

The audit's grading of the upstream operands is adopted:

Object	Current cumulative grade	Genuine residual bridge
a = 9/40	conditional geometric derivation **plus conditional action ancestry** (RRHF blind embedding)	typed normalised overlap operator, and equivalence of the original and RRHF count operands
D = diag(1,2,4)	conditional structural derivation, strengthened by an independently ordered three-level census	derive capacity balance and the stiffness map from completed H_CMR
$\eta = 3/5$	**exact conditional transport theorem** under the democratic residual measure	select the uniform residual measure; derive the η^2 serial composition
$\eta^2 = 9/25$	conditional two-filtration attenuation law	same-action composition theorem
$\chi = 2/5$	structural rank theorem; physical weight conditional on Green isotropy	derive the isotropic intermediate response, or compute the weighted trace directly
$\varphi = 2\pi/3$	strong structural C_3 holonomy with conditional action ancestry	primitive role-odd orientation carrier, branch sign and placement
b, c, A_gen	**exact consequences, not inputs**	

The upgrade over the first provenance pass is that a and φ now carry *conditional action ancestry* as well as geometric ancestry: RRHF-1 embeds the ledger in a frozen augmented parent and returns the charged tensor blindly. That is conditional action sufficiency, not primitive uniqueness.

RRHF-1's rewriting $a = 9/[2\cdot C(6,3)]$ does not upgrade this. Its numerator is a colour–triality readout where the original was a product of generation modes, and its denominator is twice a balanced-support shell where the original was available mixing channels. No typed theorem bridges those operands.

§0.2.2 The denominator agreement is a single-point coincidence

The audit observes that $6K - 2 = 40 = 2 \cdot C(K-1, (K-1)/2)$ at $K = 7$. **That is not an identity.** Checked across odd K:

K	6K − 2	2·C(K−1,(K−1)/2)	equal
3	16	4	no
5	28	12	no
7	**40**	**40**	**yes**
9	52	140	no
11	64	504	no
13	76	1848	no

The two formulas agree at exactly one value of K and diverge rapidly on either side. This **sharpens** rather than softens the audit's conclusion: the numerical coincidence of the two denominators carries no structural content, and the demand for a typed bridge between "available mixing channels" and "twice the balanced-support shell" cannot be met by pointing at the shared value 40. Either operand set must be derived, or the agreement is accidental.

§0.2.3 A typed-comparison error, corrected

That comparison was not licensed. The controlling rule — **a later calculation may downgrade an earlier result only if it evaluates the same typed object, on the same carrier, with the same physical readout** — is exactly the rule this paper failed to apply to itself. The source construction claims **C_3 covariance for the residual carrier, not D_6**, and under C_3 the objection dissolves:

rim harmonic	C_3 character (generated by r^2)
k = 1, 4	ω
k = 2, 5	$\bar{\omega}$
k = 3	trivial

so the C_3 multiplicities are (1, 2, 2), and a **C_3-invariant three-dimensional subspace does exist** — trivial $\oplus$ one ω $\oplus$ one $\bar{\omega}$. **The (3 + 2) split is C_3-covariant.** The objection is withdrawn.

§0.2.3.1 What survives as a well-posed question

One thing does survive, and it is a question rather than an objection. The K = 7 wheel that supplies the axle carries the **full D_6**, under which the same five-dimensional space decomposes as (2 + 2 + 1) and the only invariant three-dimensional subspaces are $E_1 \oplus (k = 3)$ and $E_2 \oplus (k = 3)$. So:

Is the closure-transport residual carrier D_6-covariant, or only C_3-covariant?

If it is only C_3-covariant, the audit's grading stands untouched and this paper has no business comparing the two. If it is the wheel's rim and therefore inherits D_6, the (3 + 2) split needs justification, and the choice between $E_1 \oplus (k = 3)$ and $E_2 \oplus (k = 3)$ matters, since E_1 is the axle's source and E_2 its target. That is a typing question about the source construction, not a result of this paper, and it is recorded as debt I1b in that form.

A second point is retained unchanged, because it is the audit's own open bridge rather than an objection to it: on a reducible carrier **"democratic" requires its measure named.** Equal weight over the five directions gives 3/5; the C_3 isotype multiplicities are unequal (1, 2, 2), so equal weight over isotypes gives something else. The audit already lists deriving the physical invariant measure as owed; this simply notes that the reducibility makes the specification unavoidable.

§0.2.4 Three phases that must not be conflated

The cumulative audit's fourth withdrawn comparison — that the void C_3 character and the common-mode curvature phase are *related to but not automatically identical with* the role-odd generator phase — is adopted here as a standing typing firewall. Three distinct objects carry 120°- or 150°-shaped values and are routinely confused:

Object	Value	Type	Status
O_q	e^(2πi/3)	C_3 **generation character**	recovered by the axle (§17)
arg ζ_C_3	5π/6 = 150°	**role-even common-mode curvature** phase	**source-phase candidate derived; inheritance by the physical branch coefficient not established**
φ	2π/3 = 120°	**role-odd holonomy** in the direct 1↔3 seed entry c·e^(iφ)	**inherited; typed primitive bridge open**

O_q and φ share the value 120° and are *not* the same object. **The axle's recovery of O_q does not derive φ**, and no result in this paper should be read as closing the role-odd holonomy debt (falsifier F31).

IORC-1 sharpens this and must not be read as dissolving it. That paper writes O_q(σ) = e^(iσ2π/3) *and* φ_σ = σ·2π/3 in a single display. The first is a statement about the **generation character**; the second asserts the **role-odd holonomy**. Its own closing sentence says the same-action carrier map into the role-odd operator remains open, and its ledger records the physical embedding as false. **The display runs ahead of the text**, and F31 continues to apply. The orientation chain is therefore

j_prep^σ ⟶ O_q(σ) = e^(iσ2π/3) ⟶ **[I1d open]** ⟶ φ_σ ⟶ Ω_q ⟶ sgn J,

with I1d the single typed blocker on it. **The bridge is now writable as an explicit operator statement**, which is what I1d must return from the same action:

𝕀_role-odd[O_q(σ)] = e^(iσ2π/3)E_{13} − e^(−iσ2π/3)E_{31},

anti-Hermitian as required. Note the typing: this is the **orientation** map only. The amplitude of the seed entry is $c = \chi a[a\eta^2/(D_3 - D_1)]$ and comes from the count package, so the derived object is $c \cdot \mathbb{I}$_role-odd[O_q(σ)]. Once this map is returned by the same action, IORC supplies the orientation sign and the role-odd generator receives it **without inserting φ separately**.

§0.2.7 IORC-1 numerical-certificate corrections — no change to the structural return

Re-executed from the wheel and the stated action, with no input from that paper's arrays. **The substantive return stands**; two of its printed certificates describe the wrong comparison object and understate what was achieved.

Claim	Independent value	Status
π_hub = 7/D_α, π_rim = 2(1 + cosh α)/D_α	0.189423730905787, 0.135096044849036	confirmed, residual $\leq 5.6\times10^{-17}$
rim current uniform on all six edges	spread 6.9×10^{-17}	confirmed
J_rim = 2 sinh α/(19 + 12 cosh α)	0.060231697816834	confirmed, residual 4.9×10^{-17}
spoke currents vanish	$\leq 2.1\times10^{-17}$	confirmed — detailed balance holds on the spokes, so **all** entropy production is on the rim
entropy production per event	0.693147180559945	equals ln 2 to 3.3×10^{-16}
α⋆ from the one-Bit condition	0.959001109719975	confirmed; entropy monotone on (0, 5], so the root is unique
C_3 amplitude 2J_rim, eigenvalues ±i√3 J	0.120463395633668, 0.208648721689784	confirmed
odd edge-score space one-dimensional		confirmed analytically: rotation invariance leaves a rim and a spoke parameter; reflection is odd on the rim and **even** on the spokes, forcing the spoke parameter to zero

Correction 1 — the C_3 eigencharacter check. For $J\_C_3 = a\cdot\mathrm{circ}(0, 1, -1)$ the vectors $v\_\pm = (1, \omega, \omega^2)/\sqrt{3}$ and its conjugate carry **conjugate** eigenvalues $\pm i\sqrt{3}a$. The printed residual 4.173×10^{-1} is exactly $2\sqrt{3}\cdot J\_C_3$, i.e. $\|Mv - (-\lambda)v\|$ — each labelled vector was paired with the eigenvalue of the *other* branch. Paired correctly, $\|Mv - \lambda v\| = 7.2\times10^{-17}$. **A sign or label exchange merely swaps + ↔ −; it does not invalidate the returned conjugate character pair**, and the check must pair each labelled vector with the eigenvalue obtained under the adopted rotation convention.

Correction 2 — the kernel reversal check. The physically correct comparison is not $K_- = K_+^T$ but the stationary-measure reversal

$K_- = \Pi^{-1}K_+^T\Pi$, $\Pi = \mathrm{diag}\,\pi$.

Because the hub and rim stationary weights differ, plain transposition is the wrong diagnostic; the printed 5.745×10^{-2} is $\|K_- - K_+^T\|$_max. Against the correct object the residual is 1×10^{-15}. The clean invariant check remains the current reversal $J_- = -J_+$, which the source reports at 9.0×10^{-17}.

Provenance note. α⋆, π_hub and π_rim are **identical to the archived D36/ISG-1 values**. IORC-1's contribution is the *action* whose unique minimiser is that kernel, not the numbers — which the source states plainly. Agreement with the archived branch is reproduction, not independent confirmation.

§0.2.5 How much the role-odd holonomy actually does

Diagnostic, quarantined, non-insertable. Holding everything else at its adopted value and varying φ alone:

φ	J	‖V_ub‖	γ	χ²_all
0	$+1.1545\times10^{-5}$	0.004863	15.79°	1394
π/3 = 60°	$+2.9974\times10^{-5}$	0.004859	44.80°	550
π/2 = 90° *(early candidate)*	$+3.2822\times10^{-5}$	0.004356	59.10°	194
2π/3 = 120° *(adopted)*	**$+2.9968\times10^{-5}$**	**0.003557**	**73.02°**	**11.3**
5π/6 = 150°	$+2.2176\times10^{-5}$	0.002517	86.10°	423
π = 180°	$+1.1533\times10^{-5}$	0.001312	96.08°	1646

The one-parameter diagnostic optimum is **117.169°**, so the adopted 120° sits **2.83° away in a 360° range**, with steep walls either side.

This is the same pattern already recorded for the curvature phase, which sits 1.64° from its own optimum — but the two are *different phases on different carriers*, and neither proximity is evidence for the other. Both are adopted rather than derived, so what the scan establishes is narrow and worth stating exactly: **whatever primitive construction eventually derives φ must land within a few degrees of 120°, and the observables are sensitive enough to test it at that resolution.** The scan may not be used to select φ (F4).

§0.2.6 Relation to the master reduction

The audit's closing point is adopted: the five open bridges are **not five independent programme-level debts**. Their unresolved portions are specific manifestations of the incomplete commitment-maintenance action — the event-level directing measure controls the normalised overlap and residual weighting; the post-commitment maintenance law controls stiffness and serial attenuation; the orientation-sensitive maintenance sector controls the holonomy; the fermion-dressed readout controls the typed placement of the generator. Completing H_CMR should promote all five or falsify one on its proper carrier. The debt ledger below is organised accordingly.

§C1. Why this paper begins here

VERSF is not a collection of standalone proposals, and a flavour paper written as though it were will keep rediscovering locally what the corpus proved globally. The governing question is:

Given the prior non-identifiability, tautology and singularity theorems, and given the later constitutive extensions, what flavour information is now genuinely fixed, what remains conditionally transportable, and what still requires installation?

§C2. The binding inheritance ledger

Prior result	Updated consequence for CKM
Theorem 72.1 (pre-extension non-closure) — six witness families prove a continuum of inequivalent positive boundaries preserving every quantity certified by that corpus	Binding for any object built only from the audited scalar corpus. PFDMI is a later constitutive extension and may break part of the witness freedom, but **must demonstrate that break explicitly**. P_vis^curv, the faithful internal embedding and the total norm are not yet shown to lie outside the witness family.
Occupancy witness — the history functional is minimised at zero occupancy; the nonzero branch requires an underived occupancy premise	**No channel may be removed through an unreturned occupancy premise**, in any sector, in any form.
η-continuum and whitening theorem	Absolute response normalisation cannot be read from a scalar Schur form; an identity readout is not physical content.
ASAC-1 — action rescaling leaves phases, Born weights and dimensionless ratios invariant; Λ⋆ non-identifiable	Overall action scale cannot select the phase or any dimensionless mixing ratio. Any new term must supply **new dimensionless oriented structure**, not a coefficient. The void axle is admissible under this test precisely because it is a unique representation-theoretic intertwiner.
TPB pre-temporal certificate — reversible frozen scalar kernels have $\Sigma = 0$, $\lambda\_TPB = 0$	Retained as a no-go for those kernels. Superseded at extended-action level by the PFDMI jump/no-jump compensator, finite mean waiting and isolated oriented rim mode. SI calibration and the internal endpoint frame remain open.
PFDMI-EX2F / EX2G / EX2H — the K = 7 history law selects an oriented sixfold mode;	The programme now possesses the source geometry required to test a void-to-generation

Prior result	Updated consequence for CKM
the attaching map is forced within the declared class; the committed hub-plus-rim geometry has exact D_6 stabiliser and three unique antipodal axes	descent. It supplies orientation and the sixfold/three-axis carrier, **not** the faithful internal flavour embedding or physical readout.
Explicit canonical K=7 closure-transition graph — $W_6=C_6 \vee K_1$,	V
RRHF-1 (representation-resolved history functional) — returns the charged Yukawa tensor at conditional action level, $\sigma(Y_u) = (7.22\times10^{-6}, 3.71\times10^{-3}, 0.971983)$, $\sigma(Y_d) = (1.56\times10^{-5}, 3.13\times10^{-4}, 0.016196)$, $\sigma(Y_e) = (2.99\times10^{-6}, 6.20\times10^{-4}, 0.010696)$, with the full charged tensor matching its target operator to $\approx7\times10^{-14}$ and two CY3′ routes agreeing to 1.9×10^{-14}	**The dominant inherited constraint on this paper.** $V_{CKM} = U_u^\dagger U_d$ follows from RRHF-1 without any curvature-plus-readout construction, so this paper's construction is a proposed correction to that return, not a parallel path, and its open grades bind only the correction. The Master Ledger nonetheless grades the resulting CKM as "conditional magnitude and CP reconstruction; **primitive orientation and admitted branch remain open**" — which is the same conclusion this paper reaches at §18.2 by an independent argument. Disjointness is **closed negatively**: the two share the frame architecture and the generator coefficients, and RRHF's frame sits within 3.65% of the bare shared baseline $e^{(\Omega_q)}$ while the curvature correction departs from it by 62.56% (§0.1.1). **RRHF is the controlling programme-level return; the curvature block is a candidate correction to it.**
IORC-1 (irreversible oriented record current) — a strictly convex preparation action on the same wheel whose unique minimiser is the exponential tilt $K \propto T_0 e^{(\sigma\alpha\omega_C)}$; the one-Bit condition $24\alpha \sinh \alpha/(19 + 12 \cosh \alpha) = \ln 2$ fixes $\alpha\star = 0.959001109719975$ uniquely, giving a nonzero stationary rim current $J = \sigma\cdot0.060231697816834$ whose antipodal quotient has the conjugate C_3 characters as exact eigenvectors	**Independently re-executed and confirmed** (§0.2.7). Supplies the **realised sign** of the generation character as a spontaneous first-event mark — the law is reversal-symmetric and prefers neither branch. Does **not** select the sign at law level, prove persistence, or derive the role-odd seed holonomy φ. Advances the orientation debt O1; leaves the maintenance lock and the $O_q \to \varphi$ carrier map open.
PSU-1 (primitive-source uniqueness audit) — Fold/Hilbert/Born architecture passes and all 76 XI-4 nulls are exactly classified as 72 gauge + 1 record-radial + 3 ledger, but **primitive-action uniqueness FAILS** on three explicit counterfamilies	**The most consequential inheritance for Appendix G.** All three counterfamilies re-executed and confirmed: a third symmetry-allowed Fold-edge coupling (rank-4 invariant basis, TV = 9.887×10^{-3}), an infinite Wilson-harmonic family with identical $S''(0) = Z$ (TV =

Prior result	Updated consequence for CKM
	1.212×10^{-1}), and an infinite completion-amplitude family agreeing through fourth order. **The admissible family is infinite-dimensional, so the ≥22 floor is superseded in kind** (§G.6.2). The Wilson family is structurally identical to the flavour tower of AX-1.5b, making debt A6 an instance of a programme-wide pattern (§G.6.3).
HCMR-XI-4 (neutral selection, amplitude alphabet, finite mixed-jet instrument) — Dirac branch selected by left polar modulus; charged alphabet from the ordered amplitude; four-channel finite completion instrument; absolute action unit **not** derived	Three gates close and one no-go is retained. The charged alphabet is now unique **up to block ordering** rather than up to O(36) — a correctly sized gain. The neutral selection carries the A5e-type provenance question (A5f), one internal consistency relation is off by 0.026% (A5g), and the scale arithmetic shows the added relation **raised** nullity from 2 to 3. New structural check: $\tau\_F/\Delta t_0=2$ and $T_0/\tau\_F=7=K$ exactly.
HCMR-XI-3 (source-native marked residual instrument) — a minimum-relative-entropy oriented-mark law whose Fisher is exactly $J\_D^TJ\_D$ (verified), one global normaliser over 402 marks, and a reported exact HM4 gauge-shape closure	Repairs XI-2's failure by attacking **structure**, confirming the §7.0.1.2.3 diagnosis. The gauge agreement at 4.485×10^{-16} is either an exact identity or a reproduction, depending on whether Z descends from the HM4 form — **unsettled**, debt A5e. Unambiguously new: the Fisher identity, the single-instrument typing, C_compq_Y=C_compq_B-L=0, and the branch-resolved neutral sector.
HCMR-XI-2 (microscopic matrix execution) — RSU instrument, CPR copy isometry and the record-completed 102D parent all pass at machine level, with the record block converging second-order to $\lambda\_min\to1.81881\times10^{-3}$; **same-action HM4 gauge reproduction FAILS**	Confirms the flavour position rather than changing it. The V-metric refinement is closed negatively (no positive metric exists on an oriented branch). **The reported failure diagnosis is corrected here**: a stiffness anisotropy would leave the matrices ~90% aligned, and they are 11% aligned, so the mismatch is of eigenstructure not eigenvalues (§7.0.1.2.3).
HCMR-XI-1 (record-sufficient completion and scale selection) — conditional law-level completion of H_CMR under RSU and CPR, with ξ_physical=84.137117µm selected on the post-copy branch	**Closes the record-conditioning and copy ambiguities, not the coupling ambiguity.** A later revision upgrades RSU and CPR from adopted principles to consequences of the single-source, non-contextual-weight, operational-distinguishability and record-definition theorems — accepted here, subject to §7.0.1.2's two qualifications (the likelihood

Prior result	Updated consequence for CKM
	exponent is fixed only by the single-record Born normalisation, which the derivation does not cite; and the proposed V-metric refinement introduces a non-unique object). Its completed cost carries no free coefficients, but is defined relative to a still-chosen T_0 and contains no wheel Hessian term, so the 6 quadratic and 13 cubic invariant couplings of Appendix G — including c_ax — are untouched (§7.0.1.1). The ξ selection resolves the former 0.078% discrepancy as a **UV-ruler artefact**, verified: the coherence value and the old Route-M value differ by exactly the CODATA/old Planck ratio to 1.5×10^{-8}. That is a rounding check, **not** independent route agreement, and may not be cited as the latter. S=ħM imports the action unit, consistent with ASAC-1's rescaling no-go and correctly flagged at source — so the scale closes by measurement at one point, not by derivation.
Scalar-readout no-go (§21)	The reduced scalar mechanism returns g_A23 = 0 and cannot produce the phase. **This is not a no-go for the cumulative action, and by R6 may not be quoted as one.**
Void-canvas axle theorem (§15–§17)	No odd-power equivariant descent exists — a $\mathbb{Z}_2$ parity result; the quadratic descent is unique at that order and returns exact C_3 democracy and the character chain, with the even tower above it unbounded. Dynamical insertion, internal lift, readout and the phase-carrier typing remain owed.

§C3. Grading rules

Exact · **Conditional transport** (for any admitted source, the following follows) · **Numerical certificate** (stated grid, precision, margins) · **Not identified** (proved or inherited-by-default undetermined; distinct from *not yet computed*) · **Closed negative for a named mechanism** (requires a different mechanism, not effort) · **Structurally returned, installation owed** · **Diagnostic** (computed with benchmarks in hand; never citable as derivation).

R1. No object inside the Theorem 72.1 witness freedom may be described as returnable by evaluation of the pre-PFDMI action.

R2. No normalisation may be quoted without stating which of the four of §20 it is.

R3. A no-go proved for the pre-PFDMI scalar mechanism may not be promoted to a no-go for the extended marked-history action. Conversely, **the existence of a symmetry-allowed void axle may not be promoted to a physical CKM return** until its nonzero coefficient, faithful internal lift, readout and phase-carrier typing are returned by the same action.

R4. Where a result depends on which object carries a phase, the object must be named. Two prior corrections in this series — the complement typing and the geodesic base point — were failures of exactly this discipline.

R6. Every grade in this paper is extension-scoped. "Not identified", "not returned" and "closed negative" mean *within the proposed curvature correction*. None may be quoted against the cumulative stack, whose controlling RRHF return reaches the same observables without this correction. An extension-level negative becomes a programme-level negative only by directly evaluating and contradicting the cumulative operator, which this paper does not do.

R8 — the reproduction rule. A machine-zero agreement between two quantities is a reproduction, not a corroboration, unless the two are shown to have disjoint derivation chains. The burden lies on the party citing the agreement, and it is discharged by exhibiting the two chains — not by certifying that one array was unopened while a quantity derived from it was used.

This has arisen four times in this series: the RRHF/curvature CKM agreement, the six-figure J_q coincidence, Z against the HM4 gauge form, and S_D against the polar modulus $\sqrt{(M_\nu M_\nu^\dagger)}$. The shape is identical each time — a residual at 10^{-16} between an object and a reconstruction built from its own derived quantities. Two resolved as reproductions; two remain open (A5e, A5f). Stating the rule in advance is cheaper than re-litigating it per execution.

R7. The paper discusses two calculational realisations of one shared construction — the RRHF baseline return and a proposed curvature-corrected extraction — and must not describe them as two physical routes. They are not independent paths, and RRHF controls the programme-level grade.

R5. Graph counts are typed. The twelve wheel edges, six independent cycles, one stationary Perron mode, two-dimensional $k=\pm1$ slow eigenspace, three generation axes and $1/\sqrt{3}$ branch democracy are not interchangeable numerical resources. In particular, no charged-lepton 1/12 or 7/6 factor may be transferred into the CKM amplitude or phase without an explicit operator map into the physical weak-doublet Hessian.

Part C2 — The extension's transport algebra

Position under the spine. The programme's charged-sector derivation is §C: Y_f from the action derivative, then $V_{CKM} = U_u^\dagger U_d$. This Part is not that derivation. It is the exact algebra of the proposed correction — statements about what follows *given* an admitted branch source, none of which asserts that the action returns that source. Its results are exact and construction- independent where marked **[X]**, and extension-scoped otherwise (R6).

§C4. The projected object

$$P\_W = P\_vis^{curv} P\_R P\_C,\quad H\_W = P\_W^\dagger H\_cl P\_W,\quad \Omega\_W = i\varepsilon\_W H\_W = \Omega\_even \otimes I + \Omega\_odd \otimes \tau_3 + \Omega_1 \otimes \tau_1 + \Omega_2 \otimes \tau_2,$$

with the clean flavour-frame regime in which Ω_mix is gapped. **Amplitude democracy is a pre-readout statement; branch, leakage, retention and phase are post-readout.**

§C5. The branch carrier and inherited generator

Hermitian quadratures S_ij, A_ij, planes $\mathfrak{B}\_ij$ mutually Killing-orthogonal, coordinates $x = (s_{12}, a_{12}, s_{23}, a_{23}, s_{31}, a_{31})^T$, cycle R_3 permuting planes.

$$\Omega\_q = \begin{bmatrix} 0 & a & c\cdot e^{(i\varphi)} \\ -a & 0 & b \\ -c\cdot e^{(-i\varphi)} & -b & 0 \end{bmatrix},\quad a = 9/40,\ b = 81/2000,\ c = 243/100000,\ \varphi = 2\pi/3,$$

with $\|\Omega\_q + \Omega\_q^\dagger\| = 0$. Its (2,3) entry is $b \neq 0$, so no plane may be selected as unoccupied.

§C6. The six-coordinate Hessian return

Closure, transport and record second variations each return a scalar on every branch plane, so

$$\mathbf{G\_C_3^{\mathfrak{B}} = \kappa\_C_3\, I_6,\quad [K\_even,pre^q, R_3] = 0.}$$

Coordinate note. This is the *descended*-coordinate description of the same coupling that Appendix G counts in the cubic sector on the wheel: quadratic in the source and linear in the branch is a third derivative in primitive E_1 coordinates and a second variation here. The two are related by the chain-rule bridge of §G.6.8 and are not competing locations (F7f).

§C7. Democracy and isotype independence

Theorem 1. A Hermitian block commuting with the cyclic shift is circulant, $K = dI_3 + hR_3 + h^*R_3^\dagger$, so forward coefficients are equal and $w_{12} = w_{23} = w_{31}$. Democracy is a matrix consequence, not a postulate. ■

Theorem 2 (isotype-independent equidistribution). For any one-dimensional character χ_k and any equivariant isometry into $\mathfrak{B}$, the image lies in the χ_k isotype spanned by $v\_k = (1, \omega^k, \omega^{(2k)})/\sqrt{3}$, and **every branch component has modulus $1/\sqrt{3}$**. Verified for all three k. ■

This theorem does real work in §18, where the void descent is shown to land on a *non-trivial* isotype. The amplitude survives that relocation; the phase does not.

§C8. Norm transfer

Under role-metric isotropy (Schur on the irreducible fundamental) and shared closure profile, $3|\zeta|^2 = b^2$, so the **transported** modulus is $b/\sqrt{3}$ — the modulus is required, ζ being complex. The **physical** coefficient is $|\zeta_{phys}| = \rho_{mult}^{physical} \cdot c_{ax}/\sqrt{3}$, which coincides with $b/\sqrt{3}$ only if $c_{ax} = b$ and $\rho_{mult}^{physical} = 1$ (FD-3). The interpretation of b is load-bearing: it must be the total common-mode norm *before* branch selection.

§C9. Convention invariance and the BCH residue

Theorem 3. $V = \exp(-\lambda\Omega_0 + \tfrac{1}{2}\Omega_q)\cdot\exp(+\lambda\Omega_0 + \tfrac{1}{2}\Omega_q)$ depends only on $\boldsymbol{\zeta \coloneqq \lambda z}$. Certificates bit-identical across three conventions; the historical convention error executed exactly the half-strength branch. ■

$\log V = \Omega_q - (\lambda/2)[\Omega_0, \Omega_q] + O(\Omega_0^2)$, $[\Omega_0(z), \Omega_q]_{13} = -az$, $\boldsymbol{\Delta_{13} = \tfrac{1}{2}a\zeta + O(\zeta^2)}$.

Exact commutator verified at residual 1.07×10^{-20}, diagonal purely imaginary and traceless. Curvature-induced residue $\Delta_{13}^{conv} \coloneqq (\log V(\zeta))_{13} - (\log V(0))_{13}$ has modulus 0.0026332 against leading 0.0026306, **+0.10%** — *not* the bare entry 0.0048910, which is seed-dominated.

§C10. The readout as a six-component source

$F(x|m) = F_0(x) - m\cdot M_\chi(x) + O(m^2)$, $g_{\mathfrak{B}} = \nabla_{\mathfrak{B}} M_\chi(x_0) \in \mathbb{R}^6$, and since $H_{\mathfrak{B}} = \kappa_{C_3} I_6$,

$\boldsymbol{x'(0) = g_{\mathfrak{B}}/\kappa_{C_3}}$, with $t_{12} = g_{12}/g_{23}$, $t_{31} = g_{31}/g_{23}$, $\theta_h = \arg g_{23}$, $\rho_{mult}^{reduced} = \sqrt{(\Sigma|g|^2)}/(\kappa_{C_3}\, b/\sqrt{3})$ — **a reduced-model source normalisation relative to the declared $b/\sqrt{3}$ transport, not the full-carrier physical survival factor $\rho_{mult}^{physical}$**,

inverse carrying $e^{i\theta_h}/N$ with $N = \sqrt{(1+|t_{12}|^2+|t_{31}|^2)}$ — **1/N, not N**. Bijective only on $g_{23} \neq 0$; the complementary four-dimensional locus is where the readout does not select the visible branch, so the chart presupposes that it does.

§C11–§C12. The conditional numerical return

Supplied with $\zeta_{C_3} = (b/\sqrt{3})e^{5i\pi/6}$:

$$|V_{CKM}| = \begin{bmatrix} 0.974795773 & 0.223070728 & 0.003556877 \\ 0.222929780 & 0.973991276 & 0.040537720 \\ 0.008689898 & 0.039754796 & 0.999171678 \end{bmatrix}$$

$J = 2.996767765\times10^{-5}$, $\delta = 73.054785°$, $(\alpha, \beta, \gamma) = (84.527191°, 22.452580°, 73.020229°)$ summing to 180° exactly, $\|V^\dagger V - I\|_F = 2.50\times10^{-16}$, det V = 1 identically (a theorem of tracelessness, so residuals certify arithmetic only).

§C13. Empirical confrontation

Quantity	Transport	Benchmark	Miss	Pull
‖V_us‖	0.223071	0.22431 ± 0.00085	−0.55%	−1.46σ
‖V_cb‖	0.040538	0.0411 ± 0.0012	−1.37%	−0.47σ
‖V_ub‖	0.003557	0.00382 ± 0.00020	−6.89%	−1.32σ
‖V_td‖	0.008690	0.0086 ± 0.0002	+1.05%	+0.45σ
J	2.997×10^{-5}	$(3.12 \pm 0.125)\times10^{-5}$	−3.95%	−0.99σ
α	84.53°	84.1° ± 4.15°	+0.43°	+0.10σ
β	22.45°	22.6° ± 0.44°	−0.15°	−0.34σ
γ	73.02°	65.7° ± 3.0°	+7.32°	+2.44σ

Two declared metrics: **χ²_all = 11.327 (bare-generator curvature, 150°, eight-entry set)**, **χ²_curv = 8.982** (six curvature-sensitive entries, dropping inert |V_us| and inherited |V_cb|). No percentage quoted without naming its metric.

With $\Omega_0 = 0$ the seed alone gives α = 127.98°, γ = 32.37°, J = 1.844×10^{-5}, |V_ub| = 0.003931: the residue rotates a badly distorted triangle into the correct CP sector and moves J by +63%, spending some of the previously good |V_ub|. One complex amplitude cannot independently control the CP area and the short side.

Part C3 — What the extension does and does not settle

Position under the spine. The boundary tabulated here is the extension's, not the programme's. Where the controlling RRHF return supplies the same object, the programme-level grade is the Master Ledger's (R6, and the **[E]/[P]/[X]** tags below).

§C20. Four normalisations

#	Normalisation	Fixed by	Status
1	Representation-geometric split $1/\sqrt{3}$	Theorems 2 and 6	**Exact** democracy factor, isotype-independent; **recovered in two equivalent representations, not two independent determinations of physical amplitude** (§18.3)
2	Projector convention $\rho\_read^{\Pi_{23}} = \lVert\Pi_{23}Z_0\rVert/\lVert Z_0\rVert$ (the curvature readout P_vis^curv, **not** the frozen P_Y^RRHF)	choice of readout map	Convention; = $1/\sqrt{3}$ for Π_{23}
3	Multiplicity survival ρ_mult^physical	full-carrier character census	**Not identified**

#	Normalisation	Fixed by	Status
4	Readout concentration	channel type	**Not identified**

Arithmetic correction retained. $\Pi_{23}Z_0 = Z_{23}/\sqrt{3}$, so $\rho\_read = 1/\sqrt{3} = 0.577350$, **not 1** — below the assumed band floor $\sqrt{(2/3)}$. An orthogonal projection does not rescale the component it keeps, so ρ_read measures discarded norm and does not enter the transport; only ρ_mult^physical enters the physical transport, via ζ_phys = ρ_mult^physical·(c_ax/√3)·e^(iθ_phys), which reduces to b/√3 only under the AX1 comparison values c_ax = b and unit retention. ρ_read and ρ_mult^physical are different objects: ρ_read is a discarded-norm ratio for a chosen projector, ρ_mult^physical a physical survival factor. **No readout theorem relating them has been supplied.**

The whole-commitment repair remains excluded. $A_{23} = |Z_{23}\rangle\langle Z_0|$ gives ρ_read = 1 but returns |ζ| = b, a factor √3 larger, destroying the Theorem 2 split: |V_ub| = 0.004386, |V_td| = 0.010587, $J = 3.8382\times10^{-5}$, γ = 98.37°, χ^2_all = 304.91. The transport of §11 belongs to the ordinary-projection reading.

§C21. Scalar-readout no-go — retained, and now doubly restricted

Theorem 8

The realised scalar closure-maintenance mechanism produces branch support but no relative phase current, so g_A23 = 0, h_{23} is real, and **arg ζ ∈ {90°, 270°}, never 150°**. Numerically: 90° gives $J = +4.1501\times10^{-5}$, γ = 55.41°, χ^2_all = 1530.8; 270° gives $J = -4.6562\times10^{-6}$, γ = −9.08°, χ^2_all = 9259.7. **Closed negative for that mechanism. ■**

Proposition 9

Treating the weak complement as the anti-linear structure map J = εK with $J^2 = -I$ (which is phase-independent under ε ↦ e^(iθ)ε, so the Frobenius–Schur indicator selects no phase) causes the C_3 character to cancel in the direct complement–cycle square, again yielding {90°, 270°}. **The missing phase is not recovered by retyping the old scalar composition. ■**

Corollary 9.1 (reclassification, twice restricted)

Theorem 8 and Proposition 9 were originally stated as programme-level negatives. They have since been restricted twice, and both restrictions must be carried together.

First restriction (void axle). The later axle construction returns a quadratic character descent and source-phase candidate, so the reduced scalar no-go does not exclude that different structure. **It does not return a physical phase mechanism or value**, so the no-go does not apply to the extended marked-history action.

Second restriction (cumulative reconciliation). The no-go concerns a **reduced scalar package** that does not consume RRHF-1 or the graded source-jet chain. The cumulative corpus returns the charged tensor and hence V_CKM by a route in which no scalar readout of this kind appears at all. Theorem 8 therefore has no purchase on the programme, and by R6 may not be quoted against it.

Correct scope: the reduced scalar readout cannot return the CKM phase, and retyping the complement does not repair it. Within the curvature extension the axle supplies the quadratic character descent and a source-phase candidate; the physical carrier, holonomy and value are not returned. At programme level the CKM phase is returned conditionally by RRHF-1, and what remains open there is the primitive orientation and the admitted branch — which is §18.2's finding reached independently.

§C22. Empirical status of the source-phase candidate

The axle produces a **source-locked quadratic character descent and source-phase candidate**. One carrier reading would give 150° branch-independently; the other would give 150°, 30° or 270° according to the unreturned visible component (§18.2). **The 150° numerical branch is therefore a legitimate post-freeze conditional test, not yet a unique return**, and on that footing the one-parameter diagnostic minimum near 148.362° is informative — 1.638° separation in a 360° range.

Two constraints on reading it. **Historical blindness cannot be recreated**; what is enforced is source-locking, and no CKM value entered the slow-mode selection, the D_6 calculation, the uniqueness proof or the frame map. And **J is a weak discriminant**: two of six census branches reproduce it to 1.97×10^{-4}, so the phase test must be run on |V_ub| and γ, not on J. Subject to §18.2, the phase may no longer be varied to improve the fit without abandoning the axle theorem.

§C23. Global orientation and comparison scale

Orientation. Same-action conjugation gives V^(−) = conj(V^(+)), preserving all moduli and |J| and flipping only sgn J, which does not show the two constructions physically equivalent. **The status has changed without closing.** IORC-1 supplies an action-derived mechanism by which one orientation is *realised*: the preparation law is reversal-symmetric and returns a conjugate pair, with the sign carried as a mark on the first commitment event and P(±) = ½ beforehand. So the orientation is now understood as **spontaneously realised rather than absent** — a real advance over "no mechanism exists". It is still not *determined*: the law prefers neither branch, and the post-commitment lock that would make a realised sign persist is explicitly open, with a known record-amplitude escape E_RB(L) = 9w_R e^(−2L) → 0. **The sign of J therefore remains undetermined**, and the preferred circulation must still be tied to the physical CP convention rather than to an index orientation.

Scale. |J_q| runs +5.06% between 243.722807638 GeV and 5.80 TeV against a −3.95% miss, so **the comparison scale is not fixed** and the J and γ residuals are not yet interpretable as branch properties. What is prohibited is choosing a scale, scheme or threshold because it improves agreement — not evolving in either direction.

§C24. The boundary, tabulated

Scope note (R6). Every "not identified" or "not returned" entry below is a grade **on the proposed curvature correction**. Where the controlling RRHF return supplies the same object, the programme-level grade is the Master Ledger's, not this one. Entries are marked: **[E]** = extension-scoped, may be moot if RRHF already supplies it; **[P]** = holds at programme level, corroborated by the Master Ledger; **[X]** = exact and construction-independent.

Object	Status
Scalar branch Hessian, circulant return, democracy	**Exact** in the minimal cell **[X]**
Isotype-independent equidistribution	**Exact [X]**
Canonical W_6 graph, twelve edges, six cycles, **Laplacian spectrum** {0,2,2,4,4,5,7} and the k=±1,±2 eigenspaces	**Exact graph-theoretic return [X]**
Quoted **transition** spectrum {1,1/2,1/2,-1/4,-3/28,0,0}	**UNRESOLVED** — no standard operator reproduces it; $\hat{T}$ needs an explicit definition (§14.1.1, debt A0b)
Identification of the EX2F-selected mode with the k=±1, L=2 eigenspace	**CLOSED by Stage I** — overlap 1.0000000000000000 with the canonical k=−1 harmonic, $\lVert Lv-2v \rVert = 1.75\times10^{-15}$, hub amplitude 7.2×10^{-17}, half-turn parity −I (§7.0)
Hom(Sym^k E_1,E_2) = 0 for all odd k (**$\mathbb{Z}_2$ parity, Cor AX-1.1/1.5b**); unique quadratic axle at that order (**D_6**); Q(u) democracy	**Exact source-locked group return [X]**
Whether the action installs **only** the quadratic even channel, or also a quartic	**Moot for the presently tested action** — exactly family-rotation invariant, so neither the quadratic nor the quartic family channel is exposed; reopens only after a nonzero generation-sensitive coupling is installed (§G.6.9). (Cor AX-1.5b) the Born case has an upper-bound theorem, this side has not attempted one
$\hat{h}^2$ = O_q as the statement that the descent is quadratic	**Structurally returned** — replaces the owed phase grammar
$\zeta = \lambda z$, $\Delta_{13} = \frac{1}{2}a\zeta$, det V = 1, unitarity, source bijection	**Exact [X]**
Amplification ordering $-a\zeta$ / $+b\zeta$ / 0	**Exact at leading order**

Object	Status
Freedom-budget bound (51.3% / 32.8% survive)	**Exact numerical bound** on the curvature source space **[E]**
Source isotype	**Non-trivial (ω or ω̄)** — corrects k_src = 0 **[E]**
Quadratic character descent and source-phase candidate	**derived**; physical carrier, holonomy and value **not returned [E]**
Phase *value* 150°	**Conditional candidate generated by the axle; not uniquely returned until the three selections are derived** (§18.2, §18.4); carrier **component-preserving; physical readout not derived, local metric phase-isotropic** (§G.6.10) **[E]**.
c_ax, and whether c_ax = b	**Not returned [E]** — may be moot if the RRHF route supplies the tensor directly
Complete full-carrier internal C^{50} lift (local Q-sector form derived, AXD-2)	**Not returned [E]**
P_vis^(curv), visible branch, ρ_mult^physical, concentration	**Not identified [E]** — distinct from the frozen P_Y^(RRHF) of Part I §2, which the controlling action does use
sgn J	**Undetermined [P]** — but now with a mechanism: spontaneously realised on the first commitment event (IORC-1), not law-selected; persistence open pending the maintenance lock
Comparison scale	**Not fixed [P]**
Independence of the curvature construction from RRHF	**CLOSED NEGATIVELY [P]** — shared frame architecture and shared generator (a, b, c, φ). There is one construction, not two routes (§0.1)
Whether the Ω_0 block is independent constitutive information	**OPEN [P]** — decided by the source-clean differentiation test of §0.1.3, not by matrix comparison
Inherited (a, b, c, φ)	**Conditional cumulative derivation; primitive uniqueness open [P]** — a and φ carry geometric *and* conditional action ancestry (RRHF blind embedding); b and c are exact consequences of (a, D, η, χ); the live debt is the primitive provenance of those operands (§0.2). ‖V_us‖ carries −1.46σ, inert to everything above

---

Part C4 — The extension's installation contract

§C25. The chain

canonical W_6 graph —**EX2F/EX2G mode identification**→ oriented k=±1 mode E_1 —**Q**→ generation mode E_2 —**I_int**→ faithful internal weak-doublet carrier —**P_vis^curv**→ visible CKM branch.

The required object is no longer an unspecified irreversible oriented extension: PFDMI supplies the oriented mode and the axle supplies its unique descent. What remains is installation.

§C26. What this route must return, without CKM input

Scope. These are the obligations of the **proposed curvature correction**. Several may be moot: items 2–6 concern a readout the controlling RRHF return does not pass through. The contract is retained because the curvature block may be a **differentiated component of the RRHF parent**. Since disjointness is closed negatively (§0.1), items 2–6 should be tested by the source-clean differentiation of §0.1.3 rather than treated as obligations of a potentially independent route.

Canonical-mode identification — CLOSED by Stage I (overlap 1.0000000000000000). The live contract begins at item 1.

1. **Nonzero Hessian insertion** — the coefficient of the unique quadratic channel is nonzero and regulator-stable.
2. **Faithful internal lift** — a canonical intertwiner I_int couples the spatial D_6 axle to the $\mathbb{C}^{50}$ endpoint/current carrier and removes the relevant internal flat directions.
3. **Physical norm** — the action returns c_ax and *tests* rather than assumes c_ax=b.
4. **Readout selection** — P_vis^curv returns the visible branch and separately declares multiplicity survival, concentration and discarded norm.
5. **Phase-carrier declaration** — which object inherits the branch phase, reading (i) or (ii) of §18.2, justified from the action. Under (ii) this is not separable from item 4, the readout selection.
6. **Chirality and root provenance** — that EX2G's preferred circulation and the action-to-frame type order fix their respective bits without data contact.
7. **Global orientation** — the preferred circulation tied to the physical CP sign.
8. **Comparison scale** — all observables and benchmarks transported to one declared scale and scheme.
9. **Witness-degeneracy break** — a demonstration that the full return breaks the relevant Theorem 72.1 witness freedom rather than selecting one member of it.

§C27. What ASAC-1 still forbids

Overall action scale cannot supply the phase or any dimensionless mixing ratio. The axle is admissible precisely because it supplies new dimensionless representation structure — a unique quadratic intertwiner, not a coefficient adjustment. **c_ax controls amplitude, not direction:** changing c_ax controls magnitude only; it cannot by itself select θ_phys.

§C28. The unified object — updated status

The earlier conjecture that one oriented irreversible object might underlie time, branch direction and flavour phase has received a **structural pass**: EX2F supplies the irreversible jump/no-jump process and oriented mode, EX2G the attaching map and directional preference, EX2H the D_6 wheel and antipodal axes, and the axle theorem the unique descent. The connection is no longer an analogy. What remains conjectural is the **faithful physical installation** of the spatial axle inside the weak-doublet Hessian and readout.

§C29. Prohibitions

No coefficient, isotype, root, chirality, branch, phase-carrier reading, leakage amplitude or seed revision may be selected because it improves a measured quantity. Named for detectability: ρ_mult^physical ∈ {0.904, 0.957, 0.8371, 0.826633, √3/2}; θ ∈ {148.362°, 147.72°, 145.86°}; the target source κ_C_3(b/√3)(0, e^(iπ/3), 0); the data-preferred source of §31; and **reading (ii) of §18.2 selected because it can be made to give 150° for the branch that works.** No occupancy premise may remove a channel.

The wheel invariants |E|=12, β_1=6, dim E_Perron=1, and the charged-lepton factors 1/12, 7/6, (7/6)e^(−8/3) may not be inserted into c_ax, ρ_mult^physical, ζ, the CKM seed or the phase rule by numerical analogy. They become CKM inputs only through a typed action-derived map.

Part C5 — Premises, falsifiers, debts

Position under the spine. Debts are split by whether they bind the programme or only the extension. The programme-level class is short and is the Master Ledger's; the extension class is this chapter's own, and is discharged — or made moot — by the fork of §B.

§C30. Premises

Transport. PH-CYC · WD-EQV · PH-SAME · PH-GAP · PH-INH · FW-CORE. **Axle. VX-W6** ***(DISCHARGED by Stage I**, §7.0 — the identification is executed, not premised; the verified Laplacian supplies the eigenspace and no eigenvalue of the unresolved $\hat{T}$ is used)* · VX-MODE the K=7 law selects the sixfold slow mode · VX-CIRC EX2G fixes the preferred circulation without data contact · VX-GEOM EX2H returns the D_6 hub-plus-rim carrier with three antipodal axes · VX-ROOT action-to-frame type order selects the positive half-lift · **VX-CARRIER** the declared object inherits the branch phase (§18.2, undeclared). **Not identified.** C3-MULT · AMP-RDO · CHAN-TYPE · c_ax and its identification with b.

§C31. Falsifiers

(F4, F5, F7b and F7g are reproduced in the main paper as the falsifiers governing its own claims.)

ID	Fires if
F1	An object inside the Theorem 72.1 witness freedom is described as returnable by evaluation of the pre-PFDMI action; or Theorem 72.1 is applied as though it included later PFDMI information; or it is set aside without showing which witness degeneracy the new action breaks.
F2	A channel is removed through an unreturned occupancy premise.
F3	An absolute normalisation is read from a scalar Schur form, or an identity readout is cited as physical content.
F4	**Any value is selected for its effect** — see §29.
F5	A normalisation is quoted without naming which of the four of §20 it is; or $\rho = 1$ is claimed from an ordinary Π_{23} projection; or ρ_read and ρ_mult^physical are conflated; **or a test condition is written with $1/\sqrt{3}$ in place of ρ_mult^physical·$(1/\sqrt{3})$, which presupposes ρ_mult^physical = 1 and shifts the magnitude by 15% at the band floor** (§G.6.8).
F6	An anchoring channel returning ρ_read = 1 is adopted while $\lVert\zeta\rVert = b/\sqrt{3}$ is retained.
F7	**A linear current-to-generation descent is used despite $\mathrm{Hom}(E_1, E_2) = 0$** — which holds under the half-turn parity alone, so the escape "our carrier lacks D_6" is not available (Cor AX-1.1).
F7a	Another quadratic relation in the corpus is grouped with the axle without exhibiting its group, showing the lower power forbidden, **and stating the bilinear type of the survivor** — holomorphic squares and sesquilinear squares are not the same object (Cor AX-1.3, AX-1.4).
F7b	The axle is described as the *unique* channel without the qualifier "at lowest admissible order", while the even tower above it remains unbounded (Cor AX-1.5b); **or the tower is dismissed as an irrelevant-operator ambiguity excused by universality**, which the case-3 test refutes for this sector (§G.6.4).
F7d	A quantity conjugate to the varied coupling is offered as evidence that the coupling matters, or a symmetry-forced marginal is offered as evidence that it does not (§G.6.6).
F7g	A branch is selected by taking the orthogonal complement of an "inherited" support that omits $\mathfrak{B}_{23}$ — the (2,3) entry $b = 0.0405$ is nonzero, so that complement is empty and the move is the one withdrawn in §18 (F2).
F7f	The primitive-coordinate location of c_ax (cubic in E_1, a third derivative) and its descended-coordinate location (a second variation on the branch carrier) are treated as rival claims rather than as one object in two coordinate systems related by the chain-rule bridge of §G.6.8.
F7e	The Ashkin–Teller single-fact/two-fact numerics are cited as the ground of the universality split when the structural version is available — first-fact invariance follows from unitary invariance of Haar, and multi-fact dependence from non-vanishing [U, P_A] (§G.6.7).
F7c	The two links of §G.6.5 are conflated — *whether* a quartic coefficient survives coarse-graining (open, an RG question) is treated as settling *what it does* once present (settled by direct evaluation, scheme-independent), or the reverse. An emergent-time argument bearing on the first may not be quoted against the second.
F8	**The unique quadratic axle is called a full CKM derivation** before its nonzero Hessian coefficient, faithful internal lift, readout and phase-carrier typing are returned.

ID	Fires if
F9	**A phase conclusion is stated without declaring which object carries the phase** (Rule R4, §18.2).
F10	The raw axle norm $1/\sqrt{3}$ is conflated with the physical coefficient c_ax or with b; or the two $1/\sqrt{3}$ routes are called independent determinations of the scale.
F11	The source is still described as occupying the trivial C_3 isotype after §18.1.
F12	The D36 preferred circulation is identified with the observed sign of CP violation without proving the physical orientation map; or sgn J is claimed as a prediction or as a mere convention.
F13	The 150° phase is described as an unexplained fit, or as a complete physical prediction despite §18.
F14	An extension supplies only a scale or overall coefficient, contrary to ASAC-1.
F15	A comparison is made without a declared scale, or a scale/scheme/threshold is chosen because it improves agreement.
F16	The J agreement is cited as evidence for the isotype or the phase — J is isotype-blind at 1.97×10^{-4}.
F17	A returned g_𝔅 removes substantially more of the declared sum than §32's bound permits.
F18	A single envelope percentage is quoted without naming its metric.
F19	Depth adjacency is claimed to give $g_{12} = 0$; it forbids only $\mathfrak{B}_{31}$, and no unoriented depth-local functional can distinguish the adjacent edges.
F20	Band monotonicity fails under refinement, invalidating the retention certificate.
F21	The EX2F-selected mode is called the canonical k=±1 wheel eigenmode without proving the identification, or a different eigenspace is silently substituted.
F22	The Perron mode, twelve edge links, six cycle modes, three axes or $1/\sqrt{3}$ branch split are conflated; or a charged-lepton 1/12, 7/6 or (7/6)e^(−8/3) factor is imported into CKM without a typed action-level map.
F23	An eigenvalue of $\hat{T}$ is cited as a substrate return before $\hat{T}$ is explicitly defined; or the quoted hub entry -3/28 is used without resolving the discrepancy against -3/4 (§14.1.1).
F24	The EX2F mode is identified with E_1 on antipodal parity alone, ignoring the k=3 alternating mode which shares that parity (§14.1.2b).
F25	A nonzero c_ax is claimed from the transition channel while the target eigenspace lies in ker$\hat{T}$, without naming the stationary or constitutive term that supplies it (§14.1.3).
F26	The RRHF and curvature CKM returns are cited as corroborating one another — they are **not disjoint** (§0.1); or the two calculations are described as two physical routes rather than two realisations of one construction (R7); or a matrix printed in a downstream audit is compared against this paper as though it were independent when it reproduces this paper's own curvature result; or a "not identified" grade from this paper is quoted at programme level against the cumulative stack (§0).
F27	The **AX1 diagnostic amplitude** is called independent of the shared ancestry despite being set by the shared coefficient b, **or that diagnostic ancestry is transferred to the unreturned physical amplitude** ρ_mult^physical·c_ax/√3 (§0.1.3).

ID	Fires if
F28	$A = 0.800$, b or c is listed as an independent unexplained input; or $A = 2\chi$ is cited as a derivation rather than an arithmetic coincidence in the frozen package (§0.2).
F29	The $K = 7$ agreement $6K-2 = 2\cdot C(6,3)$ is used as a structural bridge between the two doorway formulas; it holds at $K = 7$ alone (§0.2.2).
F30	$\eta = 3/5$ is used without naming the measure "democratic" refers to on a carrier with unequal isotype multiplicities; or the residual carrier's covariance group is left unstated while a decomposition-dependent claim is made about it (§0.2.3.1).
F31	The axle's recovery of the generation character $O_q = e^{(2\pi i/3)}$ is treated as deriving the role-odd holonomy φ, or the two 120° values are conflated with each other or with arg ζ = 150° (§0.2.4).
F32	A later calculation downgrades an earlier result without evaluating **the same typed object, on the same carrier, with the same physical readout** — the rule this paper itself violated at §0.2.3 before correction.

§C32. Open debts

Axle-installation class. A0 *(CLOSED — Stage I, §7.0: overlap 1.0000000000000000 with the canonical k=-1 harmonic at L=2, hub amplitude 7.2×10^{-17}, half-turn parity -I)* · **A0b** give $\hat{T}$ an explicit definition and resolve the hub-entry discrepancy of §14.1.1 · **A1** the master action contains the unique quadratic $Sym^2E_1\to E_2$ channel with nonzero regulator-stable coefficient · **A2** the faithful internal $\mathbb{C}^{50}$ lift, removing the relevant post-gauge flat directions · **A3** the physical coefficient c_ax, testing c_ax=b · **A4 — physical readout and phase carrier OPEN.** The completion jet is component-preserving, but $P_{vis}^{curv} = P_{23}$ is not established, and the local completion metric is exactly phase-isotropic. **No local phase value is selected** (§G.6.11). · **A5** provenance of the chirality and root bits (§18.4) · **A5b** cite the single-record Born normalisation that fixes the RSU likelihood exponent to $\beta=1$, which composition-compatibility alone does not (§7.0.1.2a) · **A5c** *(CLOSED NEGATIVELY — HCMR-XI-2: the oriented D36 kernel has nonreal eigenvalues, so no positive quasi-Hermitian metric exists on a single branch; §7.0.1.2.1)* · **A5d** *(ADDRESSED — HCMR-XI-3 attacks structure not coefficients and reports cosine 1.000000000000000; §7.0.1.2.4)* · **A5e** settle whether Z=(1,5/2,298) is independent of the HM4 gauge quadratic form or a Schur reduction of it (§7.0.1.2.4) · **A5f** settle the same question for the neutral selection: is S_D independent of $\sqrt{(M_\nu M_\nu^\dagger)}$, or defined as it upstream (§7.0.1.2.5)? · **A2b** reconcile the 49-against-50 field-multiplicity count between RUDM-1's twelve components and AXD-2's class weights, located at the missing N_R_triv — the Dirichlet predictive law depends on which is right (§G.6.9.1) · **A5g** reconcile λ_min of the right complement, 1.527103739469 reported against 1.527494392135 implied by the charged spectrum, or state the differing regulator levels · **A6 — quartic-channel status. Moot for the presently tested action**, which is exactly family-rotation invariant, so no family-sector object is visible — the quadratic axle included. The question **reopens only after the missing generation-sensitive completion coupling is installed**, and must then be tested against that new coupling rather than the present arrays. Chronology in §G.6.3–§G.6.9.

Readout class. R1 derive P_vis^curv and the visible branch · **R2** return multiplicity survival and concentration separately · **R3** prove the returned readout breaks the relevant Theorem 72.1 witness freedom.

Orientation class. O1 tie the preferred circulation to the physical CP sign — **partially advanced** by IORC-1, which supplies the realisation mechanism but not the law-level preference; what remains is (i) the post-commitment orientation lock that prevents erasure or reversal, and (ii) the same-action carrier map inserting the selected character into the role-odd operator (= I1d).

Programme-level class (bind regardless of route). P1 complete and uniquely derive H_CMR · **P2** independently select ξ, reconciling the four distinct length objects rather than averaging them · **I6** *(closed negatively — the routes are not disjoint)* superseded by **I6a** the two-stage test of §7: Stage I freezes the completed H_CMR and asks for the coefficient, lift, readout, retention and phase carrier without CKM input; Stage II inserts the Stage-I return into RRHF and re-differentiates. RRHF tests integration, not microscopic forcing · **I1 — Primitive charged-generator provenance.** Derive from a frozen primitive marked-history action the charged readout projector, the normalised 1↔2 overlap, the residual response metric, the one- and two-step attenuation operators, the weighted non-adjacent continuation operator, and the physical holonomy; then test post-freeze whether these return $(a, D, \eta, \chi, \varphi) = (9/40, \mathrm{diag}(1,2,4), 3/5, 2/5, 2\pi/3)$. **A = 0.800 is not a separate input** — it is the exact consequence $A = b/a^2 = 4/5$ (§0.2) · **I1b** state whether the closure-transport residual carrier is D_6-covariant or only C_3-covariant — the (3+2) split is C_3-covariant, and the question only bites if the carrier inherits the wheel's full D_6 (§0.2.3.1); and separately name the measure that "democratic" refers to on a carrier with unequal isotype multiplicities · **I1d** close the typed bridge from the generation character O_q to the role-odd holonomy φ in the seed entry — they share the value 120° and are different objects (§0.2.4). **Now the single named blocker on the orientation chain:** IORC-1 supplies the character's sign, and this map is what would carry it into φ · **I1e** restate IORC-1's two mis-specified residuals — the C_3 character check against the conjugate eigenvalue, and the conjugate-kernel check against the plain transpose rather than the π-reversal (§0.2.7) · **I1c** supply a typed bridge between "available mixing channels" (6K−2) and "twice the balanced-support shell" (2·C(6,3)), noting these agree at K = 7 only (§0.2.2) ·

Route-level class (bind this construction only). I2 the −1.46σ |V_us| miss · **I3** comparison scale, scheme and thresholds · **I4** covariance-aware test on independent observables · **I5** analytic upgrade of the retention certificate.

Retained quantitative bound. Granting the whole unreturned six-component source, **51.3% of χ^2_all and 32.8% of χ^2_curv survive**, with the largest survivor |V_us|, inert to every component. The missing installation can at best halve the declared disagreement — a falsifiable prediction about a calculation not yet performed, and simultaneously the strongest evidence the construction is not carrying hidden freedom.

Part C6 — Verdict of the curvature audit

The curvature extension is **unadmitted** here. Its exact content — the W_6 certificate, Hom_$D_6(E_1,E_2)=0$, the uniqueness of the quadratic axle and the threefold democracy of its normalised image — stands as mathematics independent of admission. Its physical content — coefficient, faithful internal lift, readout branch, norm and phase carrier — is not returned by the completed same-action Stage I projection and is decided by the sealed I6a test of §7.

PHYSICAL-INSTALLATION BLOCKERS: A1 ‖ A2 ‖ A3 ‖ A4 ‖ R1 ‖ R2 ‖ R3.

These are the debts that decide whether the axle is installed. Two others are deliberately **not** in that list, because they are different kinds of debt: **A0b**, the undefined $\hat{T}$, is a graph-operator provenance question and the axle does not depend on the disputed hub eigenvalue — Stage I identifies the source through the verified Laplacian instead; and **I3**, the comparison scale and scheme, bears on empirical confrontation rather than on installation. **A0 is closed** (Stage I).

A.1 Curvature extension (Part II)

Quantity	Expression	Value
Canonical closure graph	$W_6=C_6\vee K_1$	7 vertices, 12 edges
Cycle rank	β_1=	E
Graph-Laplacian spectrum	spec(L)	{0,2,2,4,4,5,7}
Transition spectrum, quoted	spec($\hat{T}$)	{1,1/2,1/2,-1/4,-3/28,0,0} — **not reproduced**
Closest standard operator	I-1/4L	{1,1/2,1/2,0,0,-1/4,-3/4} — six of seven entries match
Canonical E_1 candidate	k=±1 boundary harmonics	dim 2, L=2, half-turn parity -1
Canonical E_2 realisation	k=±2 boundary harmonics	dim 2, L=4, parity +1; lies in ker$\hat{T}$ **under the candidate I-1/4L only** — the quoted spectrum's zero sector cannot be identified with E_2 until $\hat{T}$ is defined
Antipode-even rim sector	k=0,2,4	= trivial ⊕ E_2, dim 3
Antipode-odd rim sector	k=1,3,5	=E_1⊕ alternating, dim 3 — parity alone insufficient
a, b, c, φ	9/40, 81/2000, 243/100000, 2π/3	0.225, 0.0405, 0.00243, 120°
‖ζ_C_3‖	b/√3	0.023382685902
ζ_C_3	−b/2 + i·b/(2√3)	−0.020250000000 + 0.011691342951 i
‖Δ_{13}‖ leading / converged	ab/(2√3) / ‖Δlog V_{13}‖	0.002630552164 / 0.0026332 (+0.10%)
⟨χ_E_1, χ_E_2⟩	Theorem 4	1.48×10^{-16}
‖L(u)‖ linear pair-sum	Theorem 4	2.05×10^{-16}

Quantity	Expression	Value
$\langle\chi\_Sym^2E_1, \chi\_E_2\rangle$	Theorem 5	1 exactly
Q(u) components / norm	Theorem 6	modulus ⅓ each / 0.577350269190
democracy spread	Theorem 6	1.11×10^{-16}
Σ_a Q(u)_a	lies in E_2	0
R_3 eigenvalue on Q(u)	§18.1	$e^{(\mp 2\pi i/3)}$ — non-trivial isotype
ρ_read under Π_{23} (declared-P_{23} reduced-model formula, **not a physical readout theorem**)	1/√3	0.577350269
$\lVert V^\dagger V - I\rVert\_F$, $\lVert\det V - 1\rVert$	machine	2.4975×10^{-16}, 2.22×10^{-16}
J: supplied / 90° / 270° / 30° / anchoring	exponentiation	2.9968×10^{-5} / 4.1501×10^{-5} / -4.6562×10^{-6} / 2.9974×10^{-5} / 3.8382×10^{-5}

Appendix G — The invariant parameter budget of H_CMR on the minimal carrier

Scope and status. This is an H_CMR calculation, not a flavour calculation. It appears here because it **quantifies the obstruction identified in §7.0.1.1** — the programme-level non-uniqueness of which the flavour blocker is one face. It derives nothing about V_CKM and installs nothing. **Grade: exact, on the declared minimal carrier.**

What it is not. It is not a derivation of H_CMR, and no such derivation is available: the programme ledger exhibits two inequivalent marked-score lifts reproducing the same Hessian to 10^{-15} while differing by ≈1.7366 in relative norm. Writing down an action would be selecting one member of that continuum. **The correct response to a proved underdetermination is to characterise the family, and that is what follows.**

G.1 The quadratic sector is exactly six-dimensional

The wheel's automorphism group is D_6 of order 12, verified to preserve adjacency. On the vertex carrier $\mathbb{R}^7$, the space of D_6-invariant **symmetric** operators has

dim=6,

confirmed two independent ways: by direct rank computation on the commutant equations $MXM^T=X$ over all twelve group elements, and by Schur — since $\mathbb{R}^7=2A_1\oplus B\oplus E_1\oplus E_2$, the symmetric commutant has dimension $\Sigma\_i\ m\_i(m\_i+1)/2=3+1+1+1=6$.

An explicit spanning set, verified invariant and of full rank 6:

#	Generator
1	hub diagonal
2	rim diagonal
3	hub–rim coupling
4	rim nearest-neighbour
5	rim next-nearest
6	rim antipodal

The list is **complete, not illustrative**: these six span the whole invariant space.

G.2 Locality ladder, and where the Laplacian sits

Interaction range	Parameters
range 1 (hub diagonal, rim diagonal, hub–rim, rim nn)	4
range 2 (adds next-nearest)	5
range 3 (adds antipodal)	6

Expanded in the basis above, the graph Laplacian is

$L=(6,3,-1,-1,0,0)$, residual 2.9×10^{-15}.

The Laplacian is not a generic invariant operator: it is the shortest-range point consistent with the symmetry, with both long-range coordinates *exactly* zero. Those two zeros are an assumption of the reference construction, not a return.

G.3 Isotypic form, and what the corpus actually constrains

Every D_6-invariant symmetric operator acts as

λ_1 on E_1, λ_2 on E_2, λ_B on the k=3 mode, a 2×2 symmetric block on $2A_1$,

i.e. 1+1+1+3=6, matching G.1. Verified by confirming that the E_1, E_2 and B harmonics are exact eigenvectors of a random invariant operator.

This is the form in which the axle's inputs appear: λ_1 and λ_2 are single real numbers. Against them the certified corpus supplies

- E_1 is the slowest non-trivial mode — an **inequality**, and an open condition;
- the axle target is E_2 — a **sector label**, not a value.

Neither fixes a number. Of the six quadratic parameters, the certified data constrain one inequality between two of them.

G.4 The cubic sector, and where c_ax lives

From the character formula $\chi\_Sym^3V(g)=1/6[\chi(g)^3+3\chi(g)\chi(g^2)+2\chi(g^3)]$ averaged over D_6, the multiplicities in $Sym^3(\mathbb{R}^7)$ are

irrep	A_1	A_2	B_1	B_2	E_1	E_2
multiplicity	**13**	3	10	4	13	14

so there are **13 independent D_6-invariant cubic couplings**. Since dim $Hom\_D_6(Sym^2E_1,E_2)=1$,

c_ax is exactly one coordinate of a 13-dimensional space.

G.4.1 A genericity remark that reframes the axle conjecture

A nonzero c_ax is generic within the unrestricted 13-dimensional invariant cubic family. That gives no derivational presumption about the physical action — AXD-2 shows the presently tested action lies on a constrained, family-rotation-invariant subspace and does not expose c_ax at all (§G.6.9). With that qualification:

- **the axle's existence is generic in the unrestricted family** — but the tested action is not a generic point of it, so Stage I's non-return is **not** evidence of an absent calculation;
- **the axle's *value* is not** — c_ax=b is a single codimension-one condition.

The unrestricted invariant family makes a nonzero channel generic, **but this carries no physical-installation presumption**. The first action-level question remains whether **Stage I-A returns c_ax ≠ 0**. Only after that passes does the sharper AX1 conjecture become the codimension-one comparison c_ax = b.

G.5 The preparation sector is rigid only relative to a choice

IORC's minimiser is unique **given** the reference kernel T_0. The invariant freedom in T_0 itself is three parameters: one from the hub (p_hub→ hub, the rim shares fixed by symmetry) and two from the rim (p_→hub, p_→self, with the two rim directions forced equal by reflection and the row summing to one).

IORC uses a=1/7, b=c=d=1/4 — the **maximum-entropy** point of that family, with row entropies ln7=1.945910 and ln4=1.386294. Natural, but a choice.

The preparation sector is rigid given T_0, not rigid absolutely.

G.6 The budget

Contribution	Parameters
D_6-invariant quadratic couplings	6

Contribution	Parameters
D_6-invariant cubic couplings	13
reference-kernel freedom in the preparation sector	3
total on the minimal carrier	**≥ 22**

against which the certified corpus supplies

Constraint	Type
one-Bit condition $\Sigma=\ln 2$	one **equation** (fixes α)
E_1 slowest non-trivial	one **inequality**
axle target is E_2	a sector **label**

One equation and one inequality against at least 22 parameters.

This is a floor, not an estimate. The physical carrier is far larger than seven vertices — the internal $\mathbb{C}^{50}$ space, the 243 defect coordinates, the role and closure factors — so the true count is worse, and every additional carrier factor multiplies rather than adds.

G.6.1 What HCMR-XI-1 does and does not remove from this budget

The conditional law-level completion of 2 August closes two ambiguities that are **not** in this table — the eleven record-conditioned tilts and the copy law — and leaves every entry of the table standing. Its completed cost M carries no free coefficients, but it is *defined relative to* T_0, which is the 3-parameter row above, and it contains no wheel Hessian term at all, so neither the 6 quadratic nor the 13 cubic couplings are addressed.

Budget after HCMR-XI-1: unchanged at ≥22. *(Superseded in kind by §G.6.2: the admissible family is infinite-dimensional, and this finite count survives only as a statement about the D_6-invariant quadratic and cubic slice.)*

The same holds after the subsequent upgrade of RSU and CPR from adopted to derived. **A principle moving from assumed to forced changes its status, not the parameter count** — the couplings it does not mention remain free. And if the proposed V-metric refinement is adopted without V itself being returned by the action, the count **rises**, since the metric operator of a quasi-Hermitian Hamiltonian is not unique (§7.0.1.2b).

Not a criticism of that paper, which does not claim otherwise — it is why the flavour blocker survives a partial H_CMR closure: **different ambiguities.**

G.6.2 PSU-1: the budget is not finite

The count of §G.6 must be revised in kind, not in size. PSU-1 exhibits three explicit counterfamilies, all independently re-executed here, that reproduce every stored low-order return while remaining physically distinct.

(a) A third Fold-edge coupling the displayed action omits. Invariance under the two global $\mathbb{Z}_2$ flips and endpoint exchange leaves only $a = \sigma_i\sigma_j$ and $b = \omega_i\omega_j$, hence four edge orbits. The invariant functions $\{1, (1-a), (1-b), (1-a)(1-b)\}$ have rank 4 — confirmed — so after removing the constant the symmetry permits **three** independent couplings. The displayed action uses two. The third, $J_{\sigma\omega}(1-a)(1-b)$, is allowed and unfixed, and it is physically visible: at $J_\sigma = J_\omega = 1$ and $J_{\sigma\omega} : 0 \to 0.3$ the normalised edge-record distribution moves by TV = $9.886644211506\times10^{-3}$, which reproduces to twelve figures.

(b) An infinite Wilson-harmonic family with identical quadratic stiffness. For $S_t(\theta) = (Z - 4t)(1 - \cos\theta) + t(1 - \cos 2\theta)$, every member is local, orientation-even, periodic and nonnegative on $0 \le t \le Z/4$ — checked — with $S''(0) = Z$ for all t, while $S''''(0) = -(Z + 12t)$ differs. The holonomy distributions differ by TV = 1.212065193×10^{-1}, reproduced. **An infinite family of inequivalent plaquette actions returns the exact same gauge Hessian.**

(c) An infinite completion-amplitude family. $f_a(\varphi) = \sin\varphi + a\sin^5\varphi$ agrees with $\sin\varphi$ through fourth order and first differs at fifth — verified symbolically — so every stored completion jet is reproduced while the finite outcome distribution moves.

None of the three is removable by the physical quotient, since each changes a normalised finite outcome distribution.

The consequence for this table. §G.6 gave a finite floor of ≥ 22 on the minimal carrier and noted the true count is worse. PSU-1 shows it is worse **in kind**:

The admissible action family is infinite-dimensional, not merely larger than 22.

The finite count remains correct as a statement about the D_6-invariant slice of the quadratic and cubic sectors. It is not a bound on the admissible family.

G.6.3 The flavour tower is an instance of the same pattern

Controlling status after AXD-2. The discussion from §G.6.3 through §G.6.8 records the chronology of A6 **before** the family-blindness result. For the presently tested action **A6 is moot**, because neither the quadratic nor the quartic family channel is visible. The question reopens only after a nonzero generation-sensitive same-action coupling is installed.

Corollary AX-1.5b found that the even tower above the quadratic axle does not terminate — dim $\mathrm{Hom}(\mathrm{Sym}^4E_1, E_2) = 2$ — and asked whether the action installs the quadratic channel alone. **PSU-1's Wilson-harmonic family is structurally the same phenomenon in the gauge sector:** higher harmonics reproduce the same lowest-order stiffness and differ only above it.

So debt A6 is not a flavour peculiarity. It is one instance of a programme-wide pattern — **matching a low-order jet does not fix the tower above it** — and PSU-1's closing judgement applies verbatim to the axle:

Setting the higher couplings to zero is a minimality prescription, not a derivation.

This paper uses the quadratic axle alone. **By that standard, using it alone is currently a minimality prescription.** The remedy PSU-1 identifies is the same in both sectors: an RG fixed point whose irrelevant directions drive the extra couplings to determined values. Absent that, neither the quartic axle channel nor the higher Wilson harmonics are excluded.

G.6.4 Uniqueness was the wrong target — and the right one gives a sector-dependent answer

A methodological correction is accepted here before the finding that follows from it. A Standard Model derivation does not require a literally unique microscopic action. It requires a single physical equivalence class: actions related by gauge or field redefinition, differing only by irrelevant operators, or flowing to the same infrared fixed point are all acceptable, exactly as many lattice actions yield one continuum theory. The correct target is therefore

$\forall\ S \in \mathcal{A}$_Fold, RG_IR(S) ~ S_SM,

a **universality** theorem rather than a uniqueness theorem. The counterfamilies of §G.6.2 damage the derivation only if they fall in the third of three cases: (1) same infrared theory, (2) differ only outside the Standard Model domain, (3) **change Standard Model observables**. Only case 3 is a derivational problem.

That reframing is right, and it is testable rather than merely reassuring. Applied to the two sectors, it gives opposite answers.

The gauge counterfamily is plausibly case 1. Adding $(1 - \cos 2\theta)$ to $(1 - \cos \theta)$ is precisely Symanzik improvement: the two differ by irrelevant operators and flow to the same continuum theory, and $S''(0) = Z$ is unchanged so the gauge coupling does not move. Universality applies and no derivational problem arises.

The flavour tower looks like case 3. If a quartic contribution survives coarse-graining with a nonzero coefficient it enters the **same 1↔3 entry** as the quadratic axle, rather than sitting in a separate operator that could flow away independently. Computed on the actual selected mode:

quantity	value
$\|Q_2(u)\|$	0.577350
$\|Q_4(u)\|$	0.096225
relative size $\|Q_4\|/\|Q_2\|$	**0.166667**
C_3 character of Q_2	$e^{(-2\pi i/3)}$

quantity	value
C_3 character of Q_4	**$e^{(+2\pi i/3)}$ — the conjugate**

The quartic image lands in the **conjugate** isotype, so it shifts the phase rather than merely rescaling the magnitude — and the phase is what carries the entire empirical gain (§5.2.3). Admixing a quartic channel at relative source strength ε, scanning its unknown phase:

ε	shift in $\lvert\zeta\rvert$	χ^2_5 range over the quartic phase
0.00	0%	5.25
0.05	0.83%	4.72 – 5.85
0.10	1.67%	4.25 – 6.51
0.30	5.00%	2.99 – 9.87

A five-percent shift in the total axle norm — produced here by a thirty-percent relative quartic source amplitude — moves the diagnostic χ^2 by roughly a factor of three. Smaller admixtures move it correspondingly less: $\varepsilon = 0.05$ is a 0.83% norm shift and a χ^2_5 range of 4.72–5.85. Source amplitude and norm displacement must not be conflated. By the three-case test, the flavour tower is a genuine derivational problem: an admissible coefficient changes a Standard Model observable.

What this makes of debt A6. It is not the general question "does the action install only the lowest channel", which universality would excuse. It is the specific question **"is the quartic coefficient zero, RG-irrelevant, or experimentally constrained?"** — the three exits the reframing itself names. Until one is shown, the quadratic-only construction remains a minimality prescription, and the correct grade for the flavour sector is unchanged by the reframing even though the gauge sector's improves.

G.6.5 The emergent-time critique, and which link it hits

A subsequent correction argues that an external-time equilibrium RG cannot overrule a framework in which time is itself produced by the commitment process being coarse-grained. **The argument is legitimate, and it is accepted — but it targets a different link from the one §G.6.4 settled.**

The chain has two links.

(L1) Does the quartic coefficient survive coarse-graining? — an **RG** question. (L2) If it survives, does it change CKM observables? — a **direct evaluation**.

§G.6.4 answered L2. That calculation performed no renormalisation at all: it set $\zeta\_total = \zeta_2 + \varepsilon\cdot$(quartic image), pushed the result through the same exponentiation used everywhere else in this paper, and read off χ^2. **It is therefore ensemble-independent** — it does not care how the coupling came to be there, only what it does once present.

The joint statement is conditional, and sharper than either half alone:

If the quartic survives, it moves the observables (L2, settled). Whether it survives is exactly what an emergent-time RG must decide (L1, open).

So the critique makes debt A6 **more** urgent, not less. It does not show the concern was misplaced; it identifies the single test that could retire it.

G.6.5.1 Two qualifications on the critique itself

(a) The proposed cancellation is a possibility, not a mechanism. With $\Gamma\_t = \Gamma\_s/\lambda$, a K_4-dependence in the proto-time rate can cancel one in the commitment rate — but only if $\Gamma\_s(K_4) \propto \lambda(K_4)$ exactly. Checked over generic alternatives: linear-versus-linear with different slopes leaves a 40% residual dependence across the tested range, and matched slopes with differing curvature still leave 36%. **Only exact proportionality removes the dependence.** Absent a structural reason for it, invoking the cancellation is a fine-tuning rather than an argument, and the reweighting must be computed rather than assumed.

(b) The emergent-depth point is stronger than the critique states. It is offered as something that "could change the dynamic exponent and the effective scaling interpretation". The sharper version is decisive: **marginality of the four-spin coupling is a known two-dimensional feature of the Ashkin–Teller model.** In two dimensions the model has a line of critical points with continuously varying exponents, along which the four-spin term is exactly marginal — a c = 1 Gaussian-line phenomenon specific to 2D. In three effective dimensions the same model has ordinary critical behaviour, with that coupling relevant or irrelevant depending on the region and **not marginal at all**.

So if the third dimension is genuinely reconstructed from the Fold sequence, **the marginal verdict may be an artefact of running the calculation on a two-dimensional lattice**, and the physical answer would be a clean relevant-or-irrelevant determination rather than a marginal line. That is a stronger reason to redo the test than the time-reweighting argument, and it also means the redone test is more likely to give a decisive answer than the original.

G.6.5.2 Status after the correction

Object	Status
Gauge harmonics	irrelevant in the declared fluctuation basin
Completion deformation	irrelevant near the completion vacuum
Mixed Fold coupling	marginal **in the pre-commitment two-dimensional interface ensemble**, which may not be the physical one
Physical-time Fold universality	**not tested**
Standard Model observable universality	**not tested**

Object	Status
L2: a surviving quartic changes CKM	**settled, and unaffected by the critique** (§G.6.4)

The earlier reading — that a marginal interface coupling constitutes a universality failure — is withdrawn as too strong. What replaces it is narrower and better posed: a parameter survives in the reversible interface theory, and it has not yet been tested through commitment.

G.6.6 ETRG-1: the emergent-time test executed

The test called for in §G.6.5 has been run. Its central move is correct and its verdict is appropriately narrowed: the earlier claim that Fold critical universality fails physically is withdrawn, leaving *pre-fact* critical-line dependence established and *observable* universality open. Three of its findings are checked here, and two are weaker than the report presents them.

Confirmed, and more general than claimed. The equilibrium transfer process is pre-temporal: its Doob kernel $P_{xy} = T_{xy}\psi_0(y)/(\Lambda_0\psi_0(x))$ with $\pi_x = \psi_0(x)^2$ gives $\pi_x P_{xy} = \psi_0(x)T_{xy}\psi_0(y)/\Lambda_0$, manifestly symmetric in x and y whenever T is. **This is an identity, not a numerical result, and it holds for every symmetric transfer matrix** — verified on random symmetric matrices at detailed-balance residuals of $10^{-17} – 10^{-18}$. The consequence is stronger than stated: an equilibrium transfer calculation could never have produced a physical-time result **for any model**, so the correction to the earlier test is structural rather than model-specific.

Likewise the cancellation of a common hazard, $p_A = \gamma w_A/\Sigma\gamma w_B = w_A$, is exact by construction.

Weaker than presented (a): the single-cell result is a symmetry tautology. The Ashkin–Teller model carries $\mathbb{Z}_2 \times \mathbb{Z}_2$ acting independently on σ and ω. Under $\sigma \mapsto -\sigma$ the mark $q = \sigma\omega$ flips, so $P(q = +1) = P(q = -1) = ½$ is **forced for every K_4**, critical or not — verified exactly at $K_4 = 0, 0.2, 0.4$. The quoted residual of 1.28×10^{-15} is machine noise on an identity. It is therefore not evidence of universality, and would hold unchanged in a model where every other quantity depended on K_4. The report calls it "an important positive result"; it carries no information about the question at issue.

Weaker than presented (b): the two-fact correlator is conjugate to the varied coupling. Since $q_iq_j = (\sigma_i\sigma_j)(\omega_i\omega_j)$, the quantity $\langle q_iq_j\rangle$ is precisely the four-spin correlator whose coefficient **is** K_4. Verified: it moves from 0.098 to 0.541 across $K_4 = 0 \rightarrow 0.4$ on a small ring. That an operator's expectation responds to varying its own coupling is the minimum possible finding, and it would be remarkable only if it did not. The informative test is a quantity **not** conjugate to K_4 — a universal amplitude ratio, or a Standard Model observable — which is exactly what the report's own §8 correctly identifies as the remaining gate.

And the emergent-depth objection is untouched. ETRG-1 runs on two-dimensional strips of widths 4–7. §G.6.5.1(b) records that marginality of the four-spin coupling is a known two-dimensional feature of this model, tied to the $c = 1$ line, and that in three effective dimensions the same coupling is relevant or irrelevant rather than marginal. **That objection is not addressed by running the same 2D calculation with commitment appended.** If the effective dimension is three, the entire marginal line — and with it the K_4 family — may not exist.

Net effect on this chapter. ETRG-1 does not change the L1/L2 split of §G.6.5. It confirms that emergent-time reparameterisation alone cannot remove a dependence once it enters relative record weights, which is the honest core of the result, and it correctly relocates the decisive question to the source-native commitment map. Before AXD-2, debt A6 remained open, with its test now specified twice over: the commitment map of ETRG-1 §8, **and** the effective-dimension question that neither test has yet addressed.

G.6.7 RUDM-1: the question reduces to a commutator

The successor calculation closes the preparation measure, the commitment projectors and the relative hazards, and in doing so **reduces the whole K_4 question to a single algebraic condition.** Independently checked here.

Bookkeeping and structure, confirmed. The twelve component dimensions sum to 49, matching the $\mathbb{C}P^{48}$ tangent, with $\Sigma P\_A = I_{98}$ on the real carrier. The Haar disintegration into Dirichlet($d_1,\ldots,d_{12}$) and the predictive law $P(A \mid \rho) = (d\_A + n\_A)/(49 + N)$ are the standard conjugate results. The two-fact law $P(A,B) = d\_A(d\_B + \delta\_AB)/2450$ normalises to 1.000000000000000 and is exactly exchangeable — verified. And the 100,000-amplitude audit residual of 2.47×10^{-4} sits at **1.29 standard errors** of the analytic Dirichlet sampling error, so the audit is measuring noise rather than a discrepancy.

The reduction, and it is the substantive result. If $[U\_K_4, P\_A] = 0$ for every A then $\lVert P\_A\ U\ z\rVert = \lVert U\ P\_A\ z\rVert = \lVert P\_A\ z\rVert$ by unitarity, so every committed weight is untouched. Verified against a mixing witness: a block-diagonal U leaves the weights fixed at 2.8×10^{-17} while a component-mixing U moves them by TV = 0.45.

The K_4 question is therefore: is the Fold-to-tangent intertwiner block-diagonal with respect to the twelve gauge × S_3 components — that is, does it lie in U(1)×U(2)×⋯×U(12) rather than the full U(49)?

That is a sharp, finite, checkable condition replacing an open-ended search, and the corpus does not yet derive the intertwiner.

It also explains the earlier split, which §G.6.6 criticised on its particulars. Haar is unitarily invariant, so *any* unitary pre-fact evolution leaves the unconditioned first-fact ensemble alone — single-fact universality holds always, for structural reasons. The record-conditioned posterior carries $\prod\lVert P\_A\ z\rVert$^(2n_A), which is **not** unitarily invariant unless the commutator vanishes — so multi-fact laws can move. **The single-fact/two-fact split of §G.6.6 was structurally correct even though its two specific demonstrations were a symmetry tautology and a self-conjugate correlator.** This is the principled version of that split, and it should be cited in preference to the Ashkin–Teller numerics.

Bearing on this chapter. The commutator condition concerns K_4 and the commitment projectors, not the axle coefficient, so debt A6 is untouched directly. **But the method transfers**, and it is the sharpest suggestion yet for how to settle A6: rather than asking whether a quartic channel survives coarse-graining, ask whether its contribution is **visible through the commitment map at all** — that is, whether it moves $\|P_A z\|^2$ for any commitment projector. If it does not, it cannot change a committed record law regardless of its coefficient, and A6 closes without an RG calculation.

G.6.8 A coordinate clarification this paper needed, and one flag on the proposed test

A subsequent retraction withdraws an audit that had concluded the programme contains no weak-doublet flavour bridge, on the ground that it confused *"the selected executable does not contain this bridge"* with *"the programme contains none"*. The first remains true; the second does not. **The correction carries a clarification this paper needs internally.**

The apparent inconsistency, and its resolution. Appendix G places c_ax in the **cubic** sector, while Part II computes the branch object as a **second** variation — the scalar projected Hessian $G_{C_3} = (\kappa_C + \kappa_T + \kappa_R)I_6$ whose commutation with the cycle forces equal branch weights. Read side by side these look contradictory. They are not: they are **the same object in two coordinate systems**, related by the descent

$E_1 \otimes E_1 \xrightarrow{Q} E_2 \xrightarrow{\mathbb{I}\_int} \mathfrak{B}\_even.$

In primitive E_1 coordinates the interaction is quadratic in the source and linear in the branch, so it appears as a mixed **third** derivative $D^3S[u, u, \mathfrak{B}]$. In descended branch coordinates the same term is a **second** variation $D^2S\_cl[\mathbb{I}\_intQ(u), \mathfrak{B}]$. The chain-rule bridge between the two is what must be supplied:

$\boldsymbol{D^3S\_primitive[u, u, \mathfrak{B}\_even] \leftrightarrow D^2S\_cl[\mathbb{I}\_int\ Q(u), \mathfrak{B}\_even].}$

Part II calculates the right-hand side in the minimal cell. Appendix G counts the left-hand side on the wheel. **Neither is wrong and neither supersedes the other**, and this paper should not be read as asserting two incompatible locations for the same coupling.

A concrete consequence for A6. In primitive coordinates the quartic question is a **fifth**-derivative question, $D^5S[u,u,u,u,\mathfrak{B}]$ — quartic in the source, linear in the branch. That is consistent with $\dim \mathrm{Hom}(\mathrm{Sym}^4E_1, E_2) = 2$ and with the retraction's own reference to "the two quartic channels", and it gives A6 a concrete order to be evaluated at rather than a vague "higher-order" description.

One flag on the proposed full-carrier test. Among its conditions the retraction lists

$P\_vis\text{^}curv\ \mathbb{I}\_int\ \hat{Q}(u) = e\text{^}(i\theta\_phys)\ (1/\sqrt{3})\ E\_r,\ \hat{Q}(u) = Q(u)/\|Q(u)\|,\ r \in \{12,23,31\}.$

That presupposes $\rho\_mult^{physical} = 1$. §20 separates four distinct normalisations, of which the representation-geometric $1/\sqrt{3}$ is exact while the multiplicity survival $\rho\_mult^{physical}$ is **not identified**, and the physical coefficient is $\zeta = \rho\_mult^{physical} \cdot (b/\sqrt{3})$. At $\rho\_mult^{physical} = 0.85$ the magnitude is 0.4907b rather than 0.5774b — a 15% difference in the quantity the whole empirical comparison rests on. Written as stated the test builds in exactly the assumption F5 forbids. It should read

$$P\_vis^{curv}\ \mathbb{I}\_int\ \hat{Q}(u) = \rho\_mult^{physical}\ e^{i\theta\_phys}\ (1/\sqrt{3})\ E\_r,\ r \text{ returned by the action}$$

with $\rho\_mult^{physical}$ **returned, not set** — which the retraction's own closing list already demands when it records $\rho\_mult^{physical}$ among the open full-carrier objects.

A recording note. The retraction states that the live site could not be reached from its execution environment, so any website-only revision may move the boundary again. That limitation belongs with the result rather than in a footnote.

G.6.9 AXD-2 executed: the intertwiner is unique, and the action is blind to it

The full-carrier calculation has been run. It returns one genuine advance and one decisive negative, and the negative resolves more than the debt it was aimed at.

The advance, confirmed. The occupied representation ledger and the twelve gauge generators force a unique local intertwiner on the quark doublet, $\mathbb{I}\_Q(q) = \Sigma\ q\_ij\ |Q\_i\rangle\langle Q\_j| \otimes I\_(colour\times weak)/\sqrt{6}$. The uniqueness is Schur — Q_L is irreducible as (3,2), so dim Hom(Q,Q) = 1 and zero for inequivalent pairs — and the $\sqrt{6}$ is the normalised identity on the 3×2 gauge factor. Verified: the quoted plane matrix is a rotation by **exactly 120.000000°** with unit determinant, and the character $-0.5 + 0.866025403784i$ reproduces $e^{2\pi i/3}$ to twelve figures. **$\mathbb{I}\_int$ is therefore no longer wholly unspecified in the Q sector**, which retires part of debt A2.

The negative, and it is stronger than "not yet returned". A finite Q-family rotation commutes with all twelve gauge generators and with **all seven terminal projectors at residual exactly zero**. By the lemma of §G.6.7 — $[U, P\_A] = 0$ implies $\|P\_AUz\| = \|P\_Az\|$ — every committed weight is untouched, so the terminal law is **exactly family-rotation invariant**. The 42 post-gauge flat directions (six branch quadratures at each of seven endpoints, verified arithmetic) are the same fact in coordinates, and at 98.6% outside the gauge span they are physical flat directions rather than gauge zeros.

The terminal current-and-record action tested for this correction is blind to rotations in family space, not merely to one channel within it. By R6 this is an extension-level negative and may not be promoted to a programme-level one: RRHF already returns a generation-resolved charged tensor and CKM frame.

This moots A6 for the present action, and that is a cleaner outcome than the debt anticipated. The cheap route of §G.6.7 asked whether a quartic channel moves $\|P\_A z\|^2$ for any commitment projector. The answer is that **no family-sector object does — including the quadratic axle the action is supposed to install.** So "does the quartic survive?" is not open for this action; it is undefined, because neither channel is visible. The question becomes live only once a family-sensitive coupling exists at all. **A6 is therefore moot for the presently tested action, and reopens only once the generation-sensitive coupling is installed** — see the controlling statement in §32.

G.6.9.1 A cross-execution discrepancy: 49 against 50

The Stage-I internal ray's class probabilities are (Q, u, d, L, e, N, H) = (0.36, 0.18, 0.18, 0.12, 0.06, 0.06, 0.04). Verified: these sum to 1 and are **exactly the Standard Model field multiplicities over 50** — (18, 9, 9, 6, 3, 3, 2)/50, with Q = 3 colours × 2 weak × 3 generations and so on. That is a clean internal check.

But RUDM-1's twelve gauge × S_3 components sum to 49, not 50. Grouped by field they give (Q, u, d, L, e, N, H) = (18, 9, 9, 6, 3, **2**, 2). Under S_3 every generation-indexed field splits as trivial(1) ⊕ standard(2), and every other field in that list carries **both** a _triv and a _std entry — e_R, L_L, d_R, u_R and Q_L all do. **N_R carries only N_R_std = 2, with no N_R_triv.**

So the discrepancy is exactly one dimension and its location is identified: either the neutral singlet's trivial component is deliberately removed for a stated reason, in which case AXD-2's weight of 3/50 is wrong, or it was omitted, in which case RUDM-1's $\mathbb{C}P^{48}$ and its Dirichlet(d_A) disintegration are built on 49 where they should be 50. **The two executions cannot both be right, and the Dirichlet predictive law P(A | ρ) = (d_A + n_A)/(49 + N) depends on the answer.** Recorded as debt A2b.

G.6.10 AXD-3: the completion jet is component-preserving, with a provisional P_{23} consequence

This is the most consequential result in the series for this chapter, and it goes against the attractive branch. The generation-sensitive completion jet called for in §G.6.9 has been derived.

What is derived, confirmed. Common Q_L ancestry forces Λ_Q(X) = X ⊕ X ⊕ 0 on $\mathbb{C}^3$_u ⊕ $\mathbb{C}^3$_d ⊕ $\mathbb{C}^3$_e, since T_u and T_d share the left carrier and T_e does not. Tensoring with the one-dimensional completion leg c_H = ½(1,1,1,1) — unit norm — gives $\mathcal{K}$_n(X) = c_H ⊗ [Λ_Q(X)T_ch$^{(n)}$], unique **up to one real scalar**. The B_{23} Gram is exactly isotropic at 0.4726386533535782·I_2: rank two, equal quadrature strength, zero overlap. **The completion tensor therefore supplies no preferred phase of its own.**

What it resolves. The §18.2 fork asked whether the branch coefficient inherits the source mode's rim-to-rim advance (reading (i), giving 150° branch-independently) or the phase of the selected descended component (reading (ii), giving one of three values). **The derived jet is component-preserving, which is reading (ii). AXD-3 establishes that the completion jet is component-preserving. It does not establish which component is physically read.** Under the provisional and unestablished choice P_vis^curv = P_{23} the selected component would carry 30° or 270°. **HCFC-1 shows that this projector is not derived and that the local metric selects no phase.** The quadratic character descent and source-phase candidate are derived; the physical carrier, comparison holonomy and phase value are not returned — **not 150°**.

And the empirical consequence, which the report does not compute.

branch	J	|V_ub|	|V_td|	χ^2_5
150° *(the branch this paper has been executing)*	$+2.997\times10^{-5}$	0.003557	0.008690	**5.3**
30° or 270° — the provisional P_{23} branch	$+2.997\times10^{-5}$ / -4.66×10^{-6}	0.006541 / 0.003372	0.004859 / 0.005763	**539.0 / 1031.6**
no axle at all (RRHF baseline)	1.860×10^{-5}	0.004198	0.005895	290.7

At the provisional P_{23} branch the axle makes the fit worse than having no axle at all. The improvement from 291 to 4.1 that this chapter has been reporting belongs to a branch the derivation does not return.

The conclusion is chirality-independent. Which of 30° or 270° the readout selects depends on the mode chirality — an open selection (VX-CIRC) — and the two conventions swap them. Verified: with component phases (0°, 120°, 240°) the B_{23} branch reads 270°, and with (0°, 240°, 120°) it reads 30°. **Neither is 150°, and both are worse than no axle**, so the finding does not rest on the unresolved chirality bit.

G.6.10.1 What this does to F4

F4 forbids selecting any value because it improves a measured quantity. **Until now the 150° branch was defensible as a declared execution under an undeclared carrier.** That defence is gone: the carrier is no longer available as an undeclared convention: the jet's component-preserving character is derived, and no local structure returns 150°. **Nor does any local structure return 30° or 270°** (§G.6.11) — so 150° is not merely unselected, it is one of three branches none of which the framework presently returns. **Continuing to feature the 150° number as the paper's headline would be selection for effect in the exact sense F4 prohibits.**

The correct handling is the one the paper already applies elsewhere: the 150° execution is retained as a **quarantined diagnostic**, clearly labelled as a branch the derivation does not return, and the provisional P_{23} branch and its χ^2 are reported alongside it wherever the improvement is quoted.

G.6.10.2 Two things the result does not settle

The coefficient is still open, and one obstacle is now quantified: the unsplit common jet has norm exactly **twice** that of the half-split odd jet — verified — because $\|X \oplus X\| = \sqrt{2}\|X\|$ while $\|-X/2 \oplus X/2\| = \|X\|/\sqrt{2}$. So c_ax = b does **not** follow from weak-complement norm preservation without separately declaring how the pre-split profile norm relates to the RRHF block norm.

A branch-independent 150° is not excluded, but nothing supplies it. It would require an additional source-phase functional attaching the rim-to-rim advance to every visible branch independently of the component read — reading (i) made explicit. No such functional is present in the frozen completion, current or readout packages. **That, not the coefficient, is now the single object on which the attractive branch depends.**

G.6.11 HCFC-1: a genuine no-go, and one theorem that reinstates a withdrawn argument

The new no-go is real and important. The frozen completion attachment on the returned B_{23} plane has Gram exactly γI_2 — at level 4, $\gamma = 0.4726386533535782$ — so for $K\_\theta = \cos\theta \cdot K\_R + \sin\theta \cdot K\_I$ the norm $\|K\_\theta\|^2 = \gamma$ for every θ. Verified: spread 1.7×10^{-16} over a 360-point sweep. **The local completion attachment can transport a phase but cannot select one**, and no value is preferred — not 150°, and equally not 30° or 270°.

But Theorem HCFC-1A reinstates an argument this paper has already withdrawn. It sets $\mathfrak{B}\_inh = \mathfrak{B}_{12} \oplus \mathfrak{B}_{31}$ and takes the orthogonal complement to be $\mathfrak{B}_{23}$, concluding P_vis^curv = P_{23}. The inherited role-odd generator has

(1,2) entry a = 0.225 · **(2,3) entry b = 0.0405** · (1,3) entry c = 0.00243,

so **$\mathfrak{B}_{23}$ is occupied**, $\mathfrak{B}$_inh is the whole carrier, and the orthogonal complement is zero rather than $\mathfrak{B}_{23}$. §18 records the identical move as withdrawn — *"set $\mathfrak{B}\_inh = \mathfrak{B}_{12} \oplus \mathfrak{B}_{31}$ and took the complement. False, since $b \neq 0$ in the (2,3) slot; and by the occupancy witness it was forbidden a priori"* — and **falsifier F2 fires on it directly.** The theorem's premise is false and its conclusion is not established.

G.6.11.1 What this does to §G.6.10

§G.6.10 recorded the carrier as resolved, returning 30°/270°. That conclusion combined the component-preserving jet with the **assumed** readout P_vis^curv = P_{23}. HCFC-1 attempts to establish that projector and fails on the occupancy witness.

So 30° and 270° were conditional on P_{23} all along, and P_{23} is not established. Combined with HCFC-1B, the honest position is stronger than either report states: **the local structure determines no phase at all.** 150° is not derived; 30° and 270° are conditional on an unestablished projector resting on a withdrawn premise.

This does **not** rehabilitate 150°. It moves the phase from *resolved against the attractive branch* to *undetermined by anything local*, with the holonomy proposal of §6 the only candidate mechanism and its loop unrealised.

G.6.11.2 What survives, and it is not nothing

Independent of the readout choice, and therefore unaffected:

Result	Status
ρ_component = 1	exact — orthogonal projection does not rescale what it keeps
ρ_mult^(AXD-3 image) = 1, with ρ_mult^physical open	exact — rank-2 image with equal singular values 0.6874872023
HCFC-1B local phase blindness	**exact, and a genuine new no-go**
FMC-1 functional shape unique up to positive scale	correct, though it is the observation that a real linear functional on a two-plane has one Riesz vector
maintenance conjugacy: μ_F > 0 ∧ H_F ≻ 0 ⟹ B_F > 0	correct — H_F ≻ 0 gives $H\_F^{-1} \succ 0$, so $u^T H\_F^{-1} u > 0$

The value ρ_read = $1/\sqrt{3}$ has correct **arithmetic** $|Q_{23}|/\|Q(u)\| = (1/3)/(1/\sqrt{3})$, verified to sixteen figures — but it is the ratio *for* P_{23}, so it is physical only if that projector is.

The report's own §10 warning should be kept prominent: the jet Gram value 0.4726… is **not** c_ax, and the pre-split versus per-block definition of b must be reconciled before c_ax = b can even be stated.

The two irreducible returns named in §12 are the right ones — realise and evaluate the loop L_ax, and calculate μ_F and H_F in physical units — **provided the readout is derived rather than argued from novelty.** That is now a third.

G.7 What this changes

It converts the obstruction from qualitative to quantitative. The programme ledger's non-uniqueness and Stage I's rank-one degeneracy are two views of a family whose dimension can now be counted, whose basis can be written down, and in which the one coordinate the flavour sector needs — c_ax — has an address.

Three consequences follow for planning.

1. **Roughly 22 independent conditions are required on the minimal carrier alone.** The one-Bit condition is one such condition, which indicates the scale of what remains.
2. **The Laplacian's two vanishing long-range coordinates are the cheapest place to look.** They are assumptions of the reference construction; deriving or refuting them is a well-posed question that does not require the full action.
3. **The axle's existence should not be the target of the next calculation** — it is generic in the unrestricted family, though the tested action is not a generic point of it. The next target is a **nonzero physically installed c_ax**, together with its complete physical lift, readout, retention and phase carrier. Whether c_ax = b is the subsequent codimension-one **Stage I-B** comparison.