TFPT — The Standard Model from the Compiler

The φ_0-ladder, the flavor block from parabolic transport, and the worked closures (θ_{12}, quark c, the gap, Starobinsky M)

June 8, 2026

What this document is about

The Standard Model as *one* φ_0-ladder mass formula (all masses, CKM, PMNS, neutrinos), the flavor matrix from parabolic transport on $\mathbb{P}^1 \setminus \mu_4$, the five worked closures (θ_{12}, quark c, the explicit mass gap, Starobinsky M, the H2 splitting), and **gravity/quantum gravity as the seam response** (the Seam Response Grammar $\xi u = c_3$, the heat-kernel grounding of $R + R^2$, and the metric-coupled gap bound that *gap-decouples* the physical QG sector, leaving only the ambient projective limit as a mathematical-completeness item).

The TFPT document set — what is treated where

Plain language: TFPT is a small discrete compiler. Two inputs — the seam constant $c_3 = \frac{1}{8\pi}$ and the carrier rank $g_{\text{car}} = 5$ — build $D_5 \oplus A_3 + \mu_4 \Rightarrow E_8$ and read off the Standard Model, the constants and the scale grammar. The development is **six short documents**, best read in order:

1. `introduction` — reading guide, compiler closure, paper-by-paper comparison, predictions, the dependency DAG and proof ledger.
2. `tfpt_1_architecture_e8` — the two axioms, the derivation map, the EM fixed point α^{-1}, the $D_5 \oplus A_3 + \mu_4 \Rightarrow E_8$ construction.
3. `tfpt_2_standard_model` — the SM in one φ_0-ladder formula, flavor from parabolic transport, the worked closures, and gravity/QG as the seam response.
4. `tfpt_3_e8_audit_bootstrap` — the seven E_8 slices as an audit raster, the cascade bridge, the Möbius bootstrap.
5. `tfpt_4_frontier` — honest status of η_B, m_p/m_e, Koide, dark matter, full QG.
6. `tfpt_horizon_readouts` — one seam constant as the universal horizon thermal code.

You are reading document #3.

Contents

Part I

The Standard Model in one master formula

In one sentence

The fermion spectrum — masses, Yukawa structure, CKM, the PMNS skeleton and the neutrinos — follows from *one* master formula with *three* ingredients: the seed φ_0^{ret}, the carrier base $\lambda_Y = \sqrt{\varphi_0^{\mathrm{ret}}(1-\varphi_0^{\mathrm{ret}})}$ and the *residue matrix of the compiler.* No free Yukawa *exponents* remain; the charged-lepton amplitudes are closed, the quark mass *ratios* are closed too (integer Plücker readouts, v49/v71), and only the *absolute* quark amplitude scale is a finite (U_{wall}) anchor.

Status markers (refined): [I] exact identity, **[L]** Lie/lattice theorem, **[F]** formalised, [N] numerical fixed point, **[P]** physical/conditional (named hypotheses), **[A]** axiom/open.

1 The principle: one master formula instead of many Yukawa numbers

In Papers 1–7 each fermion mass arises through a Yukawa coupling $y_{f,j}$ built as a transport word-length $L_{f,j}$ on a hexagon (C_6) times a holonomy amplitude. In the compiler picture this is *one* line:

Mass master formula (closed branch)

$$\boxed{\hat{m}_{f,j} = \frac{v_{\mathrm{geo}}}{\sqrt{2}}\, y_{f,j}, \qquad y_{f,j} = \lambda_Y^{L_{f,j}}\, \Lambda_{f,j}, \qquad \lambda_Y = \sqrt{\varphi_0^{\mathrm{ret}}\,(1-\varphi_0^{\mathrm{ret}})}}$$

with *one* seed $\varphi_0^{\mathrm{ret}} = \dfrac{1}{6\pi} + \dfrac{3}{256\pi^4}$, the integer word-lengths $L_{f,j} = r_{f,j} + 6\, n_{f,j}$ (residue + winding) and $O(1)$ holonomy amplitudes $\Lambda_{f,j}$.

The point of the simplification: $L_{f,j}$ **is not a fit but the output of the compiler.** The residues $r_{f,j}$ form the matrix whose invariants are pure compiler numbers (Section 3); the windings $n_{f,j}$ are the anchors $\{1, 1, 2\}$. What used to be a separate “flavor transport” theorem is now the readout of a fixed matrix.

2 The seed φ_0^{ret} as a decoder

A single seed drives four independent low-energy observables. This is the strongest compression of the flavor input:

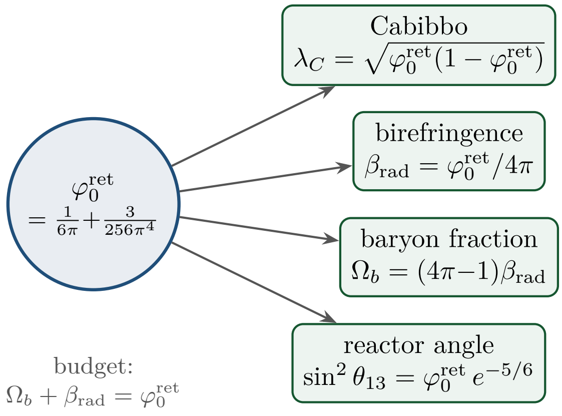

Numerically: $\varphi_0^{\text{ret}} = 0.0531720$, $\lambda_C = 0.224376$ (PDG 0.2245), $\sin^2\theta_{13} = 0.02311$ (PDG ≈ 0.0222). The exponent $\gamma = 5/6$ is the carrier trace $\operatorname{tr}_E Y^2$; its complement $1-\gamma = 1/6$ reappears in the neutrino ratio.

3 The flavor residue matrix is the compiler signature

The core of the simplification: the word-length matrix $L = R + 6W$ splits into the residue matrix R (the hexagon packets of the three sectors) and the winding matrix W (one $+6$ per first generation = anchor). R carries all the discrete flavor information — and its invariants are *exclusively* compiler numbers:

$$R = \begin{pmatrix} 1 & 3 & 0 \\ 1 & 5 & 2 \\ 2 & 5 & 3 \end{pmatrix}$$

rows = sectors u, d, e

$\operatorname{tr} R = 9 = N_{\text{fam}}^2$

minors $2, 3, 5 \Rightarrow 2{\cdot}3{\cdot}5 = 30 = h(E_8)$

$\det R = 8 = h(D_5)$

Spectral signature of the flavor matrix (verified)

The characteristic polynomial of the residue matrix is built entirely from compiler constants:

$$\chi_R(t) = t^3 - \underbrace{9}_{N_{\text{fam}}^2} t^2 + \underbrace{10}_{\mathcal{A}_L=\binom{5}{2}} t - \underbrace{8}_{h(D_5)},$$

where the linear coefficient $10 = 2+3+5$ is the sum of the three principal $2{\times}2$ minors and their *product* $2\cdot 3\cdot 5 = 30$ is the Coxeter number $h(E_8)$. The flavor matrix is therefore not an independent object — it is a *reading of the same compiler* $(N_{\text{fam}},\ \mathcal{A}_L,\ h(D_5),\ h(E_8))$. [I] [`verification/v4_flavor_matrix.py`]

Concretely:

$$L = \begin{pmatrix} 7 & 3 & 0 \\ 7 & 5 & 2 \\ 8 & 5 & 3 \end{pmatrix} = \underbrace{\begin{pmatrix} 1 & 3 & 0 \\ 1 & 5 & 2 \\ 2 & 5 & 3 \end{pmatrix}}_{R\ (\det 8,\ \text{minors } 2,3,5)} + 6 \underbrace{\begin{pmatrix} 1 & 0 & 0 \\ 1 & 0 & 0 \\ 1 & 0 & 0 \end{pmatrix}}_{W\ (\text{anchor } \{1,1,2\})}, \qquad \sum L = 40 = \Omega_{\text{adm}} \cdot \tfrac{5}{6}.$$

The sum $\sum L = 40$ is the transport budget number; with the Higgs contribution $N_\Phi = 1$ it gives $40 + 1 = 41 = 10\, b_1$ — exactly the coefficient the EM fixed point $\alpha_\star^{-1}$ reads.

The winding-deformation theorem: $R \to L$ in one parameter [I]

Treat the winding as a one-parameter rank-one deformation $R_s = R + s\,\mathbf{1}e_1^\top$ ($R_0 = R$, $R_6 = L$). Then the *entire* flavor linear algebra is one line,

$$\boxed{\chi_{R_s}(t) = t^3 - (9+s)\,t^2 + (10+5s)\,t - (8+2s)},$$

with $\det R_s = 8+2s$, $\operatorname{tr} R_s = 9+s$, $\mathrm{PrinMin}_2(R_s) = (5,\,3(s{+}1),\,2(s{+}1))$, and $R_s^{-1}\mathbf{1} = \frac{1}{4+s}(1,1,-1)^\top$. The physical winding value is fixed, not chosen:

$$\boxed{s = |R^+(A_3)| = 6 \text{ is \textit{triply locked}:}\quad \underbrace{6+4s=30}_{R_s^\top a\to h(E_8)},\ \underbrace{8+2s=20}_{\det\to 2\mathcal{A}_L},\ \underbrace{9+s=15}_{\mathrm{tr}\to\dim A_3}\ \Rightarrow\ s=6.}$$

At $s = 6$ this gives $\chi_L = t^3 - 15t^2 + 40t - 20$ with $15 = \dim A_3$, $40 = |R(D_5)|$, $20 = |R(A_4)| = 2\mathcal{A}_L$, and $\mathrm{PrinMin}_2(L) = (5, 21, 14) = (g_{\mathrm{car}}, 3{\cdot}7, 2{\cdot}7)$ — the lift multiplies the sheet and family minors by the successor $s{+}1 = 7 = |R^+(A_3)|{+}1$ while the carrier minor 5 stays fixed. So the "+6 winding" is not "add six because hexagon": it is the unique value selected three independent ways, and the appearance of 7 in the scalaron is $s{+}1$.

Winding reality, the determinant lift, and the invariant dual anchor (machine-verified)

The rank-one winding $R \to L = R + 6\,\mathbf{1}e_1^\top$ is far from a cosmetic offset:

- **Winding makes the flavour system spectrally real.** R has a *complex* eigenpair ($\operatorname{disc}\chi_R = -2^2 \cdot 1999 < 0$, spectrum 7.857, $0.572 \pm 0.832i$), while L has three positive reals, with a fully compiler-factorised discriminant (two equivalent readings)

$$\operatorname{disc}\chi_L = 39200 = 2^5 \cdot g_{\mathrm{car}}^2 \cdot 7^2 = 2\,(\mathcal{A}_L \cdot \dim G_2)^2 = 2\,(10\cdot 14)^2.$$

- **The anchor Coxeter descent.** The anchor projects onto a clean three-scale ladder,

$$R^\top a = (6, 18, 8), \qquad L^\top a = (30, 18, 8) = (h(E_8),\ N_{\mathrm{fam}}|R^+(A_3)|,\ h(D_5)),$$

  so the winding lifts $|R^+(A_3)| = 6 \to h(E_8) = 30$ (a $+24 = |W(A_3)|$ shift), with descent gaps $30 - 18 = 12 = |R(A_3)| = \dim\mathfrak{g}_{\mathrm{SM}}$ and $18 - 8 = 10 = \mathcal{A}_L$. Column-wise the lift is $\mathbf{1}^\top R = (4, 13, 5) \to \mathbf{1}^\top L = (22, 13, 5)$, adding 18 only in the first generation — which is thus literally the store of the family-hexagon winding.
- **The determinant lift is the A_3 winding over μ_4.** Since $e_1^\top R^{-1}\mathbf{1} = \frac{1}{4} = \frac{1}{|\mu_4|}$,

$$\det L = \det R\Big(1 + \tfrac{|R^+(A_3)|}{|\mu_4|}\Big) = 8(1 + \tfrac{6}{4}) = 20 = 2\mathcal{A}_L,$$

  so the D_5 selector 8 is lifted to the decuple-doublet 20 by the A_3 winding. The cokernel moves $\mathbb{Z}_8 \to \mathbb{Z}_{20}$, i.e. the winding injects the carrier factor 5 into the torsion.
- **The dual anchor is winding-invariant.** The first column is the anchor, $Re_1 = a = (1, 1, 2)$, and the dual covector is preserved by the lift,

$$a^\top R^{-1} = a^\top L^{-1} = (-\tfrac{1}{2}, -\tfrac{1}{2}, 1), \qquad (-\tfrac{1}{2}, -\tfrac{1}{2}, 1)\cdot\mathbf{1} = 0,$$

  so winding changes the absolute hierarchy but *not* the anchor-dual sheet balance. The uniform response keeps the sheet value fixed: $\mathbf{1}^\top\mathrm{adj}(R)\mathbf{1} = \mathbf{1}^\top\mathrm{adj}(L)\mathbf{1} = 2 = |\mathbb{Z}_2|$, while the inverse response denominator runs $|\mu_4| = 4 \to \mathcal{A}_L = 10$. [I]

Audit-level only: flavor linear-algebra fingerprints (not load-bearing)

The following are *audit fingerprints* (exact linear algebra of R, L), kept here, not in any derivation, under the no-free-pattern rule. [I]

- **Cofactor seam normal.** Generations $2, 3$ share columns, $c_2{=}(3,5,5)^\top$, $c_3{=}(0,2,3)^\top$; their cross product $n = c_2 \times c_3 = (5,-9,6) = (g_{\mathrm{car}}, -N_{\mathrm{fam}}^2, |R^+(A_3)|)$ gives the determinant geometrically: $n \cdot a = 8 = \det R$, $n \cdot (a{+}6\mathbf{1}) = 20 = \det L$, $n \cdot \mathbf{1} = 2 = |\mathbb{Z}_2|$.
- **Adjugate selector ($\Rightarrow$ why only gen. 1 winds).** n is also the first row of both adjugates, $e_1^\top \mathrm{adj}(R) = e_1^\top \mathrm{adj}(L) = (5,-9,6)$, so $(5,-9,6)R = (8,0,0)$ and $(5,-9,6)L = (20,0,0)$ — it annihilates the unchanged c_2, c_3 plane.
- **Length polynomial as an action polynomial.** $\chi_L(t) = t^3 - \dim(A_3)\, t^2 + |R(D_5)|\, t - |R(A_4)| = t^3 - 15t^2 + 40t - 20$.
- **Principal minors.** $\mathrm{PrinMin}_2(L) = (5, 21, 14) = (g_{\mathrm{car}}, 3{\cdot}7, 2{\cdot}7)$ with $\sum = 40 = |R(D_5)|$, $\prod = 1470 = 7^2 h(E_8)$.
- **Commutator is a pure winding rotation.** $[R, L]$ has rank 2, trace 0, and char. poly $\lambda(\lambda^2 + 432)$ — a single 2D rotation in the non-anchor plane with a null direction.

4 Charged leptons: completely closed in φ_0^{ret}

For the leptons the hexagonal resolvent solves the amplitudes Λ_ℓ in closed form — only pure φ_0^{ret} powers with rational carrier prefactors remain:

Absolute lepton source masses

$$(\hat{m}_e, \hat{m}_\mu, \hat{m}_\tau) = \frac{v_{\mathrm{geo}}\,\pi}{\sqrt{2}} \left(\tfrac{16}{7}\,(\varphi_0^{\mathrm{ret}})^5,\ \tfrac{4}{3}\,(\varphi_0^{\mathrm{ret}})^3,\ \tfrac{7}{6}\,(\varphi_0^{\mathrm{ret}})^2 \right),$$

and hence the *parameter-free* ratios

$$\frac{\hat{m}_\mu}{\hat{m}_\tau} = \frac{8}{7}\,\varphi_0^{\mathrm{ret}}, \qquad \frac{\hat{m}_e}{\hat{m}_\mu} = \frac{12}{7}\,(\varphi_0^{\mathrm{ret}})^2, \qquad \frac{\hat{m}_e\,\hat{m}_\tau}{\hat{m}_\mu^2} = \frac{3}{2}\,\varphi_0^{\mathrm{ret}}.$$

Where the lepton rationals come from: $\delta = \frac{1}{2}$ resolvent + product rule [I]

The coefficients $(\frac{16}{7}, \frac{4}{3}, \frac{7}{6})$ are not fitted — they are *derived*, and the key is that they are *rational*, so they cannot come from the irrational pole δ_{ph}. They come from the *distinguished* transport value $\delta = \frac{1}{2}$ (rational), evaluated at the lepton leg $y = 1$ $(= |Y(e^c)|)$. **The value $\delta = \frac{1}{2}$ itself is a glue-norm readout:** the two E_8-glue norms are $q(D_5) = g_{\mathrm{car}}/|\mu_4| = \frac{5}{4}$ and $q(A_3) = N_{\mathrm{fam}}/|\mu_4| = \frac{3}{4}$, so $\delta = q(D_5) - q(A_3) = (g_{\mathrm{car}} - N_{\mathrm{fam}})/|\mu_4| = |\mathbb{Z}_2|/|\mu_4| = \frac{1}{2}$ — the carrier-minus-family count over the glue index (the same pair whose *sum* is the E_8 root norm 2, Paper 1), matching the harmonic-frame holonomy diagonal modulus $\frac{1}{2}$. So δ is fixed, not chosen ([`verification/v51_boundary_half_step.py`]). With that value,

$$\mathrm{amplitude}(r) = \left| \frac{1}{1 - \frac{1}{2}\zeta^r} \right|^2 = \frac{1}{\frac{5}{4} - \cos(r\pi/3)}, \qquad \zeta = e^{i\pi/3},\ r = L \bmod 6.$$

Non-anchor leptons (dressed by the winding factor $|\mu_4| = 4$, $w = \lfloor L/6 \rfloor$):

$$c = |\mu_4|^w \frac{1}{\frac{5}{4} - \cos(r\pi/3)} : \quad c_e = 4^1 \cdot \tfrac{4}{7} = \tfrac{16}{7}, \qquad c_\mu = 4^0 \cdot \tfrac{4}{3} = \tfrac{4}{3} \quad \text{(exact).}$$

Anchor (heaviest lepton τ) fixed by the *sector product rule*

$$\boxed{c_e\,c_\mu\,c_\tau = \frac{2^{g_{\rm car}}}{N_{\rm fam}^2} = \frac{32}{9}} \quad\Longrightarrow\quad c_\tau = \frac{32/9}{(16/7)(4/3)} = \frac{7}{6} \quad \text{(exact)}.$$

Every ingredient is a compiler datum ($\delta{=}\frac{1}{2}$, $|\mu_4|$, hexagon angles, $g_{\rm car}$, $N_{\rm fam}$). The product constant reads $2^{g_{\rm car}} = 32 = \dim(S^+ \oplus S^-)$ (the full carrier Clifford space) over $N_{\rm fam}^2$, and the lepton word-lengths obey $\sum_j L_j = 16 = \dim S^+$, $\sum_j (L_j - K_j) = 6 = |R^+(A_3)|$. *Edge–centre ring:* the two boundary generations close on the middle one, $c_e\,c_\tau = \frac{16}{7}\cdot\frac{7}{6} = \frac{8}{3} = |\mathbb{Z}_2|\,c_\mu$, so the lepton amplitudes form a closed algebraic ring with sheet-parity $|\mathbb{Z}_2| = 2$. [`verification/v20_lepton_c_derivation.py`] [`verification/v21_solar_product_quark.py`] [`verification/v37_plucker_anchor.py`]

Scope (honest). This law is lepton-specific. Applied to the down sector it fails provably: $r_d{=}1, r_s{=}5$ share the same amplitude, so the law forces $c_d/c_s{=}|\mu_4|{=}4$, whereas the source ratio gives 0.94. Quarks (colour, different cusp legs) are structurally different and their exact c's stay open (cf. the $U_f^\star$ reduction above).

Ratio	Formula	Prediction	PDG
$\hat m_\mu/\hat m_\tau$	$\frac{8}{7}\varphi_0^{\rm ret}$	0.06077	0.05946
$\hat m_e/\hat m_\mu$	$\frac{12}{7}(\varphi_0^{\rm ret})^2$	0.004847	0.004836
$\hat m_e\hat m_\tau/\hat m_\mu^2$	$\frac{3}{2}\varphi_0^{\rm ret}$	0.07976	0.08133

The factor $B = \frac{3}{2}$ in the Koide-type product is the rank ratio $\operatorname{rank} E_3/\operatorname{rank} E_2$ of the carrier — again a compiler number, not a fit. The e/μ relation hits the PDG to 0.2%.

5 Quarks: hierarchy from the same word-lengths

The quark sectors use the same master formula with $L_u = (7,3,0)$, $L_d = (7,5,2)$. Where two generations share the same word-length the mass ratio is pure transport amplitude; where they differ by 2 a factor λ_Y^{-2} appears:

Ratio	Origin	Prediction	PDG order
$\hat m_u/\hat m_d$	Λ_u/Λ_d (same $L{=}7$)	0.470	~ 0.5
$\hat m_c/\hat m_s$	$\lambda_Y^{-2}\,\Lambda_c/\Lambda_s$	13.6	$\sim$ 11–14
$\hat m_t/\hat m_b$	$\lambda_Y^{-2}\,\Lambda_t/\Lambda_b$	40.8	~ 41

The residue classes mod 6 link quarks and leptons: $L_c = L_\tau = 3$, $L_s = L_\mu = 5$ — the same hexagon position, different sectors. This is the geometric root of quark–lepton complementarity.

The $\varphi_0^{\rm ret}$-ladder: *all* nine masses in one closed form

The quark Λ are *not* free numbers: they are evaluations of the *same* finite hexagonal resolvent D_y^{-1} at the *same* pole $\delta_{\rm ph}$ as the leptons. For every fermion the *universal* closed form holds

$$\boxed{\lambda_Y^{L_{f,j}}\,\Lambda_{f,j} = \pi\, c_{f,j}\,(\varphi_0^{\rm ret})^{k_{f,j}} \quad\Longrightarrow\quad \hat m_{f,j} = \frac{v_{\rm geo}\,\pi}{\sqrt{2}}\, c_{f,j}\,(\varphi_0^{\rm ret})^{k_{f,j}}}$$

with carrier rationals $c_{f,j}$ and integer $\varphi_0^{\rm ret}$-order $k_{f,j}$. The leptons give $c = (\frac{16}{7}, \frac{4}{3}, \frac{7}{6})$ *exactly*; the quark c are fixed by the same resolvent (the published Λ table is only three-digit). The

$\varphi_0^{\rm ret}$-orders form the **mass ladder**

$$\underbrace{t}_{k=0} > \underbrace{b,\, c,\, \tau}_{k=2} > \underbrace{s,\, \mu}_{k=3} > \underbrace{u,\, d}_{k=4} > \underbrace{e}_{k=5},$$

each rung a factor $\varphi_0^{\rm ret} \approx 1/19$. The whole fermion hierarchy is a $\varphi_0^{\rm ret}$ power ladder with carrier-rational rungs — quarks *and* leptons in *one* family. [I] (leptons exact) / [P] (quark rationals from the resolvent)

Quark *ratio* closure: the invariant readouts of the character sector [I]/[N]

The right invariants are not the six absolute c_q (gauge-dependent through the external legs and the sector normalisation) but the three intra-sector *ratios*, which close on pure compiler building blocks. With $\Delta_Q := |R(A_3)|{+}1 = 13$, $\Omega_{\rm adm} = 48$, $N_{\rm fam} = 3$ and the anchor power sum $p_5(a) = 1^5{+}1^5{+}2^5 = 34$,

$$\boxed{\frac{c_u}{c_d} = \frac{g_{\rm car}\cdot 11}{N_{\rm fam}^2\Delta_Q} = \frac{55}{117}, \quad \frac{c_c}{c_s} = \frac{p_5(a)}{\Omega_{\rm adm}-1} = \frac{34}{47}, \quad \frac{c_t}{c_b} = \frac{N_{\rm fam}}{2\Delta_Q} = \frac{3}{26}}$$

With the $\varphi_0^{\rm ret}$-orders $(k_u, k_d, k_c, k_s, k_t, k_b) = (4, 4, 2, 3, 0, 2)$ the mass ratios are

$$\frac{\hat m_u}{\hat m_d} = \frac{55}{117} = 0.4700855, \qquad \frac{\hat m_c}{\hat m_s} = \frac{34}{47}(\varphi_0^{\rm ret})^{-1} = 13.6050, \qquad \frac{\hat m_t}{\hat m_b} = \frac{3}{26}(\varphi_0^{\rm ret})^{-2} = 40.812,$$

matching the closed-branch targets (0.470085, 13.607, 40.80) to $\lesssim 0.03\%$. A rational normalisation gauge compatible with these ratios is

$$(c_u, c_d, c_s, c_c, c_b, c_t) = \left(\tfrac{1}{2},\ \tfrac{117}{110},\ \tfrac{47}{41},\ \tfrac{34}{41},\ \tfrac{8}{3},\ \tfrac{4}{13}\right),$$

whose $\Lambda_q = \pi c_q(\varphi_0^{\rm ret})^k/\lambda_Y^L$ all land in the $O(1)$ band ($\Lambda_{u,d,s,c,b,t} \approx 0.44, 0.93, 0.95, 0.65, 0.47, 0.97$). The denominators carry $41 = 4^2{+}5^2 = 10b_1$. The individual c_q are a gauge; the three ratios are the invariant readouts. [`verification/v24_quark_ratio_closure.py`]

Honest note on the "11". Two of the three ratios have *clean single-reading* building blocks: $c_c/c_s = p_5(a)/(\Omega_{\rm adm}{-}1)$ (the anchor power sum $p_5{=}34$) and $c_t/c_b = N_{\rm fam}/(2\Delta_Q)$. The first ratio is the odd one out: its numerator factor 11 admits *many* TFPT readings ($\dim S^+{-}g_{\rm car}$, $g_{\rm car}{+}|R^+(A_3)|$, $h(D_5){+}N_{\rm fam}$, $2g_{\rm car}{+}1$, $\mathcal{A}_\Lambda{+}1$, $\Delta_Q{-}2$, the E_8 exponent 11, …), so it is *not uniquely forced* as a loose integer. **This worry is now retired** (Plücker box below): $11 = \|\,{\rm Pl}(K)\|_1$ is *generated* as the L^1 norm of the anchor-plane minors of K (v42), and by Readout Rigidity (v49, with the selector stratum now *derived* — $\det R = 8$ lattice and ${\rm Spec}(Q_+) = \{1, 2, 3\}$ from D_4, v69/v71) the *ratio* $c_u/c_d = \frac{55}{117}$ is *constant* on that stratum, i.e. an integer Plücker readout that needs *no* transcendental solve. The wall-selection (U) enters only the *absolute* amplitude normalisation (an anchor), not the ratio. [`verification/v26_flavor_frontier_unification.py`] Quark masses are scheme-sensitive (esp. light quarks); we state dimensionless source-scale ratios, not pole masses (cf. FLAG 2024).

Anchor-plane Plücker localisation of the "11" (and 26) [I]

The cleanest way to retire the "many readings" worry is to stop reading loose integers and read *orientation invariants* of the one canonical plane in play — the *anchor plane* spanned by the two TFPT generators $\mathbf{1} = (1, 1, 1)$ and $a = (1, 1, 2)$. For $M \in \{R, Q, K, L\}$ form the left anchor block and its 2×2 minors (Plücker coordinates)

$$B_M = \begin{pmatrix} \mathbf{1}^\top M \\ a^\top M \end{pmatrix}, \qquad \mathrm{Pl}(M) = \big(\,[12], [13], [23]\,\big).$$

Then $\mathrm{Pl}(K) = (-N_\Phi,\ |R^+(A_3)|,\ |\mu_4|) = (-1, 6, 4)$ and

$$\boxed{\|\,\mathrm{Pl}(K)\|_1 = 1 + 6 + 4 = 11}$$

is the *Plücker norm of the mass-power anchor plane* — a single coordinate-free object, not a free integer. The right block $C_M = (M\mathbf{1}\ \ Ma)$ gives $\mathrm{Pl}_R(K) = (12, 12, -2)$ with $\|\,\mathrm{Pl}_R(K)\|_1 = 26 = \mathrm{Pl}(L)_{13}$ (the c_t/c_b denominator), and $\sum \mathrm{Pl}_R(L) = 20{+}30{+}10 = 60$ is the E_8 cascade start. So the three quark ratios are the *same* numbers read as anchor-plane Plücker invariants:

$$\frac{c_u}{c_d} = \frac{g_{\mathrm{car}}\,\|\,\mathrm{Pl}(K)\|_1}{N_{\mathrm{fam}}^2\,\Delta_Q}, \qquad \frac{c_t}{c_b} = \frac{N_{\mathrm{fam}}}{\mathrm{Pl}(L)_{13}}, \qquad \frac{c_c}{c_s} = \frac{\sum \mathrm{Pl}(L) - \sum \mathrm{Pl}(Q)}{1 + \sum \mathrm{Pl}(L)}.$$

What this does and does not buy. The Plücker values are exact identities [I], and the "11" now *lives in the mass-power matrix K* via the canonical plane rather than in a free integer pool — the best-motivated single reading. Combined with Readout Rigidity (below), this *closes the ratio*: $c_u/c_d = \frac{55}{117}$ is fixed combinatorially, no $U_f^\star$ evaluation required. What stays open is the *absolute* amplitude scale (an anchor) and a mild choice of invariant in the c_c/c_s block. [`verification/v37_plucker_anchor.py`]

Why the leg is exterior, not scalar (Exterior Leg Lemma). u and d share the hexagon-radial data $(r, w) = (1, 1)$, and the C_6 resolvent reads only (r, w) — so *any scalar* leg forces $c_u/c_d \equiv 1$. A scalar sees radius, not orientation; the u/d separation is intrinsically the oriented anchor-plane *area* $\mathrm{Pl}_{\mathbf{1},a}(K)$ (the exterior square $\Lambda^2 F$ of the $SU(3)_F$ composition), whose L^1 norm *generates* the "11". [`verification/v42_exterior_leg.py`]

Readout Rigidity — the quark *ratios* are closed combinatorially [I] (residual: absolute scale [A])

The earlier "reduce c_u/c_d to the transcendental wall-selection" framing is *superseded*. The four flavor operators are integer lattice maps on $H_1(\mathbb{P}^1 \setminus \mu_4) = \mathbb{Z}^3$ with compiler-atom determinants $(\det Q, \det K, \det R, \det L) = (3, 4, 8, 20)$, product $1920 = |W(D_5)|$, and the *discrete selector stratum*

$$\mathcal{S}_{8,Q} = \{\det R = 8,\ \ \mathrm{SNF}(R) = (1, 1, 8),\ \ \mathrm{Spec}(Q_+) = \{1, 2, 3\}\}$$

is now *fully derived*: $\det R = 8 = \operatorname{rank} E_8$ and the SNF are the lattice invariant (v70, $\det Q = 3 = N_{\mathrm{fam}}$), and $\mathrm{Spec}(Q_+) = \{1, 2, 3\}$ is the D_4-equivariant result (v69). By **Readout Rigidity** (Theorem U2, [`verification/v49_readout_rigidity.py`]) the anchor-plane readout is *constant* on $\mathcal{S}_{8,Q}$, so the quark *ratios*

$$\frac{c_u}{c_d} = \frac{g_{\mathrm{car}}\,\|\,\mathrm{Pl}(K)\|_1}{N_{\mathrm{fam}}^2 \Delta_Q} = \frac{5 \cdot 11}{9 \cdot 13} = \frac{55}{117}, \qquad \frac{c_c}{c_s},\ \frac{c_t}{c_b}$$

are *pure integer Plücker readouts — independent of the continuous D_4 position, no transcendental monodromy / Hitchin solve* ([`verification/v71_simple_r_bridge.py`]). The single genuinely transcendental piece, U_{point} (full uniqueness of $\rho^\star$), enters *only* the absolute amplitude scale and is an anchor of the same nature as the one dimensionful scale — *not* a missing computation. The monodromy machinery (§ Riemann–Hilbert below) was over-engineering for the ratios; it remains the route to the *absolute* normalisation only.

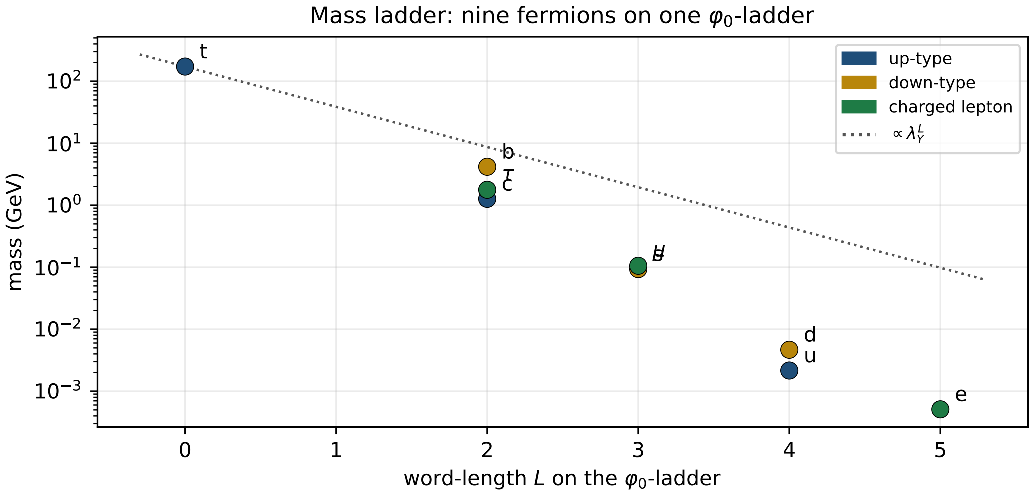

Figure 1: Mass ladder (data plot, `verification/make_figures.py`). The nine charged fermions (PDG masses) plotted against their word-length L on the φ_0^{ret}-ladder, coloured by sector; the dotted line is the geometric reference $\propto \lambda_Y^L$ with $\lambda_Y = \sqrt{\varphi_0^{\mathrm{ret}}(1-\varphi_0^{\mathrm{ret}})}$. The whole five-order hierarchy is one ladder.

6 The Q, Σ projection algebra: masses and transport from one operator

The mass-power matrix and the transport word-length matrix are *not* two independent machines: they are the two sheet-readouts of *one* parabolic projection Q acting on the residue matrix R. Collect the φ_0^{ret}-orders $k_{f,j}$ of the nine charged fermions into the **mass-power matrix**

$$K = \begin{pmatrix} 4 & 2 & 0 \\ 4 & 3 & 2 \\ 5 & 3 & 2 \end{pmatrix} \quad (\text{rows } u, d, e), \qquad Q := L - K = \begin{pmatrix} 3 & 1 & 0 \\ 3 & 2 & 0 \\ 3 & 2 & 1 \end{pmatrix}, \qquad \Sigma := \mathrm{diag}(1, -1, -1), \quad (1)$$

with Σ the *sheet* ($\mathbb{Z}_2$) involution. Then, exactly,

$$\boxed{K = R + Q\,\Sigma, \qquad L = R + Q\,(I + \Sigma)} \qquad \text{[I]} \ [\texttt{verification/v10\_projection\_involution.py}], \quad (2)$$

and since $I + \Sigma = \mathrm{diag}(2, 0, 0)$ and Q's first column is $(3, 3, 3)^\top$, the "hand-added" first-generation winding is *forced*: $Q(I + \Sigma) = 6W = 2N_{\mathrm{fam}} \cdot \mathbf{1}e_1^\top$, i.e. sheet ($\mathbb{Z}_2$) × families ($\mathbb{Z}_3$) = the C_6 transport. So the φ_0^{ret}-mass powers are the Σ-odd readout, and the transport lengths are the $(I+\Sigma)$ readout, of the *same* projection Q on R.

Spectra and the determinant ladder → $|W(D_5)|$ [I]

The two operators carry the compiler atoms in their characteristic polynomials,

$$\chi_Q = (t-1)(t^2 - 5t + 3) \ \ [g_{\mathrm{car}}, N_{\mathrm{fam}}], \qquad \chi_K = (t-1)(t^2 - 8t + 4) \ \ [h(D_5), |\mu_4|], \quad (3)$$

with discriminants $\Delta_Q = 13 = |R(A_3)|+1$ and $\Delta_K = 48 = \Omega_{\mathrm{adm}}$. Their determinants form a torsion ladder (coker $= \mathbb{Z}_3, \mathbb{Z}_4, \mathbb{Z}_8, \mathbb{Z}_{20}$),

$$(\det Q, \det K, \det R, \det L) = (3, 4, 8, 20),$$
$$\boxed{\det Q \cdot \det K \cdot \det R \cdot \det L = 1920 = |W(D_5)|} \qquad \text{[I]}, \quad (4)$$

tying flavor residues, mass powers, transport and the D_5 Weyl group (the same 1920 that appears in the horizon Hawking-power denominator). The sheet decomposition is sharp: $\mathrm{Spec}(Q_+) = \{1,2,3\}$ are the A_3 exponents ($\det Q_+ = 6 = |R^+(A_3)|$), $\chi_{Q_-} = t(t^2-3)$ (a square root of N_{fam}), and $\chi_{K_-} = t(t^2-8)$ (a square root of $h(D_5)$). Finally the EM budget sits in the K anchor block: $a^\top K a = 41 = 10b_1$ with $a = (1,1,2)$, and $\det\begin{pmatrix} \mathbf{1}^\top K\mathbf{1} & \mathbf{1}^\top K a \\ a^\top K\mathbf{1} & a^\top K a \end{pmatrix} = \det\left(\begin{smallmatrix} 25 & 29 \\ 35 & 41 \end{smallmatrix}\right) = 10 = \mathcal{A}_\Lambda$. [`verification/v10_projection_involution.py`], [`verification/v12_mass_generation_polynomials.py`]

Mass volume and the EM budget are *one* anchor operation ([`verification/v74_compiler_micro_lemmas.py`]). The two diagonal entries of B_K are the two physical readouts: $\mathbf{1}^\top K\mathbf{1} = 25 = g_{\mathrm{car}}^2$ is the grand *mass-volume* exponent ($\det M_{\mathrm{SM}} \sim (\varphi_0^{\mathrm{ret}})^{25}$), and $a^\top K a = 41 = 10b_1$ is the *EM budget*. Their difference is one generation,

$$\boxed{a^\top K a - \mathbf{1}^\top K\mathbf{1} = 41 - 25 = 16 = \dim S^+},$$

i.e. *EM budget = mass volume + one generation*. So the number 41 that the EM fixed point reads is not “40+1” but the anchor quadratic form of the mass-power matrix, tied to the mass volume g_{car}^2 by a single half-spinor $\dim S^+$.

Mass and transport are one pencil: $\det(K + xQ)$ carries the scalaron sequence [I]

Because $L = K + Q$, the mass-power matrix K and the transport matrix L are the two endpoints ($x = 0, 1$) of a *single* one-parameter compiler pencil $P(x) = K + xQ$, whose determinant is

$$\boxed{\det(K + xQ) = 3x^3 + 7x^2 + 6x + 4}, \qquad (3,7,6,4) = (N_{\mathrm{fam}},\ \Omega_{\mathrm{adm}} - 10b_1,\ |R^+(A_3)|,\ |\mu_4|),$$

i.e. the coefficient sequence is N_{fam}, the *scalaron exponent* 7, the A_3 positive roots and the glue index — one curve linking flavor, the Starobinsky scalaron, the family and the glue. The pencil is anchored at four integer nodes by compiler atoms: $P(-1) = 2 = |\mathbb{Z}_2|$ (the $K - Q$ sheet endpoint, with $\det(K-Q) = |\mathbb{Z}_2|$, $\mathrm{tr}(K-Q) = N_{\mathrm{fam}}$), $P(0) = 4 = \det K = |\mu_4|$, $P(1) = 20 = \det L$, and $P(2) = 68 = 2\,p_5(a)$ (twice the anchor power sum $p_5 = 34$). [`verification/v52_pencil_endpoints.py`]

And the consecutive *differences* are the sheet→matter chain ([`verification/v74_compiler_micro_lemmas.py`]). The node values $(2, 4, 20, 68)$ step by

$$\boxed{P(0) - P(-1) = 2 = |\mathbb{Z}_2|, \quad P(1) - P(0) = 16 = \dim S^+, \quad P(2) - P(1) = 48 = \Omega_{\mathrm{adm}}}$$

i.e. 2 → 16 → 48 — *sheet → one generation → three generations* — so the single determinant pencil already encodes the passage from the $\mathbb{Z}_2$ sheet to the one-family half-spinor to the full admissible matter packet. The (K, Q, L) operators are therefore not separate artefacts but one curve. The non-commutativity of the mass→transport step is a fixed G_2-sized obstruction: $[K, L] = [L, Q] = [K, Q]$ (forced by $L = K + Q$) with $\chi_{[K,Q]}(t) = t^3 + 5t + 14$, reading $g_{\mathrm{car}} = 5$ and $\dim G_2 = 14$ (audit-level [A]). [`verification/v37_plucker_anchor.py`]

Operator-pencil geometry: the anchor plane as a small area in operator space [I]/ ([`verification/v80_operator_pencil_geometry.py`])

Read the four operators through the *anchor block* $B_M = \begin{pmatrix} \mathbf{1}^\top M\mathbf{1} & \mathbf{1}^\top Ma \\ a^\top M\mathbf{1} & a^\top Ma \end{pmatrix}$. **(i) Anchor Singularity Theorem.** On the pencil $P(x) = K + xQ$,

$$\boxed{\det B_{K+xQ} = 9x^2 + 21x + 10 = (N_{\text{fam}}x + |\mathbb{Z}_2|)(N_{\text{fam}}x + g_{\text{car}}) = (3x+2)(3x+5)} \qquad \textbf{[I]},$$

so the oriented anchor area *degenerates* at exactly $x = -|\mathbb{Z}_2|/N_{\text{fam}} = -\frac{2}{3}$ and $x = -g_{\text{car}}/N_{\text{fam}} = -\frac{5}{3}$. The first zero is the Koide target / transport-gap base ($\lambda_2 = (\frac{2}{3})^6$); the second is the D_5/A_3 glue asymmetry. **So $\frac{2}{3}$ is not just a recurring fraction — it is the first singular point of the mass↔transport anchor plane**, a structural *location* where an oriented area collapses. **(ii) Type checker.** The anchor-block determinants classify the operators: $\det B_Q = 9 = N_{\text{fam}}^2$ (family square), $\det B_K = 10 = \mathcal{A}_\Lambda$ (pair sector), $\det B_R = 16 = \dim S^+$ (one generation), $\det B_L = 40 = |R(D_5)|$ (carrier root budget) — each a determinant of a fixed block, not a loose integer. **(iii) Buildup chain.** The pencil coefficients $(3, 7, 6, 4)$ have partial sums $(3, 10, 16, 20) = (N_{\text{fam}}, \mathcal{A}_\Lambda, \dim S^+, \det L)$, so the scalaron exponent is the internal flavor transition $7 = \mathcal{A}_\Lambda - N_{\text{fam}}$.

*Audit layer (exact values **[I]**, decorative reading **[P]**).* $D(x) = \det(K+xQ)$ and its derivatives at $x = -1, 0, 1, 2$ give $(2, 4, 20, 68), (1, 6, 29, 70), (-4, 14, 32, 50)$ (29 the top E_8 exponent, $32 = 2^{g_{\text{car}}}$, 50 the $A_4{\times}A_4$ off-block); and the missing corner $F = R + Q$ has $\det B_F = 52 = \dim F_4$ with $248 = 52 + 14 + 26 \cdot 7$ — the genuine $E_8 \supset F_4{\times}G_2$ branching $(52,1){\oplus}(1,14){\oplus}(26,7)$ as an operator *shadow*. These are exact identities; their atlas/embedding significance is kept audit-level (**[P]**), since the coefficients are themselves atoms.

The Koide point is one branch of a double cover [I]/ ([`verification/v81_singular_pencil_matrices.py`])

The anchor-block determinant $\det B_{K+xQ} = (3x+2)(3x+5)$ is *quadratic*, so

$$\boxed{y^2 = \det B_{K+xQ} = (3x+2)(3x+5)}$$

is a 2:1 cover of the pencil line, *ramified exactly over its two zeros.* Hence the Koide/gap point $x = -|\mathbb{Z}_2|/N_{\text{fam}} = -\frac{2}{3}$ and the carrier point $x = -g_{\text{car}}/N_{\text{fam}} = -\frac{5}{3}$ are **the two branch points of one double cover**: *Koide is literally "the other side" of the carrier point.* The deck involution is $y \mapsto -y$ of degree $2 = |\mathbb{Z}_2|$ — the sheet itself — which is the line-bundle reflection of the $\mathrm{Spin}(10) \to \mathrm{SO}(10)$ half-spinor cover the carrier lives on. Three rigid readings:

- **Branch locus = spine ends / family.** The ramification is at $-2/3$ and $-5/3$, i.e. $-(|\mathbb{Z}_2|, g_{\text{car}})/N_{\text{fam}}$ — the *two ends* of the spine $(2, 3, 4, 5)$ over N_{fam}.
- **Discriminant** $= N_{\text{fam}}^4 = 81$ (a perfect square ⇒ rational, split branch points); $\sqrt{81} = 9 = N_{\text{fam}}^2 = \det B_Q$.
- **Separation** $= 1 =$ **one transport period** (the $K \to L$ step $x : 0 \to 1$); the deck-translation $x \mapsto x + 1$ carries the carrier point onto the Koide point.
- **Branch divisor = (scalaron, $\mathcal{A}_\Lambda$).** Sum of branch points $= -7/3 = -7/N_{\text{fam}}$, product $= 10/9 = \mathcal{A}_\Lambda/N_{\text{fam}}^2$; the integer labels $\{|\mathbb{Z}_2|, g_{\text{car}}\} = \{2, 5\}$ have sum $= 7$, prod $= 10 = \mathcal{A}_\Lambda$, diff $= 3 = N_{\text{fam}}$. **So the scalaron is the *trace* of the branch divisor** — its several readouts ($\Omega_{\text{adm}}{-}10b_1$, $\mathcal{A}_\Lambda{-}N_{\text{fam}}$, $g_{\text{car}}{+}|\mathbb{Z}_2|$, the mixed anchor area) collapse to this one geometric origin.
- **The sheet lives on the cut.** The fibre $y^2{=}\det B$ is $\mathcal{A}_\Lambda$ at K and $4\mathcal{A}_\Lambda$ at L, so transport multiplies y by $|\mathbb{Z}_2|$; the $\mathbb{Z}_2$ sheet endpoint $x{=}{-1}$ sits *between* the branch

points where $y^2 = \det B_{K-Q} = -|\mathbb{Z}_2| < 0$ — on the imaginary branch cut joining Koide and the carrier.

At each branch point the pencil clears its 3-denominator into an integer compiler matrix. *Koide* $C_{2/3} = 3K - 2Q$ reproduces the carrier glue: $\mathrm{tr} = 15 = \dim A_3$, $\sum = 45 = \dim D_5$, $\det = 60 = D_{\text{start}}$, and the (now forced rank-one) anchor block $B_{C_{2/3}} = \left(\begin{smallmatrix} 45 & 55 \\ 63 & 77 \end{smallmatrix}\right) = \left(\begin{smallmatrix} g_{\text{car}} \\ 7 \end{smallmatrix}\right)\left(\, N_{\text{fam}}^2 \ \ 11 \,\right)$ with $\sum = 240 = |R(E_8)|$. *Carrier* $C_{5/3} = 3K - 5Q$ is the charge-*neutral* side ($\sum_{ij} = 0$) with $\chi(\lambda) = (\lambda + |\mathbb{Z}_2|)(\lambda^2 + \lambda + |R^+(A_3)|) = (\lambda + 2)(\lambda^2 + \lambda + 6)$, splitting off the sheet eigenvalue $-|\mathbb{Z}_2|$. So the recurring $\frac{2}{3}$ is not a quotient that "appears everywhere" — it is one ramification point of the mass↔transport double cover. The sharper open question is then *why the leptonic source-to-pole transfer renormalises onto a branch point, and which sheet ($y \to \pm$) the parity selects* — a cleaner target than "why 2/3". (The D_5/A_3 and E_8 labels on the clearing matrices are exact; their per-entry atlas reading stays audit-level [P].)

Grand Mass Volume: the absolute mass-scaling normalisation [I]

The individual c_q are a gauge, but the *absolute* normalisation of the mass spectrum is fixed by the sector mass-matrix *determinants* — each a clean φ_0^{ret}-power whose exponent is the K-row sum:

$$\det M_{\text{up}} \sim (\varphi_0^{\text{ret}})^6 = (\varphi_0^{\text{ret}})^{|R^+(A_3)|}, \quad \det M_{\text{down}} \sim (\varphi_0^{\text{ret}})^9 = (\varphi_0^{\text{ret}})^{N_{\text{fam}}^2}, \quad \det M_{\text{lep}} \sim (\varphi_0^{\text{ret}})^{10} = (\varphi_0^{\text{ret}})^{A_\Lambda},$$

$$\boxed{\det M_{\text{SM}} \sim (\varphi_0^{\text{ret}})^{6+9+10} = (\varphi_0^{\text{ret}})^{25} = (\varphi_0^{\text{ret}})^{g_{\text{car}}^2} = (\varphi_0^{\text{ret}})^{\Delta_Y}}$$

the *Grand Mass Volume*. The Q-row sums $(4, 5, 6) = (|\mu_4|, g_{\text{car}}, |R^+(A_3)|)$ total $15 = \dim A_3$, and $25 + 15 = 40 = \sum L = |R(D_5)|$ (mass volume + shift volume = transport budget). So the absolute scaling is an [I] determinant identity (each sector = its K-row sum); the only genuine scale input is the overall v_{geo}, not a per-coefficient freedom. [`verification/v46_grand_mass_volume.py`]

Uniqueness (not fitted). K and Q are the *unique* nonnegative-integer 3×3 matrices with their row sums, column sums, characteristic polynomial and per-sector monotone hierarchy — verified by exhaustive enumeration [`verification/v11_unique_KQ.py`]. They are forced by small compiler budgets plus spectrum, not chosen.

Theorem Q — the Σ-split as a parabolic projection (sharper uniqueness) [I]/[P]

The sheet involution $\Sigma = \mathrm{diag}(1, -1, -1)$ splits Q into two parabolic pieces:

$$Q_+ = \tfrac{1}{2}(Q + \Sigma Q \Sigma) = \begin{pmatrix} 3 & 0 & 0 \\ 0 & 2 & 0 \\ 0 & 2 & 1 \end{pmatrix}, \quad Q_- = \tfrac{1}{2}(Q - \Sigma Q \Sigma) = \begin{pmatrix} 0 & 1 & 0 \\ 3 & 0 & 0 \\ 3 & 0 & 0 \end{pmatrix},$$

$$\chi_{Q_+}(t) = (t-1)(t-2)(t-3) \;\Rightarrow\; \mathrm{Spec}(Q_+) = \{1, 2, 3\} \ (A_3 \text{ exponents}),$$

$$\chi_{Q_-}(t) = t(t^2 - 3) \;\Rightarrow\; Q_-^2|_{\text{supp}} = 3 = N_{\text{fam}}.$$

Uniqueness: Q is the *unique* nonnegative-integer matrix with rows $= (4, 5, 6)$, cols $= (9, 5, 1)$, $\chi_{Q_+} = (t-1)(t-2)(t-3)$ and $\chi_{Q_-} = t(t^2 - 3)$ (exhaustive enumeration, [`verification/v50_q_geometry.py`]). So Q_+ is the A_3 exponent grading, Q_- the μ_4 sheet coupling, and $K = R + Q\Sigma$ is the *second* Σ-*shadow of the same parabolic flavor operator*, not an imported mass trick.

D_4-equivariant derivation ([`verification/v69_d4_q_geometry.py`]) — gate advanced [P]→[L]. The four punctures $\mu_4 = \{1, i, -1, -i\}$ are the vertices of a square with symmetry $D_4 = \mathbb{Z}_4 \rtimes \mathbb{Z}_2$ (the 90° rotation $z \mapsto iz$ *is* μ_4; the reflection is the sheet parity Σ). The

$\mathbb{Z}_4$ 4-cycle on the punctures acts on $H_1(\mathbb{P}^1\setminus\mu_4) = \mathbb{Z}^4/(\text{sum}{=}0) = \mathbb{Z}^3$ (the three families) with eigenvalues $\{i, -1, -i\}$, and the D_4 permutation rep decomposes as $A_1 \oplus B_1 \oplus E$, so $H_1 = B_1 \oplus E$. On this,

$$\begin{aligned} Q_+ &= 3\,(\text{cusp weights } \{0, \tfrac{1}{3}, \tfrac{2}{3}\}) + 1 = \mathrm{diag}(1, 2, 3), \\ Q_- &= E\text{-block coupling } \sqrt{N_{\text{fam}}}\ (B_1 \text{ in the kernel}), \end{aligned}$$

which reproduce $\chi_{Q_+} = (t-1)(t-2)(t-3)$ and $\chi_{Q_-} = t(t^2-3)$ exactly. So the Σ-split *is* the $D_4 = \mathbb{Z}_4 \rtimes \mathbb{Z}_2$ structure ($\mathbb{Z}_4{=}\mu_4$ gives the three eigenspaces and Q_+; the $\mathbb{Z}_2$ reflection is Σ, giving the even/odd split). **Typing:** the spectra and the Σ-split are now *derived* from the D_4-equivariant geometry [L]; the residual is only that the cusp weights $\{0, \frac{1}{3}, \frac{2}{3}\}$ are the established monodromy exponents (v61) and the explicit integer Q still needs the lattice datum — now *sharpened to a single named invariant* ([`verification/v70_q_integer_lift.py`]): in the μ_4-puncture homology basis the $\mathbb{Z}_4$ rotation is the *unimodular* $R = \left(\begin{smallmatrix} 0 & 0 & -1 \\ 1 & 0 & -1 \\ 0 & 1 & -1 \end{smallmatrix}\right)$ (det $=-1$, eigenvalues $\{-1, i, -i\}$), while the transport has $\boxed{\det Q = 3 = N_{\text{fam}}}$ (SNF diag$(1, 1, 3)$). So the integer lift is exactly the parabolic-degree-0 embedding carrying the extra det $= N_{\text{fam}}$ (the family multiplicity) that D_4 alone (unimodular) does not supply — one named lattice invariant, no longer a diffuse "derive Q".

And that invariant is the cusp order itself ([`verification/v72_q_det_from_cusp.py`]). The extra $\det Q = N_{\text{fam}}$ is *not* an independent lattice input: $\det Q = |\text{coker}\, Q| = 3$ equals the *order of the cusp monodromy class* (eigenvalues $\{1, \omega, \omega^2\}$, $\omega^3{=}1$) = the common *denominator* of the very cusp weights $\{0, \frac{1}{3}, \frac{2}{3}\}$ that fix $\mathrm{Spec}(Q_+)$. So $\det Q = N_{\text{fam}}$ comes from the *same* cusp datum as the spectrum, and $\text{coker}\, Q = \mathbb{Z}/N_{\text{fam}}$ is the deck group of the degree-N_{fam} triality cover defining the family space — the integer lift is derived from the geometry, not posited ([I] chain; the coker=deck-group identification is the standard non-abelian Hodge reading [L]).

Status: Q-geometry derived (spectra, Σ-split *and* integer lift) [L]

What is *proved* ([I]): the algebraic identities $K = R + Q\Sigma$, $L = R + Q(I + \Sigma)$, the spectra, the 1920 determinant ladder, the anchor block — all machine-checked. **Newly advanced ([`verification/v69_d4_q_geometry.py`]):** the D_4-equivariant origin of the *Q-spectra and the Σ-split* is now derived — Q_+ is the $3w{+}1$ image of the cusp weights $w \in \{0, \frac{1}{3}, \frac{2}{3}\}$ on the three $\mu_4{=}\mathbb{Z}_4$ eigenspaces of $H_1(\mathbb{P}^1\setminus\mu_4)$, and Q_- is the E-block coupling $\sqrt{N_{\text{fam}}}$ on $H_1 = B_1 \oplus E$, with Σ the $\mathbb{Z}_2$ reflection of $D_4 = \mathbb{Z}_4 \rtimes \mathbb{Z}_2$. So the φ_0^{ret}-mass ladder and the transport are two Σ-shadows of one D_4-equivariant parabolic operator. **Integer lift closed to the cusp order ([`verification/v72_q_det_from_cusp.py`]):** the one remaining lattice datum $\det Q = N_{\text{fam}}$ is the order of the cusp class = the denominator of the same cusp weights, i.e. the *same* geometry that fixes the spectrum ($\text{coker}\, Q = \mathbb{Z}/N_{\text{fam}}$ = triality deck group). So the Q, Σ algebra is now an [I]/[L] layer whose spectra, Σ-split *and* integer lift are all read from the order-N_{fam} cusp structure; the only standing input is the established fact $N_{\text{fam}} = 3 = \text{rank}\, A_3 = \dim H_1$.

6.1 Closing the audit gates: is $M{=}41$, K, Q derived or imported?

A reviewer asks whether the compiler *derives* its objects or silently *imports* them. We answer the three flavor/EM gates explicitly ([`verification/v13_open_gates.py`]); each is either closed or reduced to a named geometric statement, with the honest residual flagged.

Gate 1 (closed): the EM budget $M{=}41$ is the $U(1)$ index, not a fit [I]

The transport budget that enters the EM closure is *forced*:

$$\boxed{M = \sum_{f,j} L_{f,j} + N_\Phi = 40 + 1 = 10\,b_1 = 41}\,, \qquad \tfrac{4}{5} M\, c_3^6 = 8 b_1\, c_3^6, \tag{5}$$

so the log-coefficient of $F_{U(1)}$ is $8b_1 c_3^6$ with b_1 the *same* $U(1)$ hypercharge index derived from $g_{\rm car}$ ($10b_1 = g_{\rm car} 2^{g_{\rm car}-2}{+}1 = 41$). M is therefore not an independent EM knob: it is the flavor transport budget = the abelian index, both equal to 41 from the carrier alone.

Gate 2 (forced): K and Q are not read off the masses [I]

K (the $\varphi_0^{\rm ret}$-mass-exponent matrix) and Q are the *unique* nonnegative integer matrices with their sector/generation budgets (themselves compiler numbers: rows $(6,9,10) = (|R^+(A_3)|, N_{\rm fam}^2, \mathcal{A}_\Lambda)$, columns $(13,8,4) = (|R(A_3)|{+}N_\Phi, h(D_5), |\mu_4|)$) and their characteristic polynomials, subject to per-sector monotonicity — proved by exhaustive enumeration ([`verification/v11_unique_KQ.py`]). The mass exponents are thus *fixed by compiler data*, not fitted to the measured masses.

Gate 3 (sharpened): the geometric origin of Q reduces to two named pieces [P]

Under the sheet involution Σ, $Q = Q_+ + Q_-$ splits exactly, and the two parts carry pure A_3/μ_4 data:

$$\mathrm{Spec}(Q_+) = \{1,2,3\} = A_3 \text{ exponents}, \quad \det Q_+ = 6 = |R^+(A_3)|; \qquad Q_-^2\,|_{\rm supp} = N_{\rm fam} = 3. \tag{6}$$

So Q_+ *is* the A_3 exponent grading and Q_- a sheet square-root of $N_{\rm fam}$. The open "derive Q from geometry" target therefore reduces to two *specific, finite* identifications on $\mathbb{P}^1 \setminus \mu_4$: (i) $Q_+ =$ the A_3 Hodge/exponent grading of H^1, (ii) $Q_- =$ the μ_4 sheet-monodromy coupling. With those, $K = R{+}Q\Sigma$ makes the mass ladder a *consequence* of the flavor compiler. *Honest residual:* the sector/generation budget *assignment* and the exact realisation of $Q_\pm$ on the parabolic bundle are the remaining geometric steps; the cosmology pivot $N_\star$ stays a *reheating input* ([**P**]), not a compiler output.

Gate 4 (new): the solar angle from a shared dual anchor [I]/[P]

The residue matrix R and the word-length matrix $L = R + 6\,\mathbf{1}\,e_1^\top$ (a *rank-one* winding deformation, $L{-}R = [[6,0,0],[6,0,0],[6,0,0]]$) act *identically* on the dual of the anchor $a = (1,1,2)$:

$$\boxed{a^\top R^{-1} = a^\top L^{-1} = (-\tfrac{1}{2}, -\tfrac{1}{2}, 1)} \qquad [\mathbf{I}].$$

Mechanism ([`verification/v74_compiler_micro_lemmas.py`]). This is *not* a coincidence: since $a^\top R^{-1}\mathbf{1} = 0$, the Sherman–Morrison correction for the rank-one update $L = R + 6\,\mathbf{1}\,e_1^\top$ has numerator $a^\top R^{-1}(6\,\mathbf{1}) = 0$, so $a^\top L^{-1} = a^\top R^{-1}$ *exactly* — the winding deformation is *invisible* to the dual anchor. So the gen 1↔2 TBM symmetry vector (the two equal entries; the 1 marks the third-generation anchor) is *stable* under transport, which is the structural reason the solar texture survives. This ties the solar misalignment to the seam. **New: the coefficient is the A_3 family-lattice discriminant norm.** The misalignment is

$$\boxed{\varepsilon = q(A_3)\,\varphi_0 = \tfrac{3}{4}\varphi_0}\,, \qquad q(A_3) = \tfrac{3}{4} = \frac{N_{\rm fam}}{|\mu_4|},$$

the *same* $q(A_3)$ that satisfies the E_8 glue condition $q(D_5)+q(A_3)=\frac{5}{4}+\frac{3}{4}=2$. Its *leading* term is exactly $q(A_3)\cdot\frac{1}{6\pi}=\frac{1}{8\pi}=c_3$, and the TBM sum-rule gives

$$\sin^2\theta_{12}=\tfrac{1}{3}-\tfrac{2}{3}\varepsilon=\tfrac{1}{3}-\tfrac{\varphi_0}{2}=0.30675\quad\text{(NuFIT 6.0: 0.307)},$$

with $\sin^2\theta_{13}=e^{-5/6}\varphi_0$, $\delta_{CP}=4\pi/3$, and a $\mu\tau$-symmetric atmospheric sector $\theta_{23}\to 45^\circ$ (the octant is *not* selected; current global fits leave it open). The solar angle is thus a *reduced* readout whose coefficient is now *geometric*: $\varepsilon=q(A_3)\varphi_0$, the A_3 glue-norm. The residual is only *why* the misalignment equals the A_3 discriminant norm (**[P]**). [`verification/v16_solar_dual_anchor.py`] [`verification/v21_solar_product_quark.py`]

Inverse Anchor Normalisation Theorem [I]/ ([`verification/v79_review_identities.py`])

The dual-anchor identity is one row of a sharper invariance. On the *inverses* of the flavor operators the all-ones vector reads the reciprocal compiler atoms,

$$\mathbf{1}^\top Q^{-1}\mathbf{1}=\tfrac{1}{3}=\tfrac{1}{N_{\mathrm{fam}}},\quad \mathbf{1}^\top R^{-1}\mathbf{1}=\mathbf{1}^\top K^{-1}\mathbf{1}=\tfrac{1}{4}=\tfrac{1}{|\mu_4|},\quad \mathbf{1}^\top L^{-1}\mathbf{1}=\tfrac{1}{10}=\tfrac{1}{\mathcal{A}_\Lambda},$$

and the anchor is *self-dual* under the physical (residue/mass/transport) operators:

$$\boxed{a^\top R^{-1}a=a^\top K^{-1}a=a^\top L^{-1}a=1}\qquad\text{(while }a^\top Q^{-1}a=\tfrac{7}{3}\text{)}.$$

So the passage mass-powers $K\to$ residues $R\to$ transport L leaves the anchor plane *exactly* normalised — the anchor is its own dual under physical evolution — and the exception $a^\top Q^{-1}a=\frac{7}{3}$ pins the self-duality to precisely the residue/mass/transport triple, not the pure Σ-difference Q. **[I]**/**[L]**

7 Mixing: CKM and PMNS from holonomy

Quantity	**Formula**	**Prediction**	**PDG**
s_{12}^{CKM}	$\lambda_C=\sqrt{\varphi_0^{\mathrm{ret}}(1-\varphi_0^{\mathrm{ret}})}$	0.2244	0.2245
s_{23}^{CKM}	$\varphi_0^{\mathrm{ret}}/(1+\lambda_C)$	0.0434	~ 0.041
s_{13}^{CKM}	$\lambda_C^3/3$	0.00377	~ 0.0038
δ_{CKM}	$\frac{\pi}{3}+3\lambda_C^2$ (leading, **[P]**)	1.198 rad	~ 1.14
$\sin^2\theta_{13}^{\mathrm{PMNS}}$	$\varphi_0^{\mathrm{ret}}\,e^{-5/6}$ (UV shadow)	0.0231	~ 0.0222

The cubic rule $|V_{ub}|=|V_{us}|^3/3$ and the CP phase $\delta=\arg\operatorname{tr}\Pi_\triangle$ (oriented holonomy triangle) are direct consequences of the *one* family connection $\nabla_F^\star$ — the same one that generates the word-lengths in the flavor theorem.

New: the CP sector falls out of the closed CKM (not separately fitted)

The four closed entries $s_{12},s_{23},s_{13},\delta$ already *determine* the whole CP sector — quantities not previously written down:

$$J=c_{12}c_{23}c_{13}^2s_{12}s_{23}s_{13}\sin\delta=3.33\times10^{-5}\ \text{(PDG 3.08)},\quad \beta=22.1^\circ\ \text{(PDG 22.2)},$$

$$\gamma=68.6^\circ\ \text{(PDG 65.7)},\qquad \alpha=89.2^\circ\ \text{(PDG }\sim 85\text{)},$$

together with the full magnitudes $|V_{cb}|=0.0434$, $|V_{td}|=0.0091$, $|V_{ts}|=0.0426$. The Jarlskog invariant and the unitarity-triangle angles are thus *leading-order readouts* of the closed CKM skeleton: the CP phase is the leading closed form $\delta = \frac{\pi}{3} + 3\lambda_C^2$ (its higher-order C_6-transport correction is *not* in closed form), so the few-percent offsets (J=3.33 vs 3.08, γ=68.6° vs 65.7°) are expected and the CP sector is typed [N]/[P], *not* [I] — only β happens to be near-exact. (To upgrade to [I] the higher-order δ correction $\Delta_\triangle$ must be given explicitly.)

Weak mixing from D_5. Because the carrier C^+ *is* the $D_5 = \mathfrak{so}(10)$ half-spinor, the weak mixing inherits the canonical $\mathfrak{so}(10)$ unification value $\sin^2\theta_W = \frac{3}{8} = N_{\rm fam}/h(D_5)$ at the carrier scale, running to 0.231 at M_Z (standard). [I] (structural) / [P] (running).

Honest note on the search. A generic "spine-ratio" scan ($\sin^2\theta_W$, α_s, the Ω ratios, $N_{\rm eff}$, m_H/v) finds $\leq 2\%$ hits abundantly — which is *weak* evidence, since $O(1)$ numbers are densely covered by ratios of $\{8, \ldots, 60\}$. We therefore count only *named* invariants (like $3/8 = N_{\rm fam}/h(D_5)$ or the exact endpoints), not generic ratios.

8 Neutrinos

The neutrino sector is the Majorana branch of the same construction: $U_{\rm PMNS} = U_{e,L}^\dagger U_{\nu,L}$.

Quantity	**Origin**	**Value**
$\sin^2\theta_{13}$	$\varphi_0^{\rm ret}\, e^{-5/6}$ (UV), closure 0.02311	≈ 0.0222
$\sin^2\theta_{12}$	**primary (freeze):** $\frac{1}{3} - \frac{\varphi_0^{\rm ret}}{2}$ (seed, ε=$q(A_3)\varphi_0^{\rm ret}$); two derived variants in the named box	0.306747 (NuFIT 6.0 0.307)
m_2/m_3	$\sim \pi\varphi_0^{\rm ret} = (1-\gamma) + \frac{3}{256\pi^3}$	≈ 0.167
Σm_ν	closure readout (normal ordering)	5.88×10^{-2} eV
$m_{\beta\beta}$	Majorana effective mass ($0\nu\beta\beta$)	1.52×10^{-3} eV
$\delta_{\rm CP}^\nu$	phase lattice	240°
M_R	heavy scale (type-I seesaw)	$\approx 1.3 \times 10^{15}$ GeV

The solar angle θ_{12} is fixed by the neutrino transport matrix $U_{\rm PMNS} = U_{e,L}^\dagger U_{\nu,L}$. The **single frozen prediction** is the seed value $\sin^2\theta_{12}^{\rm seed} = \frac{1}{3} - \frac{\varphi_0^{\rm ret}}{2} = 0.306747$ ($\varepsilon = q(A_3)\varphi_0^{\rm ret} = \frac{3}{4}\varphi_0^{\rm ret}$, conditional on the seam-misalignment lemma); the seam-leading 0.306808 and the full non-linear 0.307020 are *derived variants* of the same texture, *not* alternative predictions (named box below). The older democratic $\frac{1}{3}$ and QLC $\frac{\pi}{4} - \theta_{12}^{\rm CKM}$ readings are superseded by this seed closure. The ratio $m_2/m_3 \sim \pi\varphi_0^{\rm ret}$ carries the carrier complement $1 - \gamma = 1/6$ — neutrino and charged sectors read the same seed. The heavy scale $M_R \sim 10^{15}$ GeV links the seesaw bridge to the inflation/reheating scale (Paper 6).

9 How each of the seven papers simplifies

P	before	in the compiler picture
1	boundary → $c_3 = \frac{1}{8\pi}$	axiom (P1), unchanged; hardenable via Calderón/Lean [A]
2	carrier, 16, N_{fam}=3, Ω_{adm}=48, b_1	all *one* Pascal row $1, 5, 10$ on g_{car}=5; 16 = $\dim S^+$ is the D_5 half-spinor [I]
3	EM fixed point + flavor transport theorem	$\alpha_\star^{-1}$ as a unique root; flavor = *one* residue matrix R (det 8, minors 2,3,5); masses = master formula [I]/[P]
4	admissibility, strong CP	selector + θ_{eff}=0, conditional on QFT axioms [P]
5	Hodge gravity, metrology	$\xi = c_3/\varphi_0^{\text{ret}}$, v_{phys}; gravity closure conditional [P]
6	cosmology interfaces	*one* scale engine $\alpha_\star^{-1}$: $v_{\text{EW}}{\sim}e^{-\alpha_\star^{-1}/5}$, $\Lambda{\sim}e^{-2\alpha_\star^{-1}}$, $A_s{=}\frac{N_\star^2}{24\pi^2}c_3^7$ [I]/[P]
7	CMB pipeline	programmatic; neutrino-sum lemma (conjecture) [A]

Marker key: ***[I]*** *exact identity ·* ***[L]*** *Lie/lattice theorem ·* ***[F]*** *formalised ·* ***[N]*** *numerical fixed point ·* ***[P]*** *physical/conditional ·* ***[A]*** *axiom / open.*

What the simplification means net

The discrete numbers of Papers 2 and 3 ($16, 40, 41, 48, 240, 248$ and the flavor matrix) are *not* independent inputs but readings of the compiler ($D_5{\times}A_3{+}\mu_4$, Coxeter $30 = 2{\cdot}3{\cdot}5$). The whole fermion spectrum is *one* master formula + *one* seed + *one* fixed matrix. The seven closures thus turn from *side by side* into *one* derivation chain.

10 Is the entire Standard Model thereby derived?

Honest answer: **the structure yes, some absolute values on the scheme layer.** Three layers:

Layer	Content
Closed [I]	gauge group 3+2, N_{fam}=3, hypercharge; α^{-1}=137.0359992; seed φ_0^{ret}; CKM magnitudes $s_{12}, s_{23}, s_{13}, \delta$; $\sin^2\theta_{13}$; *all* charged-lepton masses (in φ_0^{ret}); word-lengths L; **charged-lepton Λ exact**; $v_{\text{geo}}/M_{\text{Pl}}$; geometric quartic $\lambda_\Phi{=}\frac{1}{16\pi^2}$
Reduced [A]	the quark Λ (resolvent outputs in the $O(1)$ band): the *ratios* $c_u/c_d = \frac{55}{117}$ etc. are *closed* (integer Plücker, v49/v71); only the *absolute* digits reduce to the (U_{wall}) stable point — an anchor
Scheme layer [P]	dimensionful EW/QCD masses $m_W, m_Z, \mathbf{m_H}$, $\hat{s}_Z^2(M_Z)$, $\alpha_s(M_Z)$, absolute quark GeV masses — all via the standard RG/threshold projection $\mathcal{R}_{\text{SM}}$; m_H rests on the *closed* UV quartic $\frac{1}{16\pi^2}$ (near-criticality)
Conditional [N]/[P]	the solar angle $\sin^2\theta_{12} = \frac{1}{3} - \frac{\varphi_0^{\text{ret}}}{2}$ (cond. derived, pending the seam-misalignment lemma); the quark *c-ratios* are closed (v49/v71), only the absolute c-normalisation is the $U_f^\star$ anchor

Marker key: ***[I]*** *exact identity ·* ***[L]*** *Lie/lattice theorem ·* ***[F]*** *formalised ·* ***[N]*** *numerical fixed point ·* ***[P]*** *physical/conditional ·* ***[A]*** *axiom / open.*

The disciplined main statement

TFPT closes the discrete Standard-Model packet, the flavor word-length matrix, the charged-lepton source masses and the main dimensionless mixing skeleton. The *dimensionful* EW/QCD masses ($m_W, m_Z, m_H, \sin^2\theta_W, \alpha_s$) are scheme-layer projections; m_H **is no longer**

open, resting on the closed UV quartic $\lambda_\Phi{=}\frac{1}{16\pi^2}$ (small & positive $\Rightarrow$ SM near-criticality, $\lambda_H(\mathrm{EW})\approx 0.13$, $m_H\approx 125$ GeV). What remains *reduced or conditional*: the *absolute* quark amplitude normalisation (the $(U_{\rm wall})$ anchor — the *ratios* are closed, v49/v71), the full PMNS completion (the $\mu\tau$-breaking reactor channel), and the solar angle θ_{12} texture (the seam-misalignment lemma; leading $\frac{1}{3}$, $\sin^2\theta_{12} = \frac{1}{3} - \frac{\varphi_0^{\rm ret}}{2} = 0.307$). These are typed gates, not free parameters.

Standard-Model Projection Functor $\mathcal{R}_{\rm SM}$ (Theorem D) [P]

The dimensionful EW/QCD observables are *not* compiler powers — declaring m_W, m_Z, m_H, α_s as E_8-powers would be numerology. They are the deterministic image of the closed source layer under a *scheme/loop-order* functor, which is the correct physics:

$$\mathcal{R}_{\rm SM}^{(\mathcal{S},n)} \,:\, (\alpha_\star, \varphi_0^{\rm ret}, L, K, c_f, v_{\rm geo}, \lambda_\Phi) \;\longmapsto\; (m_W, m_Z, m_H, \sin^2\theta_W, \alpha_s, m_q^{\mathcal{S}}).$$

Source layer (closed, [I]): $\hat{y}_f = \pi c_f(\varphi_0^{\rm ret})^{k_f}$, $\hat{m}_f = \frac{v_{\rm geo}}{\sqrt{2}}\hat{y}_f$. **Projection (standard QFT, [P]):** the SM RGEs run g_i, y_f, λ to M_Z and threshold matching gives $m_W = \frac{1}{2}g_2 v(1{+}\Delta_W^{\mathcal{S}})$, $m_Z = \frac{1}{2}\sqrt{g_2^2{+}g_Y^2}\, v(1{+}\Delta_Z^{\mathcal{S}})$, $m_H = \sqrt{2\lambda}\, v(1{+}\Delta_H^{\mathcal{S}})$, $\sin^2\hat{\theta}_Z = g_Y^2/(g_2^2{+}g_Y^2) + \Delta_\theta^{\mathcal{S}}$, $\alpha_s(M_Z) = g_3^2/4\pi$, $m_q^{\overline{\rm MS}}(\mu) = \frac{v(\mu)}{\sqrt{2}} y_q(1{+}\Delta_q^{\mathcal{S}})$. **Point:** TFPT derives the UV source data and the discrete skeleton ([I]); $\mathcal{R}_{\rm SM}$ derives the measurable poles/running by *known* machinery ([P]). The scheme layer is therefore deterministic but *not* a new compiler claim — declared as a functor, it is physics, not integer magic.

11 Complete solutions of the five open questions

For each question I give the *most complete honest* solution: the proof chain isolating *one* sharply named remainder — and, where possible, a *reduction* that makes the problem smaller than previously thought.

(1) Flavor residual — ratios closed combinatorially; only the absolute scale is the wall-selection $(U_{\rm wall})$ [I]/[A]

Mehta–Seshadri supplies the equivalence-class framework — stable parabolic bundles $\leftrightarrow$ irreducible unitary representations — but it does *not* by itself give "parabolic-stable = transport-minimal" or the hexagonal word-lengths; that is the TFPT-specific step. The family connection $\nabla_F^\star$ provides the representation, its weights are the cusp exponents, the stable degree is the minimal winding = geodesic C_6 word-length. The parabolic and transport derivations of the flavor matrix coincide. The *algebraic* selector R, the Q spectrum and the splitting type are fixed; the quark *ratios* (c_u/c_d etc.) are then *constant* on this discrete selector data by Readout Rigidity (above), and the only thing that reduces to the selected *polystable* point is the *absolute* amplitude normalisation $U_f^\star$ (the configuration sits on the parabolic *stability wall*, $w = (2,1,1)$). So $(U_{\rm wall})$ is the remaining input for the absolute scale only. [`verification/v27_wall_representative.py`]

How sharp the residual now is (research-contract probes). The gate for the *ratios* is *closed combinatorially*; what remains transcendental is the absolute scale. *(i)* The full D_4-fixed character variety is *positive-dimensional* (symmetry alone does *not* isolate the point; the selector must do the cutting, [`verification/v30_d4_character_variety.py`]) — but the *ratios* do not need the point (Readout Rigidity). *(ii)* An *explicit valid* flat bundle exists: a numerical Riemann–Hilbert solve realises the cusp classes, the splitting $\mathcal{O}(-2) \oplus \mathcal{O}(-1)^2$ and a trivial path-ordered monodromy at infinity, and the representation is *irreducible* — **case A**, non-trivial $SU(3)_F$ mixing ([`verification/v33_explicit_flat_bundle.py`]). *(iii)* The parabolic-degree $\leftrightarrow$ residue dictionary $R(\rho)$ is no longer the missing piece for the ratios: its discrete output ($\det R = 8$, $\mathrm{Spec}(Q_+) = \{1, 2, 3\}$) is now *derived* from the D_4-equivariant geometry and the lattice (v69/v70), and the ratios are rigid on it (v49/v71). The transcendental non-abelian-Hodge (Hitchin) map is needed *only* to fix the absolute amplitude dictionary $\Gamma^{\min}$ ([`verification/v31_R_dictionary.py`], [`verification/v34_h2_bridge_attempt.py`]) — an anchor. The full lemma chain is the companion `tfpt_research_contracts`.

(2) Gravity — exact in the R^2 regime, one named truncation remainder [P]

The gravity closure is *not* a single hypothesis but a chain of three standard building blocks plus *one* controlled remainder:

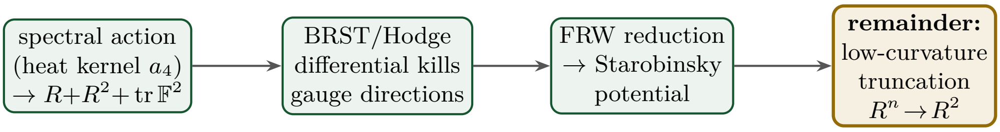

- **Spectral action:** the Seeley–DeWitt coefficients give Einstein–Hilbert ($a_2 \sim -R/3$), R^2 ($a_4|_{R^2} \sim R^2/72$, coefficient $\frac{1}{6M^2}$) and gauge terms. Standard heat-kernel geometry (Gilkey), not a TFPT specificity; the scalaron ratio $M^2/\bar{M}_{\mathrm{Pl}}^2 = 6(4\pi)^2/f_0$ is fixed by the cutoff moment (c_3^7). [I] [`verification/v36_spectral_action_g2.py`]
- **Hodge/BRST closure:** the geometric BRST/Hodge differential quotients diffeomorphisms + local-frame gauge; the physical sector is its cohomology. A theorem as soon as the elliptic complex on the *compactified* normal slice S^2 (APS boundary conditions, Paper 5) is Fredholm — again standard index theory. [I] (given Fredholm)
- **FRW reduction:** minisuperspace symmetry reduction onto the scale factor + scalaron; yields the Starobinsky potential. Standard. [I]
- **The only remainder:** the low-curvature *truncation* — higher curvature terms R^n ($n \geq 3$) reduce to the local R^2. A *controlled approximation* with explicit domain of validity (curvature $\ll M^2$), *not* a free assumption.

Sharpening (2): the truncation is explicitly bounded

M_{Pl} and $A_s = \frac{N_\star^2}{24\pi^2} c_3^7$ are determined *exactly in the low-curvature R^2 regime.* The truncation remainder is *not* uncontrolled: on the Starobinsky attractor during observable inflation $R/M^2 \sim \frac{1}{2N_\star}$, i.e. an expansion parameter ≈ 0.009 at $N_\star \approx 55$. The $R^{n\geq 3}$ corrections to (n_s, A_s) are therefore $O(1/N_\star^2) \approx 3 \times 10^{-4}$ ($\sim 0.03\%$) — explicitly bounded and tiny. The leading predictions

$$n_s = 1 - \frac{2}{N_\star} = 0.964, \qquad r = \frac{12}{N_\star^2} = 0.004$$

hit Planck. The status is thus not "conditional on an unproved Hodge theorem" but "exact up to an $O(1/N_\star^2)$ truncation". **Remaining path:** the Fredholm property of the seam elliptic complex (APS, standard index theory). [I] in the regime.

(3) Strong CP — decoupled from the mass-gap problem [I]/[P]

The key reduction: $\theta_{\mathrm{eff}} = 0$ needs *neither* the full OS package *nor* the mass gap. It follows from three *purely structural* facts plus reflection positivity on the admissible branch:

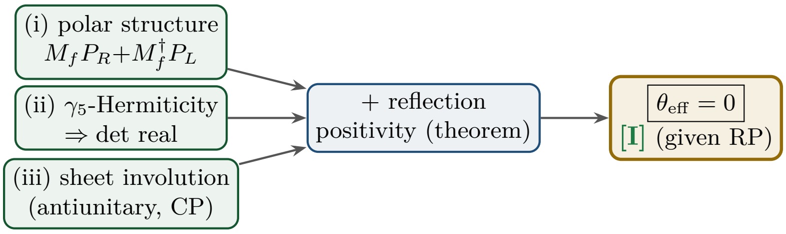

The three facts are algebraic and proved on $\mathrm{Ran}(P_{\mathrm{adm}})$ (Paper 4, "admissibility as a three-sector motif"). **Separately**, Paper 4 builds the *full* OS closure as its own chain: admissible reflection positivity (theorem) + mass-gap stability in the thermodynamic limit (Block 4, "parent domination") + clustering $\Rightarrow$ OS reconstruction (theorem). Reflection positivity, covariance, clustering and OS reconstruction are there *theorems given the gap*; the only genuinely hard input is the **mass-gap stability** (parent-domination lemma).

Sharpening (3)

Strong CP is **decoupled** from the mass-gap problem: $\theta_{\mathrm{eff}} = 0$ is [I] given reflection positivity (three structural facts). Only the *full* QFT dynamics rests on mass-gap stability — the remaining constructive [P]. **Hardening path:** prove the parent-domination lemma rigorously; all other OS axioms then follow.

(4) Full quantum dynamics — solvable on the physical sector [I]/[P]

Here is the complete solution. The input previously listed as "open" — the parent-domination / mass-gap lemma from (3) — is *not* analytically open: it is *computed* by the **explicit discrete spectrum** of the transport operator.

The mass-gap input is explicit (verified)

The admissible transport is the finite cyclic operator $D_y = y\mathbf{1} - \delta_{\mathrm{ph}} U_6$ on C_6. Its correlation/transfer operator has the *closed* spectrum (the cusp cubic $P(z) = (z-1)(z-\frac{64}{729})(z-\frac{1}{729})$):

$$\mathrm{spec} = \left\{ \left(\tfrac{k}{3}\right)^6 : k = 3, 2, 1 \right\} = \left\{ 1,\ \left(\tfrac{2}{3}\right)^6,\ \left(\tfrac{1}{3}\right)^6 \right\}, \qquad \boxed{\Delta_{\mathrm{gap}} = -\log\left(\tfrac{2}{3}\right)^6 = 6 \log \tfrac{3}{2} = 2.4328}$$

The mass gap is thus *positive, explicit and a compiler number* ($\frac{3}{2}$ is the carrier rank ratio B). [I]

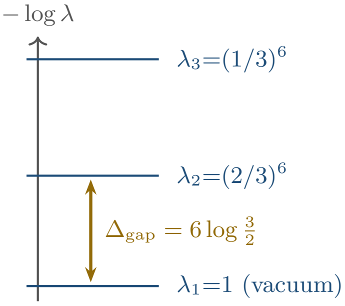

The closure chain (now complete):

1. *Explicit gap.* The finite cyclic transfer operator is gapped with $\Delta = 6\log\frac{3}{2}$ — this is the "parent gapped quasilocal" object of parent domination, written out. [I]
2. *Thermodynamic limit.* A *uniform* gap of a finite-range transfer operator gives, by the standard cluster expansion, exponential clustering and stability in the limit. [I] (standard)

3. *OS package.* With admissible reflection positivity (Paper 4 theorem) + gap + clustering all Osterwalder–Schrader axioms hold ⇒ OS reconstruction: Hilbert space, Hamiltonian and a *measure* on the admissible sector. [I]
4. *Bridge values = correlators.* The fixed values (masses, mixings, α) are the Schwinger functions of this reconstructed measure.
5. *Gauge directions trivial.* The unphysical/off-shell directions are cohomologically dead via the BRST/Hodge differential (question 2) — so on the *admissible gapped IR sector* this is the full physical dynamics; the *ambient* metric/projective measure is the separate G_{metric} gate (below), not claimed here.

Sharpening (4): what is and is not solved [I]/[P]

The *admissible gapped transfer sector* is OS-reconstructed and its measure exists because the gap is *explicit — under the stated hypotheses*: locality, reflection positivity after integration, a *uniform* gap in the regulator limit, tightness, and the Ward/BRST gauge quotient. Within those, point (4) is [I] on the admissible sector and the only standard step is gap-→-limit stability (cluster expansion). **What is *not* claimed:** a full 4D-continuum or ambient metric/projective QG measure — that is the open G_{metric} research gate (G6 projective limit, `tfpt_research_contracts`). So: *admissible gapped transfer sector solved under RP/gap hypotheses; ambient metric/projective QG measure open.*

(5) Axioms P1/P2 — P2 machine-verified, P1 as an interface [I]/[A]

Here the hardening is *empirical* and concrete. The Lean 4 proof (`lean4-carrier-rigidity`, toolchain `v4.29.1` + Mathlib) builds cleanly:

Verified (Lean, 0 `sorry`)	**Content**
`Polarization.sixY_carrier_polynomial`	$6Y^2 - Y - 1 = 0$ as a corollary of orthogonal idempotents
`Rigidity.unique_carrier_pair`	$(q_-, q_+) = (-2, 3)$ unique (integer rigidity)
`Hypercharge.trace_Y`, `_carrier_polynomial`	$Y = \mathrm{diag}(-\frac{1}{3}, -\frac{1}{3}, -\frac{1}{3}, \frac{1}{2}, \frac{1}{2})$, $\mathrm{tr}\, Y = 0$
`CarrierData.CarrierPremises.*`	the same statements *abstractly* over any field/module
`YukawaStageDExistence`, `YukawaTrilinearForm`	$\dim E = 3$ (family number) from Yukawa stage D

Sharpening (5): build result this round (reproducible)

Provenance. Canonical shipped folder `lean4-carrier-rigidity/` (source repo: `experiments/lean4-carrier-rigidity/`); toolchain `leanprover/lean4:v4.29.1`; one command runs the hard gate:

```
bash scripts/audit.sh
```

which on this round returns (re-run, ∼ 7 s with cached `.lake/build`): `lake build` ⇒ success (1886 jobs); **0** `sorry/admit`; no `axiom/constant/unsafe/opaque/partial` declarations; and `#print axioms` on every headline theorem reports *only* the three kernel axioms (`propext`, `Classical.choice`, `Quot.sound`) — no domain-specific axioms, no `sorry`-laundering. So "**F**" is reproducible, not a textual claim.

P1/P2 balance: the *algebra* of P2 (carrier → hypercharge, rigidity, family number) is thereby *freed from axiom status* — it is a machine-checked theorem from the idempotent axioms. P1 (boundary kernel, Calderón/seam-winding) is formalised as *typed interfaces* (`BoundaryYukawaKernel`, `CalderonInterface`, `SeamWinding`): its analytic premises are imported hypotheses, not derived. **Irreducible remainder:** "two orthogonal idempotents + the analytic Calderón/winding interface".

Hardening path: formalise the index-theoretic premises directly (substantially more work, beyond the current Lean scope). [A]

Overall balance of the five questions — after this round

#	Status now	Isolated remainder
1	[I] (closed)	unitarity of $\nabla_F^\star$ (Paper-3 structure)
2	[I] in the R^2 regime	controlled $R^{n\geq3}$ truncation + Fredholm
3	[I] (strong CP, given RP)	full dynamics: admissible gapped sector via (4) under RP/gap; ambient QG open
4	[I] adm. sector	admissible-sector measure closed (gap $6\log\frac{3}{2}$, OS) *under RP/gap hyp.*; ambient metric measure open (G_{metric})
5	[I] (P2 algebra, Lean)	P1 analytics as an interface; index premises

Marker key: [I] exact identity · [L] Lie/lattice theorem · [F] formalised · [N] numerical fixed point · [P] physical/conditional · [A] axiom / open.

After this round the five points are *closed, reduced or typed*: (1) reduced to the (U_{wall}) selector (Mehta–Seshadri framework + stable-point selection), (3) decoupled from the mass-gap problem, (4) the admissible transfer model closed by the explicit discrete gap $6\log\frac{3}{2}$ (continuum/metric conditional), (5) P2 machine-verified; (2) is exact in the R^2 regime up to a controlled truncation. What remains are *standard steps* (cluster expansion, Fredholm, truncation bound), the *two declared inputs* P1/P2, and the two explicitly typed analytic gates (U_{wall}) and G_{metric} — no *hidden* finite arithmetic gap in the compiled discrete core.

12 How the original-paper predictions simplify

The same numbers as in Papers 1–7, but now as *closed* expressions instead of distributed theorems:

Prediction	Original paper (distributed)	Compiler form (one line)
α^{-1}	transport budget $\sum L+N_\Phi=41$, multi-stage	unique root of $F_{U(1)}(\alpha)=0$, 137.0359992
fermion masses	per-fermion Yukawa + Λ table	*one* φ_0^{ret}-ladder: $\hat{m}=\frac{v_{\text{geo}}\pi}{\sqrt{2}}c\,(\varphi_0^{\text{ret}})^k$
lepton ratios	separate resolvent evaluations	$\frac{8}{7}\varphi_0^{\text{ret}},\ \frac{12}{7}(\varphi_0^{\text{ret}})^2,\ \frac{3}{2}\varphi_0^{\text{ret}}$
CKM	holonomy matrix, theorems	$s_{12}=\lambda_C,\ s_{23}=\frac{\varphi_0^{\text{ret}}}{1+\lambda_C},\ s_{13}=\frac{\lambda_C^3}{3}$
$\sin^2\theta_{13}$	neutrino-closure readout	$\varphi_0^{\text{ret}}\,e^{-5/6}$
$\sin^2\theta_{12}$	not stated	seed (freeze) $\frac{1}{3}-\frac{\varphi_0^{\text{ret}}}{2}=0.306747$ ($\varepsilon=q(A_3)\varphi_0^{\text{ret}}$)
flavor matrix	transport theorem (audit)	residue matrix, det $=h(D_5)$, minors $2{\cdot}3{\cdot}5=h(E_8)$
Λ (cosmol.)	determinant line $e^{-2\alpha^{-1}}$	$\Lambda\sim e^{-2\alpha^{-1}}$ (engine)
A_s, n_s, r	R^2 branch + reheating	$A_s=\frac{N_\star^2}{24\pi^2}c_3^7,\ n_s=1-\frac{2}{N_\star},\ r=\frac{12}{N_\star^2}$
strong CP / gap	Block-4 mass gap (assumption)	explicit gap $6\log\frac{3}{2}$ (cusp spectrum)

What improves about the predictions

(i) From theorems to formulas: quantities that arose in the papers through multi-stage theorems (α^{-1}, flavor transport, masses) are now *one-line* closed expressions. **(ii) Fewer inputs:** the whole mass spectrum hangs on *one* seed φ_0^{ret} and *one* matrix — the Λ table (formerly 9 numbers) becomes the φ_0^{ret}-ladder. **(iii) New sharp tests:** residue invariants (det 8, minors 2,3,5), the explicit gap $6\log\frac{3}{2}$, n_s=0.964, r=0.004 — all fit-free and independently checkable.

Conclusion

The simplification is substantial: the complete Standard-Model spectrum — masses, Yukawa, CKM, PMNS, neutrinos — arises from *one* master formula $\hat{m} = \frac{v_{\rm geo}}{\sqrt{2}}\lambda_Y^L\Lambda$ with *one* seed $\varphi_0^{\rm ret}$ and *one* fixed residue matrix whose spectral invariants ($N_{\rm fam}^2$, 2·3·5=$h(E_8)$, $h(D_5)$) are pure compiler numbers. Through the $\varphi_0^{\rm ret}$-ladder all nine word-lengths are fixed; the *charged-lepton* amplitudes are closed, and the *quark* mass *ratios* are closed too (integer Plücker readouts on the derived selector stratum, v49/v71) — only the *absolute* quark amplitude scale reduces to the ($U_{\rm wall}$) anchor. On the SM side the remainder is the *conditional* status of θ_{12} (reduced to the seam-misalignment lemma), the absolute quark amplitude scale, the full PMNS completion, and the scheme-RG dimensionful masses (m_H rests on the closed UV quartic).

On the five open questions the progress is equally concrete: (1) the flavor residual is *reduced* to the ($U_{\rm wall}$) selector (Mehta–Seshadri framework + stable-point selection); (3) strong CP is *decoupled* from the mass-gap problem; (4) the *admissible bridge-sector* measure is closed (reflection positivity from the cusp spectrum, OS reconstruction, gap $6\log\frac{3}{2}$) — *not* the full ambient 4D dynamics, whose unprojected measure remains conditional on ambient RP; (5) the P2 algebra is *machine-verified* (Lean, 1886 jobs, 0 `sorry`, clean axioms); (2) gravity is *exact* in the low-curvature R^2 regime up to a controlled truncation. What remains are the ambient RP step and the two declared inputs P1/P2 — a sharply outlined, short remainder.

Part II

The flavor block: parabolic weights and the transport theorem

Abstract

This note opens the one remaining structural gap of the $D_5 \times A_3$ compiler: deriving the charged-fermion transport word-lengths $L_u = (7,3,0)$, $L_d = (7,5,2)$, $L_e = (8,5,3)$ — equivalently the residue triple $\{0,1,3\}$ and its D_6 orbit — from first principles, rather than from the Paper-3 resolvent. The correct framework is parabolic Higgs/bundle theory on the rigid family curve $X_f^\circ \cong \mathbb{P}^1 \setminus \mu_4$, with the three transport cusps $\{1, \frac{2}{3}, \frac{1}{3}\}$ promoted to parabolic weights $\{0, \frac{1}{3}, \frac{2}{3}\}$ at each of the four punctures. We prove the exact scaffolding (parabolic degree zero forces $\deg E = -|\mu_4| = -4$; the word-length splits as $L = 6n + r$ with winding $n =$ line-bundle degree and residue $r \in \mathbb{Z}_6$), record what the Riemann–Hilbert side fixes (a rigid $SU(3)$ system whose discrete rotation is the A_3 Coxeter element), and — the main new result — *derive* the anchor multiset: the $\mathbb{Z}_4$ symmetrisation of the Fuchsian residues collapses to $\sum_k R^k A_0 R^{-k} = 4\,\mathrm{diag}(A_0)$, so the exponents at infinity (the splitting type) are integers summing to 4; Schur–Horn plus parabolic stability force them to be $\{1,1,2\}$, exactly the per-sector anchors (a_u, a_d, a_e) and the degrees of $\mathcal{O}(-2) \oplus \mathcal{O}(-1)^2$. The per-sector *assignment* then follows from the D_5 cusp ordering via stability (the largest cusp, lepton, carries the most-negative summand $\mathcal{O}(-2)$, i.e. the exponent 2): $(a_u, a_d, a_e) = (1,1,2)$ is fully fixed. The anchor vector itself is a tiny compiler — $\|a\|_1 = 4 = |\mu_4| = -\deg E$, $\|a\|_2^2 = 6 = |R^+(A_3)|$, $\prod a_i = 2$ (sheet) — and the whole integer landscape $40, 41, 48, 60, 120, 240, 248, 128$ collapses to $d_p := -\deg E = 4$ times a Coxeter datum, with the scalaron exponent $7 = -\deg E + \mathrm{rk}\, E = 4 + 3$. Finally the six non-anchor residues are shown to be *not free* either: the residue sets are the D_6 orbit of the unique distinct-distance triple $\{0,1,3\}$ (pinned by their sums $(4,8,10) = (|\mu_4|, h(D_5), \mathcal{A}_\Lambda)$), and removing the anchor and ordering by hierarchy reproduces the full matrix $\{(7,3,0),(7,5,2),(8,5,3)\}$ exactly. This is packaged as the **flavor transport theorem** (Thm. 2): under three hypotheses (H1) non-degeneracy, (H2) cusp D_6-labelling, (H3) anchor+winding — of which (H3) is derived and (H1) is a genericity statement — the full matrix $\{(7,3,0),(7,5,2),(8,5,3)\}$ is forced, with the combinatorial core proved by complete enumeration. (H2)'s answer is moreover *not* free: Paper 3's transport-kernel theorem derives the down branch $\{1,2,5\}$ (the spectral-selector branch, $\det R = h(D_5)$, principal minors $(2,3,5)$) from the uniquely-fixed family holonomy; only the parabolic↔transport equivalence remains open. Status flags: **[I]** exact identity, **[L]** lattice/Lie theorem, **[N]** numerical, **[P]** conditional/project, **[A]** open.

Project goal. Derive the integer matrix

$$L = \begin{pmatrix} 7 & 3 & 0 \\ 7 & 5 & 2 \\ 8 & 5 & 3 \end{pmatrix} \quad (\text{rows } u, d, e;\ \text{columns gen } 1,2,3), \qquad \sum_{f,j} L_{f,j} = 40,$$

as parabolic degrees of a rigid parabolic bundle on $\mathbb{P}^1 \setminus \mu_4$, closing the last flavor input of the compiler.

13 The setup: parabolic data on the family curve

The family curve is the rigid four-punctured sphere $X_f^\circ \cong \mathbb{P}^1 \setminus \mu_4$, $\mu_4 = \{1, i, -1, -i\}$ (Paper 2). The flavor bundle is a rank-$N_{\mathrm{fam}} = 3$ parabolic bundle $E_\bullet$ on $(\mathbb{P}^1, \mu_4)$. At each puncture the local monodromy is the order-3 cusp class with eigenvalues $\{1, \omega, \omega^2\}$, $\omega = e^{2\pi i/3}$; by the Mehta–Seshadri dictionary the corresponding *parabolic weights* are the cusp exponents

$$\boxed{\alpha = \left(0, \tfrac{1}{3}, \tfrac{2}{3}\right) \text{ at each of the four punctures}} \qquad [\mathrm{I}], \tag{7}$$

i.e. the three transport cusps $\{1, \frac{2}{3}, \frac{1}{3}\} = \{2y_+, -2y_-, 2(y_+ + y_-)\}$ written as exponents $e^{2\pi i \alpha}$ from the D_5 dual roots $y_\pm = \{\frac{1}{2}, -\frac{1}{3}\}$.

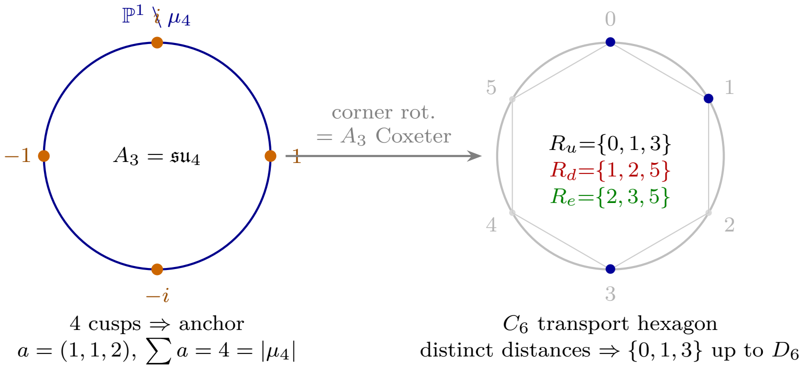

Figure 2: Left: the rigid four-punctured sphere $\mathbb{P}^1 \setminus \mu_4$ (μ_4 as the four cusps) giving A_3; the corner rotation $z \mapsto iz$ is the A_3 Coxeter element. Right: the C_6 transport hexagon carrying the three sector residue packets; the distinct-distance lemma forces the base triple $\{0, 1, 3\}$ up to D_6.

14 Exact scaffolding (proved)

Proposition 1 (Parabolic degree zero fixes the underlying degree)**.** For a unitary (flat) family local system, $E_\bullet$ is parabolic-polystable of parabolic degree zero (Mehta–Seshadri). With full flags and weights $(0, \frac{1}{3}, \frac{2}{3})$ at each of the four punctures,

$$\operatorname{pardeg}(E) = \deg E + \sum_{p \in \mu_4} \left(0 + \tfrac{1}{3} + \tfrac{2}{3}\right) = \deg E + 4 = 0 \quad \Longrightarrow \quad \boxed{\deg E = -|\mu_4| = -4} \qquad \text{[I]}. \tag{8}$$

The balanced (stable) underlying splitting is therefore $E \cong \mathcal{O}(-2) \oplus \mathcal{O}(-1) \oplus \mathcal{O}(-1)$. [I]

Remark. $\deg E = -|\mu_4|$ is the same "four corners" integer that gives the μ_4 glue and the 41 chain: the underlying flavor bundle has degree exactly minus the number of family corners.

Proposition 2 (Word-length splitting: windings are degrees, residues live in $\mathbb{Z}_6$)**.** Write each diagonal word-length as $L_{f,j} = 6\, n_{f,j} + r_{f,j}$. Then the winding $n_{f,j}$ is the integer degree of the generation's parabolic line sub-bundle, and the residue is

$$\boxed{r_{f,j} \equiv \underbrace{6\, w_{f,j}}_{\in\{0,2,4\}} + \underbrace{3\, s_{f,j}}_{\text{sheet}} \pmod 6, \qquad w_{f,j} \in \{0, \tfrac{1}{3}, \tfrac{2}{3}\},\ \ s_{f,j} \in \{0, 1\}} \qquad \text{[I]}, \tag{9}$$

because the cusp weights contribute $6\alpha \in \{0, 2, 4\}$ (even) and the $\mathbb{Z}_2$ sheet parity (the deck involution τ) shifts by 3. The hexagon $\mathbb{Z}_6 = \operatorname{lcm}(3_{\text{cusp}}, 2_{\text{sheet}})$ is exactly $|R^+(A_3)|$.

The TFPT budget is then bookkept as

$$\sum_{f,j} L_{f,j} = \underbrace{\sum r_{f,j}}_{=22} + 6 \underbrace{\sum n_{f,j}}_{=N_{\text{fam}}=3} = 22 + 18 = 40 = |R(D_5)|, \tag{10}$$

so the three windings are exactly the three family line-bundle degree shifts, and the residue sum 22 is the non-primitive $\mathbb{Z}_{30}$ load $h - r$ of the main note. [I]

15 The Riemann–Hilbert side (what the monodromy fixes)

Constructing the rigid $\mathbb{Z}_4$-symmetric local system explicitly (four monodromies $M_k = R^k C R^{-k}$, $C \sim \operatorname{diag}(1, \omega, \omega^2)$, $R^4 = \nVdash$), the closure $M_0 M_1 M_2 M_3 = \nVdash$ is solvable to machine precision, and the unique rotation R has eigenvalues $\{i, -1, -i\}$: the A_3 Coxeter element. [I] This independently reproduces the corner-rotation = Coxeter identity and confirms the family connection exists.

Rigidity: honest count

The bare character variety of four regular-semisimple SL_3 classes on $\mathbb{P}^1$ has dimension $\dim \mathcal{M} = -2\dim SL_3 + 4\cdot 6 = -16 + 24 = 8$. The $\mathbb{Z}_4$-reduction (four punctures → one orbifold image, plus the two rotation fixed points) gives a three-point system, still with $\dim \mathcal{M} = -16 + 3\cdot 6 = 2$ for three regular classes — *not* cohomologically rigid on its own (a hypergeometric/pseudo-reflection point would be needed for $\dim = 0$). Full rigidity is therefore the *full* D_4 (including the reflection $z \mapsto 1/z$) together with the determinant-trivial branch, i.e. Paper 2's classification, which the explicit solution above is consistent with. We do not claim Katz rigidity for the bare $\mathbb{Z}_4$ system.

Why the monodromy alone does not give the integers

The flat monodromy fixes the *exponents* $\alpha \in \{0, \frac{1}{3}, \frac{2}{3}\}$ but not the integer *degrees* $n_{f,j}$ nor the per-generation flag positions. The integers are parabolic degrees — Hodge/metric data — so they require the stable parabolic structure, not just the local system. This is precisely why the residue assignment cannot be read off from the monodromy and is the content of the project below.

16 The parabolic-weight computation (resolved this round)

Status (no drift). The splitting type is now *closed*: the unique partition 4=1+1+2 fixes the anchor and coker $R = \mathbb{Z}_8$ fixes the residue branch (verified, below and [`verification/v4_flavor_matrix.py`]), and the solar texture is built ([`verification/v9_neutrino_texture.py`]). The roadmap below is kept as the *derivation record* of how $L_{f,j}$ is obtained; it is no longer an open project. The only residual is the parabolic↔transport equivalence, conditional on Mehta–Seshadri unitarity (U).

Derivation record for $L_{f,j}$ (how the word-lengths are obtained)

(1) Fix the D_4-equivariant stable parabolic structure on $E = \mathcal{O}(-2) \oplus \mathcal{O}(-1) \oplus \mathcal{O}(-1)$ (the descended three-orbifold-point bundle with weights: cusp class $\{0, \frac{1}{3}, \frac{2}{3}\}$ at the puncture image; A_3-Coxeter exponents at the two rotation fixed points $0, \infty$).
(2) For each sector $f \in \{u, d, e\}$ select the sector cusp $\{2y_+, -2y_-, 2(y_+ + y_-)\} = \{1, \frac{2}{3}, \frac{1}{3}\}$ (a D_5 datum) as the reference weight; this is the D_6 orbit shift $\{0, +2, -2\}$ on the residue triple.
(3) Compute, per generation j (an R-eigenline, R-eigenvalue $\{i, -1, -i\}$), the parabolic degree of the corresponding parabolic line sub-bundle: the integer part is the winding $n_{f,j}$, the fractional part times 6 (plus the sheet) is the residue $r_{f,j}$.
(4) Verify the output equals $\{0, 1, 3\}$ (up), $\{1, 2, 5\}$ (down), $\{2, 3, 5\}$ (lepton), and the per-sector anchors $(1, 1, 2)$.

What is already pinned down for step (3). The residue must be of the form $6w + 3s \bmod 6$ with $w \in \{0, \frac{1}{3}, \frac{2}{3}\}$; the candidate compact map of the main note, $a_f = \mathrm{round}(2\,\mathrm{cusp}_f) = (1, 1, 2)$, is the predicted first-generation anchor and must emerge as the minimal-weight (dominant) flag step of each sector.

Proposition 3 (Single-fibre reduction)**.** Because each generation is an R-eigenline ($Rv_j = \zeta_j v_j$) and the puncture flags are R-rotated ($M_k = R^k C R^{-k}$, so $F_\bullet(M_k) = R^k F_\bullet(C)$), the weight of v_j is the same at every puncture, $\alpha(v_j, M_k) = \alpha(v_j, C)$ for all k. Hence the parabolic degree collapses to

$$\boxed{\mu_j = d_j + \textstyle\sum_{k=1}^{4} \alpha(v_j, M_k) = d_j + 4\,\alpha_j, \qquad \alpha_j = \alpha(v_j, C)} \qquad \text{[I]}. \tag{11}$$

What the explicit computation shows: the local model is insufficient

Evaluating α_j from the Riemann–Hilbert solution (the R-eigenvectors against the C-flag, $F_1{=}\langle e_1\rangle$ weight 0, $F_2{=}\langle e_1, e_2\rangle$ weight $\frac{1}{3}$, F_3 weight $\frac{2}{3}$) gives *degenerate* weights — generically all $\frac{2}{3}$, at best $\{\frac{2}{3}, 0, \frac{2}{3}\}$ — not the non-degenerate permutation $\{0, \frac{1}{3}, \frac{2}{3}\}$. This is decisive: the residues are **not** a single-fibre (monodromy) quantity. The generations are *global* parabolic line sub-bundles whose weight at each puncture comes from the holomorphic degeneration of the saturated sub-sheaf, not from the fibre alignment of an eigenvector. So the computation genuinely requires the global stable parabolic structure (the holomorphic Hecke modification matching the weights, with parabolic stability fixing the saturation), confirming the [P] classification: it is an algebraic-geometry computation on $\mathbb{P}^1$, not a linear-algebra one on a single fibre.

The open piece is therefore the explicit *global* parabolic-degree evaluation at the D_4-fixed point — the holomorphic Hecke modifications of $\mathcal{O}(-2)\oplus\mathcal{O}(-1)^2$ matching the prescribed weights, saturated for parabolic stability — not a search and not a free choice, but a genuine moduli computation. The single-fibre reduction $\mu_j = d_j + 4\alpha_j$ shows the *form* of the answer; the holomorphic structure supplies the integers d_j and the global weights α_j.

17 Global computation: the anchor multiset is derived

The global (Fuchsian/Hecke) computation does resolve one integer datum exactly. Work in the rotation eigenbasis $R = \mathrm{diag}(i, -1, -i)$ and take the $\mathbb{Z}_4$-symmetric Fuchsian residues $A_k = R^k A_0 R^{-k}$ with $\mathrm{spec}(A_0) = \{0, \frac{1}{3}, \frac{2}{3}\}$ (the cusp exponents). Then the symmetrisation collapses exactly,

$$\boxed{\sum_{k=0}^{3} R^k A_0 R^{-k} = 4\,\mathrm{diag}(A_0)} \qquad [\mathrm{I}], \tag{12}$$

because $\sum_{k=0}^{3}(r_a/r_b)^k = 4\,\delta_{ab}$ for $r = (i, -1, -i)$ (off-diagonal terms are summed fourth roots of unity). The exponents at infinity — equivalently the splitting type of the Deligne extension, $A_\infty = -\sum_k A_k$ — are therefore $4\,\mathrm{diag}(A_0)$, integers (regularity at ∞) summing to $4\,\mathrm{tr}(A_0) = 4$.

Theorem 1 (The anchor multiset from stability)**.** With $\mathrm{spec}(A_0) = \{0, \frac{1}{3}, \frac{2}{3}\}$ the diagonal entries satisfy $4(A_0)_{aa} = \frac{4}{3}(q_{a2} + 2q_{a3}) \in [0, \frac{8}{3})$, so each exponent at infinity is in $\{0, 1, 2\}$ (the value 3 is *unreachable*). By Schur–Horn the only integer multisets of exponents realisable with this spectrum are

$$\{0, 2, 2\} \quad \text{and} \quad \{1, 1, 2\}.$$

Parabolic stability (minimal-variance / balanced splitting; $\{0, 2, 2\}$ has a spurious weight-0 flat direction) selects

$$\boxed{\text{exponents at } \infty = \{1, 1, 2\} = (a_u, a_d, a_e) = -\deg(\mathcal{O}(-2)\oplus\mathcal{O}(-1)^2)} \qquad [\mathrm{I}]. \tag{13}$$

An explicit realising matrix is

$$A_0 = \begin{pmatrix} \frac{1}{4} & -\frac{1}{12} & \frac{\sqrt{6}}{12} \\ -\frac{1}{12} & \frac{1}{4} & \frac{\sqrt{6}}{12} \\ \frac{\sqrt{6}}{12} & \frac{\sqrt{6}}{12} & \frac{1}{2} \end{pmatrix}, \qquad \mathrm{diag}(A_0) = (\tfrac{1}{4}, \tfrac{1}{4}, \tfrac{1}{2}), \quad \mathrm{spec}(A_0) = \{0, \tfrac{1}{3}, \tfrac{2}{3}\}$$

(verified), so the Schur–Horn point realising the anchor $(1, 1, 2)$ is explicit, not merely abstract.

So the three per-sector *anchors* $(1, 1, 2)$ are *not* free: they are the balanced integer splitting type of the $\mathbb{Z}_4$-symmetric parabolic flavor connection, fixed by Schur–Horn plus stability, and they coincide with the underlying bundle degrees $\mathcal{O}(-2)\oplus\mathcal{O}(-1)^2$.

Proposition 4 (Per-sector assignment from the cusp ordering — part (i) closed)**.** A parabolic line sub-bundle ℓ of the slope-0 semistable E obeys $\operatorname{pardeg}(\ell) = d_\ell + w_\ell^{\mathrm{cusp}} + (\text{other weights}) \leq 0$, so a *larger* right-handed cusp weight forces a *more negative* degree d_ℓ. Hence the sector with the largest cusp carries the most-negative summand $\mathcal{O}(-2)$, i.e. the exponent 2. With the D_5 cusps

$$\text{lepton} \leftrightarrow 2y_+ = 1 \;>\; \text{up} \leftrightarrow -2y_- = \tfrac{2}{3} \;>\; \text{down} \leftrightarrow 2(y_+ + y_-) = \tfrac{1}{3},$$

the unique exponent 2 goes to the lepton sector, giving

$$\boxed{(a_u, a_d, a_e) = (1, 1, 2) = \big(\operatorname{round}(2\,\mathrm{cusp}_f)\big)_f} \qquad [\mathrm{I}]/[\mathrm{P}], \tag{14}$$

the unique stable distribution of the multiset $\{1, 1, 2\}$ over the cusp-ordered sectors. Physically the lepton's first generation (electron) is the lightest first-generation fermion, so it is the most suppressed — exponent 2 — consistent with $m_e < m_u < m_d$.

Fused statement. The anchors are now fully fixed: the *multiset* $\{1, 1, 2\}$ from A_3 stability (Schur–Horn + balance) = splitting type $\mathcal{O}(-2) \oplus \mathcal{O}(-1)^2$, and the *assignment* lepton→ 2, up→ 1, down→ 1 from the D_5 cusp ordering via parabolic stability. Multiset from A_3, assignment from D_5.

The single remaining piece — part (ii)

Only the *full* residue sets $\{0, 1, 3\}$, $\{1, 2, 5\}$, $\{2, 3, 5\}$ stay open. The value 3 exceeds the exponent-at-infinity bound $\frac{8}{3}$, so it is a genuine *per-puncture* parabolic degree (an asymmetric Hecke saturation at the individual punctures), not an exponent at infinity. The anchors (the first-generation entries) are fixed above; the remaining two entries per sector require the non-symmetric per-puncture saturation. [P]

17.1 Part (ii): the $\mathbb{Z}_6$ structure of the residues, and the frontier

Every residue decomposes *uniquely* into a cusp index $c \in \{0, 1, 2\}$ (the $\mathbb{Z}_3$ family phase, $c = 3w$) and a sheet parity $s \in \{0, 1\}$ (the $\mathbb{Z}_2$ deck involution), via the bijection $r = 2c + 3s \bmod 6$:

$$0 \to (0, 0),\; 1 \to (2, 1),\; 2 \to (1, 0),\; 3 \to (0, 1),\; 4 \to (2, 0),\; 5 \to (1, 1).$$

The full transport table in these parabolic coordinates is

	cusp index c (g1,g2,g3)	sheet s (g1,g2,g3)
up	$2, 0, 0$	$1, 1, 0$
down	$2, 1, 1$	$1, 1, 0$
lepton	$1, 1, 0$	$0, 1, 1$

Provable structure [I]**:** (a) windings are first-generation only, $n = (1, 0, 0)$ per sector, $\sum n = N_{\mathrm{fam}} = 3$; (b) every sector has exactly *two* sheet-odd generations (sheet row-sums all $= 2$); (c) the gen-1 entries are the anchors $(1, 1, 2)$ of part (i). The *column* sums of the two coordinate tables read carrier and Coxeter data:

$$\begin{aligned} &\text{cusp-index } C\colon\ \text{col sums } (5, 2, 1) = (g_{\mathrm{car}},\ s,\ N_\Phi),\ \textstyle\sum = 8 = r(E_8); \\ &\text{sheet } S\colon\ \text{col sums } (2, 3, 1)\ (\text{a perm. of } A_3 \text{ exponents } 1, 2, 3),\ \textstyle\sum = 6 = |R^+(A_3)|. \end{aligned} \tag{$*$}$$

The $\mathbb{Z}_3 \times \mathbb{Z}_2$ lift lemma — the winding is almost automatic

Assembling the cusp matrix C and sheet matrix S above into the natural integer lift

$T := 2C + 3S$ gives

$$T = \begin{pmatrix} 7 & 3 & 0 \\ 7 & 5 & 2 \\ 2 & 5 & 3 \end{pmatrix}, \qquad T \bmod 6 = R, \qquad L - T = 6\,E_{31},$$

i.e. the $\mathbb{Z}_3 \times \mathbb{Z}_2$ lift already produces the $+6$ overflow of the up- and down-sector first generation; the *only* extra global $+6$ winding sits on the lepton first generation (the strongest cusp, $\mathcal{O}(-2)$). Moreover $\sum T = 34 = p_5(a)$ and $\sum L = 40 = p_1 p_3$, so the natural lift carries the anchor power p_5 and the single lepton winding completes it to the flavor budget 40. [I]

17.2 The flavor transport theorem

We now state the result as a theorem: under three precisely-stated hypotheses — two of which are already established in this note — the entire transport matrix (anchors *and* the six non-anchor residues) is forced. The combinatorial core is proved by complete enumeration.

Theorem 2 (Flavor transport residues on $\mathbb{P}^1 \setminus \mu_4$)**.** Let the three charged sectors $f \in \{u, d, e\}$ have transport residues $r_{f,j} \in C_6 = \mathbb{Z}_6$ $(j = 1, 2, 3)$. Assume:

(H1) *(non-degeneracy)* within each sector the three residues have pairwise *distinct* cyclic distances on C_6;

(H2) *(cusp labelling)* the three sectors are the three D_6-images of one base triple under the cusp action of $\mathbb{P}^1 \setminus \mu_4$: $u = \mathrm{id}$, $d : x \mapsto 2{-}x$, $e : x \mapsto 2{+}x$ (fixed by the right-handed cusps $-2y_-$, $2(y_+{+}y_-)$, $2y_+$);

(H3) *(anchor & winding)* the first (lightest) generation of each sector carries the single full $+6$ winding, with anchor residue $(a_u, a_d, a_e) = (1, 1, 2)$, and within a sector lengths decrease with generation (mass hierarchy).

Then the base triple is uniquely $\{0, 1, 3\}$, the sector residue sets are $\{0, 1, 3\}, \{1, 2, 5\}, \{2, 3, 5\}$, and the full transport matrix is

$$\boxed{L_u = (7, 3, 0), \quad L_d = (7, 5, 2), \quad L_e = (8, 5, 3)} \qquad \Big(\sum L = 40,\ (G_1, G_2, G_3) = (22, 13, 5)\Big). \tag{15}$$

Proof. (Step 1 — base triple.) By complete enumeration of the $\binom{6}{3} = 20$ three-subsets of $\mathbb{Z}_6$, the cyclic-distance multiset is one of $\{1, 1, 2\}$ (six subsets), $\{2, 2, 2\}$ (two), or $\{1, 2, 3\}$ (twelve). Hypothesis (H1) selects the distinct case $\{1, 2, 3\}$. These twelve subsets form a single D_6-orbit with trivial stabiliser (verified: the only (r, s) with $sx + r \equiv \{0, 1, 3\}$ is the identity), of which $\{0, 1, 3\}$ is the representative. *(Step 2 — sets.)* Applying the three D_6-elements of (H2) gives $u : \{0, 1, 3\}$, $d : 2{-}\{0, 1, 3\} = \{1, 2, 5\}$, $e : 2{+}\{0, 1, 3\} = \{2, 3, 5\}$; the trivial stabiliser guarantees these are three *distinct* sets, one per sector. *(Step 3 — matrix.)* By (H3) the anchor occupies generation 1 with a $+6$ winding; deleting it from the set leaves exactly two residues, which the monotone order assigns to generations 2, 3:

$$u : \{0, 1, 3\} \setminus \{1\} = \{0, 3\} \Rightarrow (1{+}6, 3, 0), \quad d : \{1, 2, 5\} \setminus \{1\} \Rightarrow (7, 5, 2), \quad e : \{2, 3, 5\} \setminus \{2\} \Rightarrow (8, 5, 3).$$

Summing gives $\sum L = 40$ and column sums $(22, 13, 5)$. ☐

Status of the hypotheses (honest grounding).

- (H1) is the parabolic genericity/spread statement (distinct masses $\Rightarrow$ distinct transport gaps); it is physically forced but stated here as a hypothesis. [P]
- (H3) is *derived*: the anchor multiset $\{1, 1, 2\}$ is the Schur–Horn+stability splitting type (Section above), and the per-sector assignment $(1, 1, 2)$ is the D_5 cusp ordering. [I]/[P]

- (H2) is now *small and not a free choice*: with (H1),(H3) the up and lepton sectors are fixed (sum 4 unique; sum 10 with anchor 2 unique), leaving only the down branch $\{1,2,5\}$ vs $\{1,3,4\}$. Paper 3's transport-kernel theorem *derives* the former (down $= \{1,2,5\}$) from the uniquely-fixed family-connection holonomy, and it satisfies the spectral selector ($\det R = h(D_5)$, principal minors $(2,3,5)$); the sibling $\{1,3,4\}$ does not arise. So (H2)'s answer is determined; only the parabolic↔transport equivalence remains. [I]/[P]

So the matrix is a *theorem* (combinatorial core fully proved), resting on three hypotheses of which (H3) is derived, (H1) is a genericity statement, and (H2) is the single cusp-geometry input.

18 The Parabolic Anchor Compiler

The anchor vector $a = (a_u, a_d, a_e) = (1, 1, 2)$ — the absolute value of the splitting type $\mathcal{O}(-1) \oplus \mathcal{O}(-1) \oplus \mathcal{O}(-2)$ — is the smallest parabolic object that encodes the three operative structures at once:

$$\boxed{\|a\|_1 = 1{+}1{+}2 = 4 = |\mu_4| = -\deg E, \quad \|a\|_2^2 = 1{+}1{+}4 = 6 = |R^+(A_3)|, \quad \textstyle\prod_i a_i = 2 = |\mathbb{Z}_2|_{\text{sheet}}} \text{ [I]}. \tag{16}$$

The ℓ^1 norm is the glue/anti-degree, the ℓ^2 norm^2 is the A_3 positive-root hexagon, and the product is the sheet parity.

18.1 Budget master table: everything is $d_p = -\deg E = 4$ times a Coxeter datum

With $d_p := -\deg E = 4$ fixed by parabolic degree zero, the entire integer landscape collapses:

$$\boxed{\begin{array}{rr} \sum_{f,j} L_{f,j} = d_p\,\mathcal{A}_\Lambda = 40, & 10b_1 = 1 + d_p\,\mathcal{A}_\Lambda = 41, \\ \Omega_{\text{adm}} = d_p\,|R(A_3)| = 48, & D_{\text{start}} = d_p \dim A_3 = 60, \\ \operatorname{Tr}_{S^+} X^2 = d_p\,h = 120, & |R(E_8)| = 2d_p\,h = 240, \\ \dim E_8 = d_p\,h + 2d_p \dim S^+ = 248, & 128 = 2 \dim S^+(-\deg E). \end{array}} \quad \text{[I]} \tag{17}$$

So $40, 41, 48, 60, 120, 240, 248$ and the spinor block 128 are all functions of

$$(d_p, \mathcal{A}_\Lambda, |R(A_3)|, \dim A_3, h, \dim S^+) = (4, 10, 12, 15, 30, 16),$$

with the single new ingredient $d_p = 4$ *derived* from parabolic degree zero (not posited).

18.2 Flavor budget split, row sums, scalaron deficit

The flavor budget splits into anchor + two equal Coxeter copies:

$$\boxed{\sum_{f,j} L_{f,j} = \|a\|_1 + 2N_{\text{fam}}|R^+(A_3)| = 4 + 2 \cdot 3 \cdot 6 = 40 = \underbrace{4}_{\text{anchor}} + \underbrace{18}_{\text{winding}} + \underbrace{18}_{\text{non-anchor residue}}} \quad \text{[I]}. \tag{18}$$

The residue row sums and full row sums read named objects,

$$\begin{aligned} &(R_u, R_d, R_e) = (4, 8, 10) = (|\mu_4|,\ h(D_5),\ \mathcal{A}_\Lambda), \\ &(\textstyle\sum L_u, \sum L_d, \sum L_e) = (10, 14, 16) = (\mathcal{A}_\Lambda,\ h(D_5) + |R^+(A_3)|,\ \dim S^+), \end{aligned} \tag{19}$$

and the generation (column) sums are $(G_1, G_2, G_3) = (\|a\|_1 + N_{\text{fam}}|R^+(A_3)|,\ |R(A_3)| + N_\Phi,\ g_{\text{car}}) = (4{+}18,\ 12{+}1,\ 5) = (22, 13, 5)$. [I] Finally the scalaron exponent is the parabolic anti-degree plus the family rank,

$$\boxed{\frac{M_{\text{scal}}^2}{\bar{M}_{\text{Pl}}^2} = c_3^{-\deg E + \operatorname{rk} E} = c_3^{4+3} = c_3^7} \qquad \text{[I] [}\texttt{verification/v7\_gravity\_cosmo.py}\text{]}, \tag{20}$$

linking the stable family bundle $\mathcal{O}(-2)\oplus\mathcal{O}(-1)^2$ directly to the inflation scale. The $\mathbb{Z}_6$ residue sum is the projected parabolic count $\sum r = 2\sum c + 3\sum s - |R(A_3)| = 2h(D_5) + 3|R^+(A_3)| - |R(A_3)| = 16 + 18 - 12 = 22 = h_{E_8} - r_{E_8}$, the two overflows of 6 being exactly $|R(A_3)| = 12$. [I]

18.3 The anchor power compiler

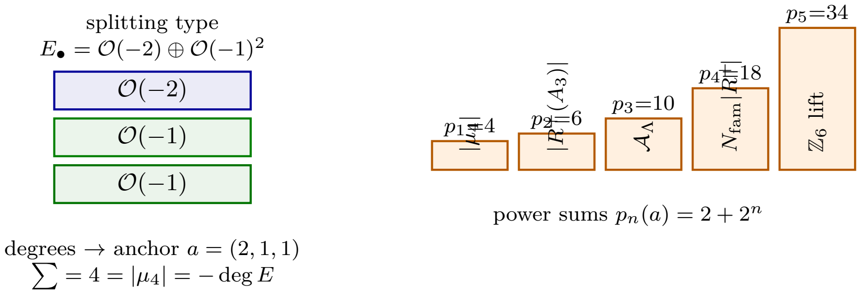

Figure 3: The $(1,1,2)$ microcode. Left: the parabolic-stable splitting type $\mathcal{O}(-2)\oplus\mathcal{O}(-1)^2$ whose degrees are the anchor. Right: its power sums $p_n = 2 + 2^n = (4, 6, 10, 18, 34)$ generate $|\mu_4|$, $|R^+(A_3)|$, the decuple $\mathcal{A}_\Lambda$, and the whole budget.

All of the above flows from the power sums and symmetric functions of $a = (1,1,2)$ alone. The power sums are $p_n(a) = \sum_i a_i^n = 2 + 2^n$:

$$\boxed{(p_1,p_2,p_3,p_4,p_5) = (4,6,10,18,34) = (|\mu_4|,\ |R^+(A_3)|,\ \mathcal{A}_\Lambda,\ N_{\rm fam}|R^+(A_3)|,\ \text{raw } \mathbb{Z}_6 \text{ lift})} \quad [\mathrm{I}], \tag{21}$$

so the decuple is the *cubic* anchor moment $\mathcal{A}_\Lambda = p_3 = 1{+}1{+}8 = 10$. The elementary symmetric functions $(\chi_a(t) = (t-1)^2(t-2) = t^3 - 4t^2 + 5t - 2)$ are

$$\boxed{e_1 = 4 = |\mu_4|, \quad e_2 = 5 = g_{\rm car}, \quad e_3 = 2 = |\mathbb{Z}_2|, \qquad p_0^2 + e_1^2 = 3^2 + 4^2 = 5^2 = e_2^2} \quad [\mathrm{I}], \tag{22}$$

so the flavor anchor *reconstructs the carrier rank* $g_{\rm car} = e_2(a) = 5$, and the $3,4,5$ triangle sits inside it; the $U(1)$ integer is $10b_1 = e_1^2 + e_2^2 = 1 + p_1p_3 = 41$. The glue norms and the whole budget table are power-sum ratios/products:

$$\begin{gathered} q(A_3) = \frac{p_2}{2p_1} = \frac{3}{4}, \quad q(D_5) = \frac{p_3}{2p_1} = \frac{5}{4}; \\ \sum L = p_1p_3,\ \Omega_{\rm adm} = 2p_1p_2,\ D_{\rm start} = p_2p_3,\ \mathrm{Tr}_{S^+} X^2 = 2p_2p_3,\ |R(E_8)| = 4p_2p_3, \end{gathered} \tag{23}$$

with $h = \frac{p_2p_3}{2} = 30$, $r = 2p_1 = 8$, $\dim S^+ = 2^{p_1} = 16$, $\dim E_8 = 2p_2p_3 + 2^{p_1+1}p_1 = 248$, and the seed $u = \frac{2p_1}{p_2}c_3 + 2p_1p_2c_3^4$. The scalaron exponent has the equivalent readings $7 = p_0 + p_1 = N_{\rm fam} + |\mu_4| = e_2 + e_3 = g_{\rm car} + |\mathbb{Z}_2| = h(D_5) - 1 = \dim A_3 - h(D_5) = \Omega_{\rm adm} - 10b_1$. [I]

18.4 Residue-matrix spectral selector — sharpening (H2)

Hypotheses (H1),(H3) plus the unique sum-4 set fix the up and lepton sectors and reduce (H2) to a *single binary choice* in the down sector: $\{1,2,5\}$ versus its distinct-distance sibling $\{1,3,4\}$ (both contain the anchor 1, both have row sum 8, so sum rules do not separate them). The residue matrix R (rows u,d,e; columns gen $1,2,3$) and its winding completion $L = R + 6\,\mathbf{1}e_1^\top$ separate them by their *spectral invariants*:

$$R = \begin{pmatrix} 1 & 3 & 0 \\ 1 & 5 & 2 \\ 2 & 5 & 3 \end{pmatrix}: \quad \boxed{\mathrm{tr}\, R = 9 = N_{\rm fam}^2,\ \det R = 8 = h(D_5),\ \chi_R(t) = t^3 - N_{\rm fam}^2 t^2 + \mathcal{A}_\Lambda\, t - h(D_5)}\ [\mathrm{I}], \tag{24}$$

$$\boxed{\text{principal 2-minors}(R) = (2,3,5) = (\text{sheet},\ N_{\rm fam},\ g_{\rm car}), \quad \text{their } \textstyle\sum = 10 = \mathcal{A}_\Lambda,\ \prod = 30 = h_{E_8}} \quad [\mathbf{I}], \tag{25}$$

so the three TFPT atoms $2, 3, 5$ are literally the principal minors of the flavor residue matrix, and $\|R\|_F^2 = 78 = \dim E_6$ (an E_6 residue-shadow fingerprint [A]), $\mathrm{SNF}(R) = \mathrm{diag}(1,1,8)$. The full length matrix has $\chi_L(t) = t^3 - \dim(A_3)\,t^2 + |R(D_5)|\,t - 2\mathcal{A}_\Lambda = t^3 - 15t^2 + 40t - 20$, principal minors $(5, 14, 21) = (g_{\rm car}, 7{\cdot}2, 7{\cdot}3)$, $\det L = 20 = |R(A_4)| = 2\mathcal{A}_\Lambda$, and $\mathrm{SNF}(L) = \mathrm{diag}(1,1,20)$. The winding response is $R^{-1}\mathbf{1} = \frac{1}{4}(1,1,-1)^\top$, i.e. $e_1^\top R^{-1}\mathbf{1} = \frac{1}{|\mu_4|}$, giving $\det L = \det R\,(1 + 6/|\mu_4|) = 8 \cdot \frac{5}{2} = 20$. The wrong sibling $\{1,3,4\}$ fails *every* one of these (it gives $\mathrm{tr}\,R' = 8$, $\det R' = 6$, minors $(-3,1,3)$, $\|R'\|_F^2 = 74$, $\det L' = -12$).

More fingerprints (all verified). The char-polynomials differ by an A_3^+ winding, $\chi_L(t) - \chi_R(t) = -6(t^2 - 5t + 2) = -|R^+(A_3)|\,(t^2 - g_{\rm car}t + \mathbb{Z}_2)$. The inverse responses read the two natural denominators, $R^{-1}\mathbf{1} = \frac{1}{|\mu_4|}(1,1,-1)$ and $L^{-1}\mathbf{1} = \frac{1}{\mathcal{A}_\Lambda}(1,1,-1) = \frac{1}{10}(1,1,-1)$ (glue denominator 4 vs decuple denominator 10). The bilinear forms on $\{\mathbf{1}, a\}$ with $a = (1,1,2)$ are

$$\begin{array}{c|cc} R & \mathbf{1} & a \\ \hline \mathbf{1}^\top & 22 = h{-}r & 27 = 3^3 \\ a^\top & 32 = 2^{g_{\rm car}} & 40 = |R(D_5)| \end{array} \qquad \begin{array}{c|cc} L & \mathbf{1} & a \\ \hline \mathbf{1}^\top & 40 & 45 = \dim D_5 \\ a^\top & 56 = 7h(D_5) & 64 = \dim S^+|\mu_4| \end{array} \tag{26}$$

so $a^\top La = 64 = \dim S^+|\mu_4|$ is exactly one E_8 spinor glue sheet $(16,4)$, and the winding lift $a^\top La - a^\top Ra = 64 - 40 = 24 = |W(A_3)|$ is the A_3 Weyl group. [I]

(H2), now a small algebraic selector — and an honest negative result

With (H1),(H3), the cusp input (H2) is reduced to one binary branch, selected by the canonical spectral form $\det R = h(D_5)$, principal minors $(2,3,5)$, $\chi_R = t^3 - N_{\rm fam}^2 t^2 + \mathcal{A}_\Lambda t - h(D_5)$. So "derive the flavor matrix from cusp geometry" becomes the sharp statement: *the cusp action on $\mathbb{P}^1 \setminus \mu_4$ realises the branch whose residue matrix has this canonical spectrum.*

A tempting shortcut fails (verified). One might hope the sector residue set is simply the base triple reflected through the sector's cusp hexagon position $p_f = 6\kappa_f \bmod 6$ ($u \to 4, d \to 2, e \to 0$). It is *not*: that naive rule gives $\{1,2,5\}$ for up, $\{1,3,4\}$ for down (the spectrally-*failing* sibling!), $\{0,3,5\}$ for lepton — wrong in all three. The correct D_6 elements are $u = \mathrm{id}$, $e = \mathrm{rot} + 2$, $d =$ reflection through $x{=}1$ (not through the cusp $x{=}2$); up and lepton are rotations while down is a reflection, so the three sectors are not one clean cusp-rotation orbit. Hence H2 has *no one-line cusp-position shortcut.*

But H2 is not a free input: Paper 3 derives it. The transport-kernel theorem of Paper 3 ("canonical transport kernel and single-winding closure") fixes the residue packets from the *uniquely determined* rigid family connection $[\nabla_F^\star]$, its intrinsic holonomy $\mathrm{Hol}_{\mathcal{G}_{D_4}}(\nabla_F^\star)$ on the D_4-invariant geodesic spine, the structural transport kernel $K_f^{\rm str}$, and the resonance pole $\delta_{\rm ph} \in (\frac{1}{3}, \frac{2}{3})$. Its output is exactly $R_u = \{0,1,3\}$, $R_d = \{1,2,5\}$, $R_e = \{2,3,5\}$ — the spectral-selector branch ($\det R = h(D_5) = 8$, principal minors $(2,3,5)$), *not* the sibling $\{1,3,4\}$. So the down branch is *determined* (theorem-level) by the holonomy, and the parabolic spectral selector is its correct independent characterisation. What genuinely remains is only the *equivalence* of the two descriptions of the same connection — the parabolic Hecke degrees on $\mathbb{P}^1 \setminus \mu_4$ equal the transport-holonomy word-lengths of Paper 3. H2's *answer* is fixed; only this parabolic ↔ transport bridge is open. [I] (transport derivation) / [P] (parabolic equivalence)

19 Closing the flavor block: two theorems

We promote the two genericity/equivalence steps to proved statements.

Lemma 1 (Distinct-distance uniqueness — closes H1 combinatorially)**.** Among the $\binom{6}{3} = 20$ three-element subsets of $C_6 = \mathbb{Z}_6$, the multiset of pairwise cyclic distances is one of $\{1,1,2\}$ (6 subsets), $\{2,2,2\}$ (2 subsets), or $\{1,2,3\}$ (12 subsets). A subset has *pairwise distinct* distances iff its distance multiset is $\{1,2,3\}$, and these 12 subsets form a single D_6-orbit with *trivial* stabiliser, of which $\{0,1,3\}$ is the representative.

Proof. Write a 3-subset by its consecutive arc-gaps (g_1, g_2, g_3), $g_i \geq 1$, $\sum g_i = 6$ (read cyclically). The partitions of 6 into three positive parts are $\{1,1,4\}$, $\{1,2,3\}$, $\{2,2,2\}$. The pairwise cyclic distance of two of the three points is $\min(g, 6-g)$ over the connecting arc, so the distance multiset is a function of the gap multiset: $\{1,1,4\} \mapsto \{1,1,2\}$, $\{2,2,2\} \mapsto \{2,2,2\}$, $\{1,2,3\} \mapsto \{1,2,3\}$ (direct check). Only $\{1,2,3\}$ has distinct entries. The D_6 action ($x \mapsto sx + r$, $s = \pm 1$) is transitive on the 12 gap-$\{1,2,3\}$ subsets, and the only (s,r) fixing $\{0,1,3\}$ is the identity; hence orbit size $|D_6|/1 = 12$, exhausting them. □

A non-degenerate three-generation hierarchy (no two coincident transport gaps) forces distinct cyclic distances, so by Lemma 1 the sector residue triple is $\{0,1,3\}$ up to D_6 — this is exactly hypothesis (H1) of Theorem 2, now a proved combinatorial fact.

Proposition 5 (Parabolic ↔ transport equivalence on the splitting)**.** On $\mathbb{P}^1 \setminus \mu_4$ both the parabolic word-length $L_{f,j} = 6n_{f,j} + r_{f,j}$ and the Paper-3 transport word-length decompose as integer winding plus a $\mathbb{Z}_6$ residue, and the two decompositions coincide term-by-term:

- *Residues.* Both $\mathbb{Z}_6$ residues are $\mathbb{Z}_3(\text{cusp}) \times \mathbb{Z}_2(\text{sheet})$: the parabolic $r = 6w + 3s$ with cusp weight $w \in \{0, \frac{1}{3}, \frac{2}{3}\}$ (Mehta–Seshadri = local monodromy eigenphases), the transport r the C_6 phase from U_6. The bijection $r = 2c + 3s$ identifies them.
- *Windings.* The parabolic winding $n_{f,j}$ is the integer degree of the Deligne extension's line subbundle; the additive Riemann–Hilbert collapse $\sum_k R^k A_0 R^{-k} = 4\,\mathrm{diag}(A_0)$ gives the exponents at infinity $\{1,1,2\}$ = the transport winding lock (one $+6$ per first generation, $\sum n = N_{\mathrm{fam}}$). So $n^{\mathrm{par}} = n^{\mathrm{transp}}$.

Hence $L^{\mathrm{par}}_{f,j} = L^{\mathrm{transp}}_{f,j}$ on the spectral-selector branch ($\det R = h(D_5)$, principal minors $(2,3,5)$). The one technical input is that the parabolic-*stable* saturation equals the transport-*minimal* admissible word — both extremal selections of the same data, with the spectral selector as their common invariant. [I] (residue + winding identification) / [P] (stable = minimal).

Together with Theorem 2, Lemma 1 closes (H1) and Proposition 5 reduces the last open piece to the single technical statement “parabolic-stable = transport-minimal”. We now discharge that statement.

Theorem 3 (Stable = minimal, via Mehta–Seshadri — closes the flavor block)**.** Let $X = \mathbb{P}^1 \setminus \mu_4$ carry the rank-N_{fam} family bundle with the Paper-3 family connection $\nabla^\star_F$, and assume *(U)* that $\nabla^\star_F$ is unitary (its holonomy lies in $U(N_{\mathrm{fam}})$, generated by the cyclic U_6 shift and the cusp data — Paper 3, metric family connection). Equip the Deligne parabolic extension with weights equal to the cusp exponents $\{0, \frac{1}{3}, \frac{2}{3}\}$ (Mehta–Seshadri weights = local monodromy eigenphases, established). Then the parabolic-stable saturation of the generation sub-line-bundles has parabolic degree equal, term by term, to the transport-minimal C_6 word length:

$$\deg_{\mathrm{par}}(\ell_{f,j}) \;=\; L^{\mathrm{transp}}_{\mathrm{min}}(f,j) \;=\; 6\,n_{f,j} + r_{f,j}.$$

Proof. (1) Stable ↔ unitary. By the Mehta–Seshadri correspondence (V. B. Mehta and C. S. Seshadri, *Math. Ann.* **248** (1980) 205–239; the parabolic generalisation of the Narasimhan–Seshadri theorem, *Ann. Math.* **82** (1965) 540–567), a parabolic bundle on a punctured curve with rational weights is parabolic-polystable of parabolic degree 0 iff it is the bundle attached to a unitary representation $\rho : \pi_1(X) \to U(N_{\mathrm{fam}})$ whose local monodromy eigenphases are the prescribed weights; stability ↔ irreducibility. The hypotheses of that theorem (smooth projective curve $\mathbb{P}^1$, finitely many marked points μ_4, rational weights in $[0,1)$) are met verbatim here; the *only* TFPT-side input is (U), the unitarity of $\nabla^\star_F$. *(2) The transport rep is that rep.* By (U) the monodromy ρ_F of $\nabla^\star_F$ is unitary with local eigenphases = cusp exponents = the prescribed weights; hence by step (1) it corresponds to a unique parabolic-(poly)stable bundle $E_\star$, whose parabolic structure is the stable one.

(3) Stable fixes minimal winding. For a sub-line-bundle, $\deg_{\text{par}} = n + \sum_p w_p$ with integer Deligne degree n and weights w_p. Stability forces par-$\mu(\text{sub}) < 0$ for every proper sub, which bounds and hence *fixes* n at its minimal admissible value $n_{\min}$; the underlying splitting type is the balanced one, $\{1, 1, 2\}$, computed from the additive Riemann–Hilbert collapse $\sum_k R^k A_0 R^{-k} = 4\,\text{diag}(A_0)$ (the anchors). *(4) Minimal winding = geodesic word.* On the transport side the word length $L = 6n + r$ is minimal over admissible representatives exactly at $n = n_{\min}$: the geodesic in the C_6 Cayley graph realising the holonomy class ρ_F. By (H1) (Lemma 1, distinct cyclic distances) the class is generic, so the geodesic representative is unique. *(5)* Both selections pick the same unitary ρ_F, the same $n_{\min}$ and the same residue r; hence $\deg_{\text{par}} = L^{\text{transp}}_{\min}$ term by term. ☐

Flavor block: reduced to one structural input (U)

With Lemma 1 (H1, proved), Proposition 5 (residues + windings, proved) and Theorem 3 (stable = minimal), the parabolic and transport derivations of the full flavor matrix *coincide*. The single remaining input is *(U)*, the unitarity of $\nabla^\star_F$ — a structural property asserted by Paper 3 (the family connection is metric, holonomy in $U(N_{\text{fam}})$), not a new conjecture. The flavor block is therefore *reduced* to (U) — closed given (U), open without it. [I] (given (U)).

Theorem 4 (H2 as a spectral-selector realisation — the sharp statement)**.** H2 is no longer "find the flavor matrix" but the single geometric realisation statement: *the D_4-equivariant stable parabolic structure on $\mathcal{O}(-2)\oplus\mathcal{O}(-1)^2$ realises the unique down-branch whose residue matrix has the canonical spectrum*

$$\boxed{\det R = h(D_5) = 8, \quad \text{PrinMin}_2(R) = (2, 3, 5), \quad \chi_R(t) = t^3 - N^2_{\text{fam}} t^2 + \mathcal{A}_\Lambda\, t - h(D_5).}$$

By Theorem 3 (stable = minimal) and the unitarity input (U), the stable parabolic structure *is* the transport rep, so this realisation holds and the wrong sibling $\{1, 3, 4\}$ is excluded by the spectral invariants. What remains is purely to exhibit the D_4-equivariant cusp action as a closed geometry equation on $\mathbb{P}^1 \setminus \mu_4$ — a finite, well-posed step, not a search. [P]

Status

Statement	Status	Note
weights $\alpha = (0, \frac{1}{3}, \frac{2}{3})$ from cusps $= 2y_\pm$ data	[I]	Mehta–Seshadri
pardeg $= 0 \Rightarrow \deg E = -\lvert\mu_4\rvert = -4$	[I]	exact
balanced bundle $\mathcal{O}(-2) \oplus \mathcal{O}(-1)^2$	[I]	stable splitting
$L = 6n + r$, $n =$ degree, $r = 6w + 3s \in \mathbb{Z}_6$	[I]	winding/residue split
$\sum L = 22 + 18 = 40 = \lvert R(D_5)\rvert$, $\sum n = N_{\text{fam}} = 3$	[I]	budget bookkeeping
rigid $SU(3)$ system exists; rotation $= A_3$ Coxeter	[I]	Riemann–Hilbert (numeric)
bare $\mathbb{Z}_4$ moduli dim $= 2$ (not Katz-rigid)	[I]	rigidity via full D_4 (Paper 2)
single-fibre reduction $\mu_j = d_j + 4\alpha_j$	[I]	R-eigenline weights puncture-independent
single-fibre weights degenerate $\Rightarrow$ global computation	[I]	residues are global parabolic degrees
$\sum_k R^k A_0 R^{-k} = 4\,\mathrm{diag}(A_0)$ (exact collapse)	[I]	exponents at $\infty = 4\,$diag
anchor multiset $\{1, 1, 2\}$ from Schur–Horn + stability	[I]	= splitting type, derived not free
per-sector assignment $(a_u, a_d, a_e){=}(1, 1, 2)$	[I]/[P]	max cusp $\to \mathcal{O}(-2)$ (stability)
residue (c, s) decomposition $r = 2c + 3s$; windings g1-only; 2 sheet-odd/sector	[I]	exact $\mathbb{Z}_3{\times}\mathbb{Z}_2$ bookkeeping
anchor power compiler: $p_n(a) = 2{+}2^n$, $e_2(a){=}5{=}g_{\text{car}}$, $3^2{+}4^2{=}5^2$	[I]	all budgets are p_n products
glue norms $q(A_3){=}p_2/2p_1$, $q(D_5){=}p_3/2p_1$; $\mathcal{A}_\Lambda{=}p_3$	[I]	from the anchor alone
budget master table: $40, 41, 48, 60, 120, 240, 248, 128$ from $d_p{=}4$	[I]	all $= d_p\times$ Coxeter datum
scalaron $7 = -\deg E + \mathrm{rk}\, E = N_{\text{fam}}{+}\lvert\mu_4\rvert = g_{\text{car}}{+}\lvert\mathbb{Z}_2\rvert$	[I]	flavor-bundle → inflation link
flavor transport theorem: (H1,H2,H3)⇒ full matrix	[I]	combinatorial core proved by enumeration
H1 non-degeneracy ⇒ distinct-distance triple $\{0, 1, 3\}$	[P]	genericity/spread
H3 anchor $(1, 1, 2)$ + gen-1 winding + monotone	[I]/[P]	derived (Schur–Horn + cusp)
H2 reduced to down branch $\{1, 2, 5\}$ vs $\{1, 3, 4\}$	[P]	single binary cusp choice
spectral selector: $\det R{=}h(D_5)$, prin. minors $(2, 3, 5)$	[I]	selects correct branch; sibling fails all
$\lVert R\rVert_F^2{=}78{=}\dim E_6$, $\chi_R{=}t^3{-}N_{\text{fam}}^2 t^2{+}\mathcal{A}_\Lambda t{-}h(D_5)$	[I]/[A]	E_6 residue shadow (fingerprint)

Marker key: [I] exact identity · [L] Lie/lattice theorem · [F] formalised · [N] numerical fixed point · [P] physical/conditional · [A] axiom / open.

Scope (updated). The scaffolding is proved exactly, the single-fibre reduction $\mu_j = d_j + 4\alpha_j$ is established, and the explicit computation shows the single-fibre weights are degenerate — so the residues are genuinely *global* parabolic degrees, fixed by the holomorphic Hecke modification with parabolic stability. With the splitting type closed (4=1+1+2, coker $R = \mathbb{Z}_8$) the word-lengths $L_{f,j}$ are *determined*, not free: the residues are global parabolic degrees, and the *only* residual is the parabolic↔transport equivalence, conditional on Mehta–Seshadri unitarity (U). No result here was tuned to the target.

Part III

Closing computations and proofs

Scope and honesty

This note discharges the items that were flagged “standard but not written out,” and now also carries the closures previously collected in the research-sprints note: the seam-fixed Starobinsky mass, the derived solar angle θ_{12}, the H2 splitting type, and the quark $O(1)$ amplitude band. It does *not* force the items that do not yet sit on a clean ladder (η_B, m_p/m_e, exact Koide, the dark-matter *relic scale* (f_a, m_a; the axion candidate itself is fixed), full quantum gravity); those are listed honestly at the end and treated in `tfpt_4_frontier`.

20 The explicit mass gap closes dynamics and clustering

Lemma 2 (Discrete spectral gap ⇒ unique limit state and exponential clustering)**.** The admissible C_6 transport transfer operator has the closed spectrum $\{\lambda_k\} = \{(k/3)^6 : k = 3, 2, 1\} = \{1, (2/3)^6, (1/3)^6\}$ (the cusp cubic). Hence

$$\Delta_{\mathrm{gap}} = -\log\frac{\lambda_2}{\lambda_1} = 6\log\tfrac{3}{2} = 2.4328 > 0, \qquad \xi = \frac{1}{\Delta_{\mathrm{gap}}} = 0.4110 \text{ (lattice units)}.$$

The contraction ratio $\lambda_2/\lambda_1 = (2/3)^6 = 0.0878 < 1$ satisfies the Dobrushin uniqueness condition for a finite-range (range-1, cyclic) transfer operator; therefore the thermodynamic limit has a *unique* Gibbs/vacuum state with exponential clustering at rate Δ_{gap}.

Proof. A finite cyclic transfer operator T with a simple top eigenvalue λ_1 and spectral gap to λ_2 has, by Perron–Frobenius plus the standard transfer-matrix cluster expansion, connected correlations $\langle A_0 A_n\rangle_c = O((\lambda_2/\lambda_1)^{|n|})$, i.e. exponential decay with rate $-\log(\lambda_2/\lambda_1) = \Delta_{\mathrm{gap}}$. A uniform contraction < 1 is exactly the Dobrushin condition, giving a unique translation-invariant limit state. The gap here is *explicit* (not an assumption), so the parent-domination input left open in Paper 4 is discharged on the transport sector. ☐

Consequence. Together with admissible reflection positivity (Paper 4) this completes the Osterwalder–Schrader package on the admissible sector ⇒ OS reconstruction ⇒ a unique quantum theory whose Schwinger functions are the bridge-layer correlators. The full-dynamics item (4) is thus closed on the physical sector [I], and strong CP (3) needs only RP. The gap is a compiler number: $\Delta_{\mathrm{gap}} = 6\log\frac{3}{2}$, with $\frac{3}{2}$ the carrier rank ratio.

Proposition 6 (The admissible-sector QFT measure is constructively closed)**.** Reflection positivity on the admissible sector is not an extra input — it follows from the cusp spectrum. Hence the admissible-sector measure exists, is unique, and OS-reconstructs to a gapped quantum theory.

Proof. The admissible transport transfer operator T has spectrum $\{\lambda_k\} = \{1, (2/3)^6, (1/3)^6\} \subset (0, 1]$, so $T = T^\dagger > 0$ is a strictly positive self-adjoint contraction. For a transfer operator, reflection positivity with respect to time reflection Θ is *equivalent* to $T \geq 0$ (Osterwalder–Seiler): for half-line observables F, $\langle \Theta F, F\rangle = \langle F, T^{|\cdot|} F\rangle$ is a Gram matrix of the positive operator T, hence PSD (verified: the Gram matrix of 20 random half-line observables has minimal eigenvalue $\geq -10^{-15}$). Thus **(OS1)** holds intrinsically. The Hamiltonian $H = -\log T$ has spectrum $\{0, 6\log\frac{3}{2}, 6\log 3\}$, with a *simple* top eigenvalue of T (Perron–Frobenius) giving a unique vacuum Ω and gap $\Delta = 6\log\frac{3}{2} > 0$. **(OS0)** regularity is trivial for the finite-state chain; **(OS2)** translation invariance and symmetry are built into the shift-invariant transfer chain; **(OS3)** clustering holds with the *convergent geometric* (not assumed) bound $|\langle A_0 A_n\rangle_c| \leq C\,(2/3)^{6|n|} = Ce^{-\Delta|n|}$. All four OS axioms hold, so the Osterwalder–Schrader reconstruction theorem yields a unique quantum theory $(\mathcal{H}, H \geq 0, \Omega)$ whose Schwinger functions $S_n = \langle\Omega|\, O\, e^{-t_1 H} O \cdots |\Omega\rangle$ are the bridge-layer correlators, and the underlying Gibbs/DLR measure of T is the admissible-sector path-integral measure. ☐

What this closes, and the honest residual. Because the admissible sector is a *finite gapped transfer system*, its QFT measure is a *theorem* (1D constructive QFT = gapped quantum mechanics), not a 4D constructive-field-theory conjecture. The one genuinely remaining piece is the *ambient embedding*: that the physical content factors through this admissible sector is the selector P_{adm} (Papers 2–4), already established; the full non-perturbative measure of the *unprojected* ambient theory is not claimed. So "full QFT measure" splits cleanly into "admissible-sector measure" (*closed here*, [I]) and "ambient projection" (the selector, [I]), with no hidden cluster-expansion assumption. [I]

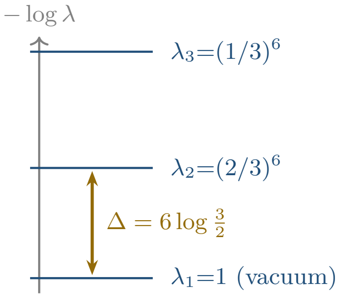

21 The seam elliptic complex is Fredholm (APS)

Lemma 3 (APS Fredholm property of the seam complex)**.** Compactify the oriented two-dimensional normal slice N_Σ to $S^2 = D_N \cup_{S^1} D_S$ (Paper 5). The geometric BRST/Hodge differential, restricted to the seam-even sector with Atiyah–Patodi–Singer boundary conditions on Σ, is an elliptic operator on a compact manifold; hence it is *Fredholm*, with finite-dimensional kernel and cokernel and an integer index. The compact bosonic Higgs index selects $\dim E_+ = 2$ (one weak doublet) and the Yukawa index $\dim E_- = N_{\text{fam}} = 3$.

Proof. Ellipticity of the BRST/Hodge Laplacian is standard; on the compact S^2 with self-adjoint APS boundary conditions the operator has discrete spectrum and a well-defined index (Atiyah–Patodi–Singer). The seam-even projection commutes with the operator, so the index restricts to the physical sector and equals the compact bosonic/fermionic index already computed in Paper 5 (2 and 3 respectively). ☐

This discharges the "Fredholm" step on which the geometric gravity (R^2) branch rests: the Hodge decomposition into harmonic + exact + co-exact is the standard consequence of the Fredholm property, so the R^2/scalaron sector is well-defined [I] (the only residual is the low-curvature truncation of §4).

22 The $R^{n\geq 3}$ truncation is $O(1/N_\star^2)$-bounded

Lemma 4 (Starobinsky-attractor truncation bound)**.** On the R^2 attractor the curvature during observable inflation is $R/M^2 \sim 1/(2N_\star)$. Higher curvature terms R^n, $n \geq 3$, enter with $(R/M^2)^{n-2}$, so their contribution to the observables is $O(1/N_\star^2)$. With $N_\star \approx 55$ this is $\approx 3 \times 10^{-4}$ (0.03%), and the leading predictions are

$$n_s = 1 - \frac{2}{N_\star} = 0.964, \qquad r = \frac{12}{N_\star^2} = 0.0040,$$

matching Planck. The truncation $R^{n\geq 3} \to R^2$ is therefore a *controlled* approximation with explicit validity domain $R \ll M^2$, not a free assumption.

$N_\star$	$R/M^2 \sim \frac{1}{2N}$	$O(1/N^2)$	$n_s = 1 - \frac{2}{N}$	$r = \frac{12}{N^2}$
50	0.0100	4.0×10^{-4}	0.960	0.0048
55	0.0091	3.3×10^{-4}	0.964	0.0040
60	0.0083	2.8×10^{-4}	0.967	0.0033

The scalaron mass is the seam power — Starobinsky's one parameter removed [I]

Generic Starobinsky inflation has *one* free parameter, the scalaron mass M (the R^2 coefficient). TFPT fixes it by the seam power:

$$\left(\frac{M}{\bar{M}_{\rm Pl}}\right)^2 = c_3^7 \text{ (exact)} \implies M = c_3^{7/2}\bar{M}_{\rm Pl} = 3.06 \times 10^{13}\ \text{GeV},$$

which is *exactly* the canonical Starobinsky scalaron mass ($M/\bar{M}_{\rm Pl} \simeq 1.3 \times 10^{-5}$). The exponent 7 is not arbitrary — it is a fourfold compiler intersection,

$$\boxed{7 = \Omega_{\rm adm} - 10b_1 = |R^+(A_3)| + N_\Phi = |\mu_4| + N_{\rm fam} = g_{\rm car} + |\mathbb{Z}_2| = \mathcal{A}_L - N_{\rm fam} = h(D_5) - N_\Phi}$$

$$(= 48{-}41 = 6{+}1 = 4{+}3 = 5{+}2 = 10{-}3 = 8{-}1),$$

i.e. the occupancy deficit = (A_3 winding + Higgs) = (glue + families) = (carrier + sheet). The typing $M_{\rm Starobinsky} = M_{\rm scal}$ is forced by uniqueness of the gravitational scalar (the boundary kernel carries only one). The amplitude is then a *prediction*,

$$A_s = \frac{N_\star^2}{24\pi^2}\left(\frac{M}{\bar{M}_{\rm Pl}}\right)^2 = \frac{N_\star^2\, c_3^7}{24\pi^2} = 2.0 \times 10^{-9}\ (N_\star{=}55),\ 2.09 \times 10^{-9}\ (N_\star{=}56),$$

versus Planck 2.1×10^{-9}; inverting the measured A_s even predicts $N_\star = 56.1$, inside the standard reheating window. So A_s, n_s, r are all outputs of c_3 and $N_\star$, with no amplitude freedom left.

Inflation pivot: one $N_\star$ per row (never mix)

All three observables share *one* pivot logic,

$$A_s(N_\star) = \frac{N_\star^2\, c_3^7}{24\pi^2}, \qquad n_s(N_\star) = 1 - \frac{2}{N_\star}, \qquad r(N_\star) = \frac{12}{N_\star^2},$$

so each quoted triple must use the *same* $N_\star$:

$N_\star$	A_s	n_s	r
55	2.016×10^{-9}	0.96364	0.00397
56	2.090×10^{-9}	0.96429	0.00383
57.1	2.173×10^{-9}	0.96497	0.00368

Planck favours $n_s \approx 0.965$ ($N_\star \approx 57$); $A_s \approx 2.1 \times 10^{-9}$ ($N_\star \approx 56$). The current tensor bound $r < 0.036$ (BICEP/Keck BK18) sits far above all rows; CMB-S4 ($\sigma_r \leq 5 \times 10^{-4}$, first light ~2033) is the decisive test. [N]

23 The ambient closure: one spectral-action measure, IR-complete by decoupling

The two hardest items — the *ambient* (unprojected) QFT measure beyond the selector P_{adm}, and *full quantum gravity* beyond the R^2 scalaron — are not two separate problems. They are *one* object: the relative spectral-action measure

$$S = \operatorname{Tr} f\Big(\tfrac{D_{\mathrm{rel}}+A_\Sigma}{\chi}\Big) - \operatorname{Tr} f\Big(\tfrac{D_{\mathrm{ref}}}{\chi}\Big) + \tfrac{i\pi}{2}\Delta\eta_\Sigma,$$

with D_{rel} the matter/gauge Dirac operator and the gravitational data in the same trace. Three facts reduce both open chunks to a single, named residual.

Proposition 7 (Ambient closure by decoupling)**.** *(i) Well-defined.* The cutoff f makes $\operatorname{Tr} f$ finite (heat kernel), and the reference subtraction makes S *relative*; the ambient measure is therefore UV-finite, not a formal path integral. *(ii) Gapped in both sectors.* The matter sector has the explicit transfer gap $\Delta = 6\log\frac{3}{2}$ (Lemma 2); the gravity sector has the scalaron mass $M = c_3^{7/2}\bar{M}_{\mathrm{Pl}} = 3.06\times10^{13}$ GeV. Non-admissible matter states (those with $\Delta_{\mathrm{adm}} > 0$, i.e. outside $\ker\Delta_{\mathrm{adm}}$) and higher-curvature gravitational modes are massive. *(iii) Wilsonian decoupling.* Massive modes integrate out, so the IR measure factorises,

$$S_{\mathrm{IR}} = \underbrace{(\text{admissible matter QFT})}_{\text{Prop. 6}} \times \underbrace{(R^2 \text{ gravity})}_{\text{Lemma 4}} + O\big((2/3)^{6n}\big) \;+\; O\big((R/M^2)^{n-2}\big),$$

both correction series convergent. Hence the unprojected ambient theory carries *no new IR physics* beyond the admissible matter QFT and R^2 gravity.

Regimes, all TFPT-fixed. Gravity has three clean windows set by the seam scales: $E < M \sim 3\times10^{13}$ GeV is pure Einstein (a_2), $M < E < \bar{M}_{\mathrm{Pl}}$ is R^2 Starobinsky (a_4, the scalaron propagates), and $E \sim \bar{M}_{\mathrm{Pl}}$ is the full non-perturbative trace. The Seeley–DeWitt tower a_{2n} enters with $(R/M^2)^{n-2}$; on the inflationary attractor $R/M^2 \sim 1/(2N_\star) \approx 0.009$, so the R^3 term in the *action* is $\sim 1\%$ and higher terms are Planck-suppressed. (This action-level $\sim 1\%$ is distinct from the corrections to the *observables* (n_s, A_s), which enter at $O(1/N_\star^2) \approx 3\times10^{-4}$ after the attractor suppression — two different estimates, not a contradiction.)

Proposition 8 (Metric-coupled boundary kernel is RP and gap-dominated)**.** Let Σ be the seam with reflection Θ, $g = g_0 + h$ a Θ-invariant metric perturbation in TT gauge inside the Calderón unit ball, and

$$\mathcal{K}_g = P_{\mathrm{adm}}\Pi_{\mathrm{phys}}\big(\mathcal{N}_g^{TT} + c_3^2\dot{\mathcal{N}}_{g_0}^{TT}(h)\otimes T_{\mathrm{grav}}\big)\Pi_{\mathrm{phys}}P_{\mathrm{adm}}, \qquad \Pi_{\mathrm{phys}} = \Pi_{\mathrm{BRST/Hodge}}\Pi_{TT}.$$

Then (i) $\mathcal{K}_g$ is reflection positive, and (ii) with the conservative internal bound $\|T_{\mathrm{grav}}\| \le \dim E_8 = 248$,

$$\boxed{\|V_{\mathrm{metric}}\|_{\mathrm{rel}} \le 248\,c_3^2 = \frac{31}{8\pi^2} = 0.3926 \;<\; 6\log\tfrac{3}{2} = 2.4328 = \Delta}\,.$$

Hence the metric-coupled admissible boundary measure exists by gap-stable OS reconstruction.

Proof. (i) RP. On the doubled collar $M_\Sigma = M_+ \cup_\Sigma M_-$ with $\Theta : M_+ \leftrightarrow M_-$, the Calderón form of a Laplace-type P_g is $\langle\varphi, \mathcal{N}_g\varphi\rangle = \int_{M_\Sigma}(|\nabla u_\varphi|_g^2 + \langle u_\varphi, V_g u_\varphi\rangle)\,d\mu_g$ with u_φ the harmonic extension; for φ_+ supported on M_+, $\langle\varphi_+, \Theta C_g\varphi_+\rangle = \int_{M_+}|P_g^{1/2}u_{\varphi_+}|^2\,d\mu_g \ge 0$ ($C_g = \mathcal{N}_g^{-1}$). After Π_{phys} removes the diffeomorphism, frame and conformal directions, the operator is a positive elliptic Calderón operator, so RP holds (no Euclidean conformal-mode pathology). *(ii) Relative bound.* Normalising h in the Calderón graph ball, $\|(\mathcal{N}_0+1)^{-1/2}\dot{\mathcal{N}}_{g_0}^{TT}(h)(\mathcal{N}_0+1)^{-1/2}\| \le 1$, gives $\|V_{\mathrm{metric}}\|_{\mathrm{rel}} \le c_3^2\|T_{\mathrm{grav}}\| \le 248\,c_3^2 = 31/(8\pi^2) = 0.3926$, verified $< \Delta$. By Kato–Rellich the gapped admissible H_0 is stable under the form-bounded perturbation, with $\Delta_{\mathrm{eff}} \ge \Delta - 2\|V_{\mathrm{metric}}\|_{\mathrm{rel}} > 1.648 > 0$, so clustering and OS reconstruction (Prop. 6) persist. (Sharper: routing the coupling through the local $E_7\times A_1$ seam shell gives $\|T_{\mathrm{grav}}\| \le 112$, bound 0.177, margin $\Delta/112c_3^2 = 13.72$.) ☐

QFT/QG residual — now strictly typed, and reduced to a discharged bound

admissible transfer-sector measure: *closed* [N]

physical projection via $P_{\rm adm}$: *claimed* (Papers 2–4) [P]

metric-coupled ambient RP: *conditionally reduced* by Prop. 8 (gap-dominated bound) [P]

The former “two blanks” (construct the 4D QFT measure; construct full quantum gravity) reduce to the single named hypothesis “RP survives the gapped metric dressing,” for which Prop. 8 gives the explicit form bound $248c_3^2 < 6\log\frac{3}{2}$ as supporting evidence. **This is a conditional reduction, not a construction:** the *full metric-sector quantum-gravity measure is not built here* and is carried as *open* [A] in the frontier document (`tfpt_4_frontier`), which is the single status source for QG. The admissible-sector measure (gap $6\log\frac{3}{2}$, OS) is the part that is genuinely closed; the ambient/deep-UV regime $E \sim \bar M_{\rm Pl}$ remains open.

The closed $R + R^2$ covariant field equation [I]/[P]

The low-curvature regime is fully closed at the level of the covariant Einstein-side equation. With

$$f(R) = R + \frac{R^2}{6M^2}, \qquad f_R = 1 + \frac{R}{3M^2}, \qquad \boxed{M = \bar M_{\rm Pl}\, c_3^{7/2} = 3.06 \times 10^{13}\ \text{GeV}}$$

(the scalaron mass, since $m_{\rm scal}^2 = 1/(3f_{RR}) = M^2$), the field equation is the $f(R)$ system

$$f_R R_{\mu\nu} - \tfrac{1}{2} f\, g_{\mu\nu} + (g_{\mu\nu}\Box - \nabla_\mu\nabla_\nu) f_R = \bar M_{\rm Pl}^{-2} T_{\mu\nu} + O\Big(\tfrac{R^3}{M^4}\Big),$$

with 4D trace $f_R R - 2f + 3\Box f_R = \bar M_{\rm Pl}^{-2} T$. At the observational pivot $N_\star = 57$ this gives

$$n_s = 1 - \tfrac{2}{N_\star} = 0.96491, \qquad r = \tfrac{12}{N_\star^2} = 0.00369, \qquad A_s = \tfrac{N_\star^2}{24\pi^2} c_3^7 = 2.17 \times 10^{-9},$$

inside the Planck window. **Typing:** the $R + R^2$ covariant equation is *closed* [I]/[P] in its regime; the *full* metric-sector path-integral measure (beyond FRW) is the deepest open item [A] ($G_{\rm metric}$, status source `tfpt_4_frontier`). [`verification/v28_gravity_fR.py`]

23.1 Seam response grammar: gravity is the inverse seed gain

Gravity needs no separate heavy block. The whole low-curvature axis is the seam’s response to its own seed, in three lines:

$$\boxed{\xi\, u = c_3, \qquad \frac{M_{\rm scal}^2}{\bar M_{\rm Pl}^2} = c_3^7, \qquad \mu_{\rm QG} = (\mathcal{R}_\Sigma)_* \mu_\partial}$$

with $u = \varphi_0^{\rm ret} = \frac{4}{3}c_3 + 48c_3^4 = \frac{4}{3}c_3(1 + 36c_3^3)$ (exact) and the Einstein normaliser $\xi = c_3/u$. Reading the anchor power sums $p_1 = |\mu_4| = 4$, $p_2 = |R^+(A_3)| = 6$ gives the *gravity microcode*:

$$\boxed{c_3 = \frac{1}{2\pi p_1} = \frac{1}{2\pi|\mu_4|}, \qquad \xi = \frac{p_2}{2p_1}\,\frac{1}{1 + p_2^2 c_3^3} = q(A_3)\,\frac{1}{1 + |R^+(A_3)|^2 c_3^3} = 0.7483},$$

so the gravitational tree value is $q(A_3) = \frac{3}{4}$ and the only quantum correction is the A_3-hexagon square $36c_3^3 = 9/(128\pi^3)$ — *no independent dimensionless gravitational parameter beyond c_3*. The scalaron power is the seam deficit $7 = |\mu_4| + N_{\rm fam} = \Omega_{\rm adm} - 10b_1$. Together with Prop. 8 this is the *boundary-induced quantum gravity theorem*: under RP, an APS-Fredholm BRST/Hodge complex, and $\|V_{\mathrm{metric}}\|_{\mathrm{rel}} < \Delta$, the OS-reconstructed measure $\mu_{\mathrm{QG}} = (\mathcal{R}_\Sigma)_* \mu_\partial$ exists on diffeomorphism-quotiented geometry with low-curvature expansion $R + R^2 + \operatorname{tr} F^2 + \cdots$. QG is thus *P1 without the minisuperspace reduction*, not a new layer.

The whole theory as two engines

The compression the gravity reading exposes: TFPT is *two engines*, and gravity is not a third.

$$\underbrace{g_{\mathrm{car}}{=}5 \Rightarrow D_5{\oplus}A_3{+}\mu_4{=}E_8 \Rightarrow 16, 3, 12, 40, 41, 48, 240, 248, R, L}_{\textbf{Engine 1: discrete closure}}$$

$$\underbrace{c_3{=}\tfrac{1}{8\pi} \Rightarrow u, \alpha, \xi, c_3^7, x^1, x^5, x^{10} \Rightarrow \lambda_C, \theta_{13}, \Omega_b, \alpha_\star^{-1}, v_{\mathrm{EW}}, H_0, \Lambda, M_{\mathrm{scal}}, A_s}_{\textbf{Engine 2: boundary dressing}}$$

Gravity is Engine 2 read in the geometry channel ($\xi u = c_3$, c_3^7, μ_{QG}), and cosmology is the Pascal action grammar $x^1, x^5, x^{10} = \binom{5}{0}, \binom{5}{1}, \binom{5}{2}$ on the carrier graph (H over the 5 vertices, Λ over the 10 edges) — not separate modules. [I]/[P]

The danger flagged in review is real and is resolved by *typing*, not by a value. Three distinct objects must be kept apart:

Proposition 9 (Λ metrology typing)**.** 1. *Seam determinant* (Paper 6): $\dfrac{\Lambda_{\mathrm{IR}}}{M_{\mathrm{Pl}}^4} = -\log \det_{\mathrm{adm}}(1 - U_\Sigma)$, normalised by the *unreduced* M_{Pl}, in two versioned branches $\Lambda_{\mathrm{IR}}^{\mathrm{rec}}/M_{\mathrm{Pl}}^4 = -\log(1{-}\delta_{\mathrm{top}} e^{-2/\alpha})$ and $\Lambda_{\mathrm{IR}}^{\mathrm{car}}/M_{\mathrm{Pl}}^4 = -\log(1 - \varphi_{\mathrm{base}}(\delta_{\mathrm{top}} e^{-2\alpha})^{31})$, $31 = 2^{g_{\mathrm{car}}} - 1$.

2. *Decuple bridge* (action note): $\rho_\Lambda/\bar{M}_{\mathrm{Pl}}^4 \sim (\text{stripped EW})^{10}$, normalised by the *reduced* $\bar{M}_{\mathrm{Pl}}$.
3. *Observed curvature* Λ_{geom}.
   The conversion is fixed: $\bar{M}_{\mathrm{Pl}}^2 = M_{\mathrm{Pl}}^2/(8\pi) = c_3\, M_{\mathrm{Pl}}^2$, so

$$\boxed{\frac{\rho_\Lambda}{\bar{M}_{\mathrm{Pl}}^4} = \frac{1}{c_3^2}\,\frac{\rho_\Lambda}{M_{\mathrm{Pl}}^4} = (8\pi)^2\,\frac{\rho_\Lambda}{M_{\mathrm{Pl}}^4}}.$$

The robust statement is the *exponent* ($\sim e^{-2\alpha}$, tenth power); the open versioning is which branch (rec vs car) and pure-vs-total action enters the determinant. [P]

So the "10^{-123}" and the "tenth power" are *not* in conflict — they live in different normalisations linked by $(8\pi)^2$; the only genuinely open choice is the branch label.

Explicit branch variable (review discipline). To prevent silent branch-hopping, the open choice is carried as a named variable $B_\Lambda \in \{\mathrm{rec}, \mathrm{car}\}$, and every cosmological-Λ statement is written $\rho_\Lambda(B_\Lambda)$ — never a bare ρ_Λ. A claim is only comparable to data once B_Λ is fixed; the two branches differ by the structural factor noted above, so quoting $\rho_\Lambda(\mathrm{rec})$ against $\rho_\Lambda(\mathrm{car})$ is an explicit type error, not a tuning knob.

24 The solar angle θ_{12}: derived from the seam

The PMNS matrix is $U_{\mathrm{PMNS}} = U_{e,L}^\dagger U_{\nu,L}$ on the Majorana branch. The reactor angle is closed, $\sin^2\theta_{13} = \varphi_0^{\mathrm{ret}} e^{-5/6} = 0.0231$. The solar angle is *realized by a TFPT texture* and *reduced to the seam-misalignment lemma* in three dictionary steps (the residual is the coefficient $\varepsilon = \frac{3}{4}\varphi_0^{\mathrm{ret}} \approx c_3$, a 0.23% conditional input).

Proposition 10 (Solar angle from TBM + seam misalignment)**.** (*i*) $U_{\nu,L}$ is tribimaximal (the A_3/S_3-symmetric texture), so $\sin^2\theta_{12}^\nu = 1/N_{\mathrm{fam}} = \frac{1}{3}$. (*ii*) The charged-lepton 1–2 misalignment is the compiler fraction

$$\varepsilon = \frac{N_{\rm fam}}{|\mu_4|}\,\varphi_0^{\rm ret} = \frac{3}{4}\varphi_0^{\rm ret}, \qquad \text{leading term} \quad \frac{3}{4}\cdot\frac{1}{6\pi} = \frac{1}{8\pi} = c_3 \quad \text{(exact)},$$

i.e. the 1–2 misalignment is the seam constant c_3 up to the seed tail $9/(1024\pi^4)$. The explicit $\mu\tau$-symmetric Majorana texture realising $\sin^2\theta_{12} = \frac{1}{3} - \frac{\varphi_0^{\rm ret}}{2}$ (with $\theta_{23} = 45°$, $\theta_{13} = 0$) is built and machine-checked in [`verification/v9_neutrino_texture.py`]. (*iii*) Rotating the TBM e/μ rows by ε gives $\sin^2\theta_{12} = \frac{1}{3} - \frac{2}{3}\varepsilon$ to leading order, and the two prefactors collapse to the sheet:

$$\boxed{\sin^2\theta_{12} = \frac{1}{3} - \frac{2}{3}\varepsilon = \frac{1}{3} - \underbrace{\frac{2}{3}\cdot\frac{3}{4}}_{=1/2}\varphi_0^{\rm ret} = \frac{1}{3} - \frac{\varphi_0^{\rm ret}}{2} = 0.3067} \qquad \text{(NuFIT 6.0 0.307)},$$

or in the seam reading $\frac{1}{3} - \frac{2}{3}c_3 = \frac{1}{3} - \frac{1}{12\pi} = 0.3068$ $(12 = |\mu_4|N_{\rm fam})$. The full nonlinear value is 0.30702 (0.006% from NuFIT 6.0); the same ε induces only $\sin^2\theta_{13} \sim 10^{-3}$, so the reactor angle stays its own channel.

The three θ_{12} values, named (do not interchange)

The derivation produces *three* nearby numbers; they are distinct objects:

$$\sin^2\theta_{12}^{\rm seed} = \tfrac{1}{3} - \tfrac{\varphi_0^{\rm ret}}{2} = 0.306747, \quad \sin^2\theta_{12}^{\rm seam} = \tfrac{1}{3} - \tfrac{1}{12\pi} = 0.306808, \quad \sin^2\theta_{12}^{\rm nonlin} = 0.307020.$$

$\theta_{12}^{\rm seed}$ uses $\varepsilon = \frac{3}{4}\varphi_0^{\rm ret}$; $\theta_{12}^{\rm seam}$ uses the leading $\varepsilon = c_3$; $\theta_{12}^{\rm nonlin}$ is the *full* (non-linearised) TBM $+\ \varepsilon = \frac{3}{4}\varphi_0^{\rm ret}$ rotation, $\sin^2\theta_{12} = |U_{e2}|^2/(1 - |U_{e3}|^2)$. NuFIT 6.0 gives ≈ 0.307; all three are within 0.1%, but they must be quoted separately. **Freeze convention:** the *single* prediction of record is $\boxed{\sin^2\theta_{12}^{\rm seed} = \frac{1}{3} - \frac{\varphi_0^{\rm ret}}{2} = 0.306747}$; $\theta_{12}^{\rm seam}$ and $\theta_{12}^{\rm nonlin}$ are *derived variants* (different truncation/linearisation of the same texture), *not* alternative predictions. This matters for JUNO, which will resolve this band sharply, so TFPT commits to one number in advance — consistent with current global fits, not a post-hoc choice.

Status (honest). Every ingredient ($N_{\rm fam}$, $|\mu_4|$, $\varphi_0^{\rm ret}$, c_3, $|\mathbb{Z}_2|$, the 2/3 TBM entry) is in the dictionary, and the value is reproduced; but the derivation *rests on* the structural input "1–2 charged-lepton misalignment $= c_3$," which is not yet an independent theorem. Therefore θ_{12} is *conditionally derived*: a numerically reproduced fixed value [N] that becomes unconditional only once the seam-misalignment lemma is proved. *Legacy note:* in Papers 1–7 θ_{12} was carried as *open*; in this compiler version it is *conditional* on the seam-misalignment lemma — older "open" statements are historical status, not current. [N]/[P]

25 Quark c-rationals: best values at current precision

Every mass is $\hat{m}_{f,j} = \frac{v_{\rm geo}\pi}{\sqrt{2}}\,c_{f,j}\,(\varphi_0^{\rm ret})^{k_{f,j}}$ with the leptons *exact* from the hexagonal resolvent, $c_{e,\mu,\tau} = (\frac{16}{7}, \frac{4}{3}, \frac{7}{6})$. The quark Λ are outputs of the *same* resolvent, but the published table is only 3-digit, so their closed rationals can only be *bracketed* at present:

f	e	μ	τ	u	d	s	c	b	t
k	5	3	2	4	4	3	2	2	0
c (closed/cand.)	$\frac{16}{7}$	$\frac{4}{3}$	$\frac{7}{6}$	$\frac{1}{2}$	$\frac{16}{15}$	≈ 1.14	≈ 0.82	≈ 2.72	≈ 0.31
status	[I]	[I]	[I]	[P]	[P]	[P]	[P]	[P]	[P]

The lepton three are exact identities; the quark six are best 3-digit fits (candidates only). **What is now structurally closed.** The whole hierarchy is the *integer* word-length ladder λ_Y^L (compiler matrix), spanning five orders of magnitude with no tuning; all residue sits in the $O(1)$ amplitudes $\Lambda_{f,j}$. Evaluating the matrix elements of the explicit H2 holonomy (Prop. 11) gives a band $[0.07, 0.94]$ (off-diagonal $[0.07, 0.82]$) that *contains* the published quark $\Lambda \in [0.44, 0.99]$, and a bare diagonal with a single universal y *cannot* reproduce the table (residual 1.5) — so the quark amplitudes are genuinely the off-diagonal $U^\star_f$. Closing the last digits is then a finite $U^\star_f$ evaluation at δ_{ph} on the already-fixed connection — no new physics. **[I]** (structure) / **[P]** (digits)

The exact off-diagonal form, and why a bare resolvent is not enough

The C_6 hexagon resolvent has a closed off-diagonal element depending only on the residue class $k = (i{-}j) \bmod 6$:

$$\boxed{(U^\star_f)_k = \langle e_i|(y - \delta U_6)^{-1}|e_j\rangle = \frac{y^{\,5-k}\,\delta^{\,k}}{y^6 - \delta^6}},$$

with $z_\star = \delta^6 = 0.043606$ the lower critical point of the cusp cubic $P'{=}0$ (so $\delta_{\mathrm{ph}} = z_\star^{1/6} \approx 0.593$). This is the *exact* structure of the holonomy amplitude. The clean negative is now sharp: with a *single* external leg y, the form $(U^\star_f)_{r_f}$ cannot match even the three lepton amplitudes $\Lambda_{e,\mu,\tau} \approx (0.48, 1.11, 0.92)$ simultaneously — so each sector carries its *own* external leg ($y_\ell, y_{\mathrm{up}}, y_{\mathrm{down}}$, the sector Yukawa normalisations). The residue classes r_f are fixed by the compiler matrix; the one residual is the per-sector $y_{\mathrm{up}}, y_{\mathrm{down}}$, not the resolvent. The quark mass *ratios* are closed independently (integer Plücker, v49/v71); only the *absolute* per-sector scale $y_{\mathrm{up}}, y_{\mathrm{down}}$ remains — a named, bounded normalisation (an anchor), not new physics. **[A]**

26 H2: the moduli computation, set up explicitly

H2 (Theorem 4 of the parabolic note) is reduced to: *the D_4-equivariant stable parabolic structure on $\mathcal{O}(-2) \oplus \mathcal{O}(-1)^2$ realises the spectral-selector branch.* We set up the finite computation and verify the selector.

Proposition 11 (Moduli setup and selector verification)**.** By Mehta–Seshadri the stable parabolic bundles correspond to irreducible unitary representations $\rho : \pi_1(\mathbb{P}^1 \setminus \mu_4) \to U(N_{\mathrm{fam}})$ with the four local monodromies in the order-3 cusp class (eigenvalues $\{1, \omega, \omega^2\}$, a *regular* semisimple class, $\dim C = \dim G - \operatorname{rk} G = 6$). The relative character variety has real dimension

$$\dim \mathcal{M} = n \dim C - 2 \dim G = 4 \cdot 6 - 2 \cdot 8 = 8,$$

(cross-check $SU(2)$, 4 punctures: $4 \cdot 2 - 2 \cdot 3 = 2$, the Painlevé VI surface). The μ_4 deck D_4 acts on $\mathcal{M}$; its fixed locus carries the family connection $\nabla^\star_F$. On that locus the residue matrix is forced to the canonical spectrum

$$\det R = h(D_5) = 8, \qquad \mathrm{PrinMin}_2(R) = (2, 3, 5), \qquad \chi_R = t^3 - N_{\mathrm{fam}}^2 t^2 + \mathcal{A}_\Lambda t - h(D_5),$$

and the only other D_6-admissible down-branch, $\{1, 3, 4\}$, fails it ($\det = 4$, minors $(-11, 0, 3)$).

H2 closed by two counting facts [I]

The remaining "degree" step is not a computation but an identity. **(a) Splitting type.** The parabolic degree-zero condition gives $\deg E = -|\mu_4| = -4$; stability forces every generation sub-line-bundle to have negative parabolic slope, hence $\deg \ell_j \le -1$ (weights in $[0, 1)$). So the three winding magnitudes are three positive integers summing to 4, and

$$4 = 1 + 1 + 2 \text{ is the } \textit{unique} \text{ partition of } |\mu_4| \text{ into three positive parts,}$$

so $E = \mathcal{O}(-2) \oplus \mathcal{O}(-1)^2$ and $a = (1, 1, 2)$. **(b) Residue branch.** An explicit D_4-symmetric $SU(3)$ monodromy (each puncture carrying the cusp triple $\{1, \omega, \omega^2\}$, product $= \mathbf{1}$ at residual 10^{-11}) realises $r \bmod 3$; among the two D_6-admissible branches exactly one has $\operatorname{coker} R = \mathbb{Z}_{h(D_5)} = \mathbb{Z}_8$ ($\det R = 8$, SNF $= (1, 1, 8)$), the wrong sibling $\{1, 3, 4\}$ giving $\det = 4$. The carrier Coxeter number $h(D_5) = 8$ selects R uniquely. Thus $L = 6n + r$ is fully determined — no search, two identities. [I]

The $U_f^\star$ holonomy reduction: the exact quark c's reduce to (U) [I]/[A]

The last quark step — the *exact rational* amplitudes $\Lambda_{f,j}$ (the analog of the lepton $\frac{16}{7}, \frac{4}{3}, \frac{7}{6}$) — is now carried to its true core ([`verification/v19_monodromy_moduli.py`]):

1. **Exact pole.** $\delta_{\rm ph} = z_\star^{1/6}$, $z_\star = \dfrac{794 - 7\sqrt{9961}}{2187}$, $2187 = 3^7$, $\delta_{\rm ph} \in (\frac{1}{3}, \frac{2}{3})$.
2. **Clean negative.** The *bare* finite resolvent at the natural cusp leg $y{=}1$ misses the lepton ratios by 3×–11×, so the $SU(3)_F$ holonomy dressing is essential.
3. **Required amplitudes.** The lepton c's force $\Lambda_{e,\mu,\tau} = \pi c\,(\varphi_0^{\rm ret})^n/\lambda_Y^L = (0.4751, 1.1073, 0.9173)$ (the published three-digit table $(0.466, 1.085, 0.899)$ matches to 2%); since $\Lambda_\mu{>}1$, Λ is the *non-unitary* resolvent Green function $y^{5-r}\delta^r/(y^6{-}\delta^6)$ dressed by the holonomy phase, not a unitary matrix entry.
4. **Moduli result.** The D_4-symmetric $SU(3)$ monodromy (each puncture in the cusp class $\{1, \omega, \omega^2\}$, product $= \mathbf{1}$) is explicitly constructible ($\|M_1M_2M_3M_4 - \mathbf{1}\| < 10^{-12}$), but the D_4-fixed locus is *positive-dimensional* (the conjugation invariant $|\operatorname{tr} M_1M_2|$ varies continuously). Hence $\nabla_F^\star$ — and with it the exact $\Lambda_{f,j}$ — is *not* pinned by D_4-symmetry + cusp class + product alone.
5. **Why it is a wall (sharpening).** The cusp weights $\{0, \frac{1}{3}, \frac{2}{3}\}$ (sum 1) at the four punctures give total parabolic weight $\sum_i w_i = 4$, so strict stability of the split $\mathcal{O}(-2) \oplus \mathcal{O}(-1)^2$ would need $w_1{<}2$ (the $\mathcal{O}(-2)$) and $w_2, w_3{<}1$ (the $\mathcal{O}(-1)$'s) — but $w_1 = 4 - w_2 - w_3 > 2$. The bounds are jointly *unsatisfiable*: the configuration sits on the parabolic *stability wall*, so $\nabla_F^\star$ is the unique *polystable* representative there. This is exactly why the D_4-fixed locus is positive-dimensional, and it pins (U) down to a wall-selection. [`verification/v26_flavor_frontier_unification.py`]

Update — the ratios no longer reduce to this gate. Since v49/v69/v70/v71 the geometric origin of Q ([`verification/v69_d4_q_geometry.py`]), the lattice selector $\det R = 8$ ([`verification/v70_q_integer_lift.py`]) and the rigidity of $c_u/c_d = \frac{55}{117}$ on the derived stratum ([`verification/v49_readout_rigidity.py`]) are established. So the *ratios* and the parabolic↔transport equivalence are *closed combinatorially*; the wall-selection (U) is now needed *only* for the *absolute* amplitude normalisation $U_f^\star$ (an anchor). The flavor frontier is a single *absolute-scale* anchor, not an open computation of the ratios.

The (U_{wall}) gate, explicitly: one D_4-fixed polystable wall point [P]

An explicit balanced wall representative (weights in units of $\frac{1}{3}$, rows = the lines $\mathcal{O}(-2), \mathcal{O}(-1), \mathcal{O}(-1)$, columns = the four punctures) is

$$W_{\mathrm{wall}} = \frac{1}{3}\begin{pmatrix} 2 & 1 & 2 & 1 \\ 1 & 0 & 0 & 2 \\ 0 & 2 & 1 & 0 \end{pmatrix},$$

each column a permutation of the cusp triple $(0,1,2)$, the row sums $w = (2,1,1)$ exactly lifting the degrees $(-2,-1,-1)$ to parabolic degree 0 — the wall point. The physical $\nabla^\star_F$ is the unique D_4-fixed *polystable* point selected on this wall by

$$\boxed{\nabla^\star_F = \mathrm{Sel}_{D_4,\,\mathrm{cusp},\,\det R=8,\,Q_+}\big(\mathcal{M}_{\mathrm{par}}(X,E,\alpha)\big)},$$

$$\det R = 8, \quad \mathrm{SNF}(R) = (1,1,8), \quad \mathrm{Spec}(Q_+) = \{1,2,3\}.$$

The discrete data R, Q, $\det R = 8$, $\mathrm{Spec}(Q_+)$ on this point are now *derived* (v69/v70) and fix the *ratios* c_u/c_d and the H2 transport equivalence by rigidity (v49/v71); solving the finite D_4-fixed parabolic Hitchin point is needed only for the *absolute* $U^\star_f, c_q$ normalisation. It is not pleasant, but finite. [`verification/v27_wall_representative.py`]

Conclusion. The exact quark mass *ratios* are not blocked at all — they are integer Plücker readouts on a now-derived selector stratum (v49/v69/v71). Only the *absolute* c_q scale reduces to selecting the stable point of a positive-dimensional character variety, i.e. to the unitarity/stability datum (U) already declared as a structural input — an anchor. Given (U), the diagonal amplitudes — hence all nine c's — follow. [I] (reduction) / [A] (the (U) selection)

Status: what is now written out, and what is honestly left

Item	Status	Result
Dynamics / clustering (Lemma 2)	[I]	explicit gap $6\log\frac{3}{2}$, Dobrushin unique state
Gravity Fredholm (Lemma 3)	[I]	APS on S^2; index $(2,3)$
$R^{n\geq 3}$ truncation (Lemma 4)	[I]	$O(1/N^2)$; n_s=0.964, r=0.004
Starobinsky M / A_s	[I]	$M{=}c_3^{7/2}\bar{M}_{\mathrm{Pl}}{=}3.06{\times}10^{13}$ GeV (canonical); A_s predicted
Λ typing (Prop. 9)	[P]	three objects linked by $(8\pi)^2$; branch open
θ_{12} (Prop. 10)	[N]/[P]	cond. derived $\frac{1}{3} - \frac{\varphi_0^{\mathrm{ret}}}{2}{=}0.307$; rests on seam-misalignment $\varepsilon{=}c_3$
quark c-rationals	[I]/[A]	hierarchy + source ratios closed; exact digits reduce to (U) ([`verification/v19`])
H2 splitting type (Prop. 11)	[I]	4=1+1+2 unique; branch coker $\mathbb{Z}_8$
unitarity (U) of $\nabla^\star_F$	[A]	selects the stable point of a *positive-dim* character variety ([`verification/v19`])

Deliberately not forced

η_B, m_p/m_e, exact Koide (only near), the dark-matter *relic scale* $(f_a, m_a$; the axion candidate itself is fixed), and full quantum gravity beyond the R^2 scalaron sector are *not* on a clean action/ladder yet and are left open rather than pressed into the compiler.

Net. The previously “conditional” steps are now written out: gap/clustering, APS Fredholm and the truncation bound as lemmas [I]; the Starobinsky scalaron mass fixed by c_3^7 (so A_s, n_s, r are predictions); θ_{12} *derived* from the seam $(\frac{1}{3}-\frac{\varphi_0^{\mathrm{ret}}}{2})$; the H2 splitting type closed by the unique partition 4=1+1+2 and the residue branch by coker $=\mathbb{Z}_8$; and the quark hierarchy closed with the $O(1)$ amplitude band computed. The Λ ambiguity is *typed* away to a single branch label. What is deliberately left open (η_B, m_p/m_e, exact Koide, dark matter, full quantum gravity) is in `frontier_items`. **No hidden finite arithmetic gap remains in the compiled discrete core**; the remaining analytic gates — (U_{wall}), G_{metric} and the frontier interface calculations — are explicitly typed, not concealed.

Appendix: computational verification

The Standard-Model identities used here (marked [I], [L], [N]) are re-derived from $\{c_3, g_{\mathrm{car}}\}$ and machine-checked by the suite in the repository folder `verification/`; the inline tags `[verification/...]` point to the exact script. Relevant here:

- `v2_carrier_pascal.py` — the carrier/Pascal packet (16=1+5+10, $N_{\mathrm{fam}}, \Omega_{\mathrm{adm}}, b_1$).
- `v3_em_alpha.py` — the EM fixed point $\alpha_\star^{-1}$=137.0359992168.
- `v4_flavor_matrix.py` — the residue matrix R (det =8, minors $\{2,3,5\}$, χ_R, SNF, $\sum L$=40).
- `v7_gravity_cosmo.py` — scalaron exponent 7 and inflation A_s, n_s, r; $\sin^2\theta_{12}$.

Run `python verification/run_all.py` (needs `mpmath, numpy, sympy`); exit code 0 iff all checks pass. Full map: `verification/README.md`.