TFPT — E_8 Audit, Cascade Bridge and Bootstrap

The seven E_8 slices as an audit raster, the old cascade as the same even-integer spine, and the Möbius self-consistency loop

June 8, 2026

What this document is about

E_8 as an *audit container*, not a mystery: the seven maximal slices of 248 as a falsification raster (every load-bearing number must appear in at least one projection), the bridge showing the old E_8 orbit cascade $D=60-2n$ is the same even-integer spine as the compiler, and the Möbius bootstrap in which $g_{\text{car}} = 5$ and the "8" in c_3 are overdetermined E_8-closure fixed points (only π irreducible).

What E_8 is *not* (and why the known no-go results do not apply)

E_8 here is the unimodular *audit / compiler hull* that classifies the admissible discrete charge and residue structures — **not** an unbroken physical gauge group. The Standard Model is a *readout* after projection (not a direct E_8 gauge theory), the Lorentz signature appears only after projection, and fermions are not "everything in the adjoint". Consequently the standard objections to literal E_8 world-formulas do *not* bite: Distler–Garibaldi (no chiral fermions / representation obstructions for E_8 as a 4d gauge group) and Coleman–Mandula (no nontrivial mixing of spacetime and internal symmetry) both constrain E_8 as a *spacetime gauge symmetry*, which TFPT does not claim. The defensible statement is: *TFPT is the unique parabolic compilation of the anchor* $a = (1, 1, 2)$ *into the* E_8 *audit hull* — a discrete classifier, attacked on its own (combinatorial) terms, not on gauge-theoretic ones.

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 #4.

Contents

Part I

The E_8 secondary-branching atlas

Purpose and discipline

This note does *not* introduce new physics. It builds an *audit raster*: the claim that *every load-bearing TFPT number should appear in at least one E_8 branching projection.* A number that appears in none is suspect; a number appearing in several independent projections becomes structurally strong. All arithmetic identities below are exact [I]; the *physical readings* are audit-level [A] unless tied to a derivation. (All 77 numeric claims were machine-checked with exact integer/rational arithmetic.)

1 The compiler core (where this note sits)

The strongest current form of the theory is a five-layer compiler:

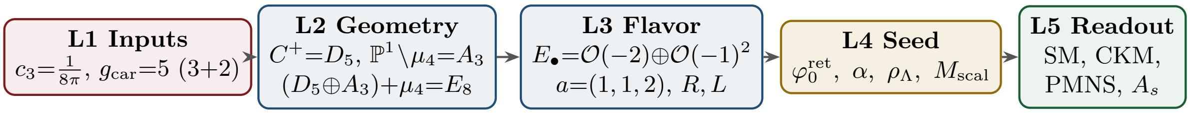

E_8 is now a *closure* object, not an input: the carrier is the $D_5 = \mathfrak{so}_{10}$ half-spinor, the four punctures $\mathbb{P}^1 \setminus \mu_4$ give $A_3 = \mathfrak{su}_4$, and the common discriminant group $\mathbb{Z}_4$ glues them, with the norm test $q(D_5) + q(A_3) = \frac{5}{4} + \frac{3}{4} = 2$. This note charts the *secondary* maximal subalgebras of E_8 as audit windows onto the TFPT modules.

2 The $(1, 1, 2)$ anchor microcode

The parabolic anchor $a = (1, 1, 2)$ (splitting type of $E_\bullet = \mathcal{O}(-2) \oplus \mathcal{O}(-1)^2$) is a miniature compiler. Its power sums $p_n(a) = 1^n + 1^n + 2^n = 2 + 2^n$ generate the discrete budget:

n	$p_n = 2 + 2^n$	**TFPT reading (verified)**
1	4	$\lvert\mu_4\rvert = -\deg E_\bullet$ (glue index)
2	6	$\lvert R^+(A_3)\rvert$ (positive A_3 roots) $= \mathbb{Z}_6$ cusp order
3	10	$\mathcal{A}_\Lambda = \binom{5}{2} = \dim V_{D_5}$ (action ladder / pair sector)
4	18	$N_{\mathrm{fam}}\,\lvert R^+(A_3)\rvert = 3 \cdot 6$
5	34	$2h(D_5) + 3\lvert R^+(A_3)\rvert = 16 + 18$; and $34 - \lvert R(A_3)\rvert = 22 = \sum R$

and the budget compresses to products of power sums:

$$q(A_3) = \frac{p_2}{2p_1} = \frac{3}{4}, \quad q(D_5) = \frac{p_3}{2p_1} = \frac{5}{4}, \quad \sum L = p_1 p_3 = 40,$$

$$\Omega_{\mathrm{adm}} = 2p_1p_2 = 48, \quad D_{\mathrm{start}} = p_2p_3 = 60, \quad \lvert R(E_8)\rvert = 4p_2p_3 = 240.$$

The first three power sums alone reach E_8: $p_1p_2p_3 = 4{\cdot}6{\cdot}10 = 240 = \lvert R(E_8)\rvert$, and $p_1p_2p_3{+}(p_4{-}p_3) = 240 + 8 = 248 = \dim E_8$.

The anchor difference ladder: the binary carrier spine [I]

While the *products* of the power sums give the large budgets, the *differences* give the binary carrier spine in one line:

$$\boxed{p_n(a) = 2 + 2^n \implies \Delta p_n = p_{n+1} - p_n = 2^n} : \quad (2, 4, 8, 16) = (|\mathbb{Z}_2|,\ |\mu_4|,\ h(D_5),\ \dim S^+).$$

So the single anchor $a = (1, 1, 2)$ generates the big numbers multiplicatively and the doubling spine $2, 4, 8, 16$ additively — one vector, two readings.

The anchor reconstructs the carrier rank — and the 41-chain

The characteristic polynomial of the anchor vector is $(t-1)^2(t-2) = t^3 - 4t^2 + 5t - 2$, with elementary symmetric data

$$e_1(a) = 4 = |\mu_4|, \qquad e_2(a) = 5 = g_{\mathrm{car}}, \qquad e_3(a) = 2 = |\mathbb{Z}_2|_{\mathrm{sheet}}.$$

So the parabolic flavor geometry reads the carrier rank $g_{\mathrm{car}} = 5$ *backwards*: $D_5 \Rightarrow a$ but also $a \Rightarrow g_{\mathrm{car}}$. Moreover $3^2 + 4^2 = 5^2$ sits as $(p_0, e_1, e_2) = (3, 4, 5)$, giving the *anchor form* of the abelian index

$$\boxed{10\, b_1 = e_1(a)^2 + e_2(a)^2 = 4^2 + 5^2 = 41}$$

alongside the known $41 = 40 + 1 = \sum L + N_\Phi$. [I]

Pascal closure: $g_{\mathrm{car}} = 5$ is the unique solution (kills alternatives) [I]

The carrier rank is not merely a good fit — it is the *unique* root of an exact closure. The even-exterior carrier dimension 2^{g-1} must equal the truncated Pascal sum $\binom{g}{0} + \binom{g}{1} + \binom{g}{2}$ (the $1+g+\binom{g}{2}$ that builds $\dim S^+$):

$$\boxed{2^{\,g-1} = \binom{g}{0} + \binom{g}{1} + \binom{g}{2} \iff g = 5}$$

(g=3 : 4≠7; g=4 : 8≠11; **g=5 : 16=16**; g=6 : 32≠22; g=7 : 64≠29). At $g = 5$ the row is $\binom{5}{0}, \binom{5}{1}, \binom{5}{2} = 1, 5, 10$, the same Pascal triple that reads family, action ladder and the E_8 root count $16 \cdot 5 \cdot 3 = 240$. This is the fourth independent lock on $g_{\mathrm{car}} = 5$ (with rank-fill, Coxeter-match, and integer-glue/norm), and unlike those it explicitly *kills* the neighbours $g = 4, 6$.

The seed is an anchor expression

The retained seed is exactly the two-term anchor combination of c_3:

$$\boxed{\varphi_0^{\rm ret} = \frac{2p_1}{p_2}\,c_3 + 2p_1p_2\,c_3^4 = \frac{4}{3}c_3 + 48\,c_3^4 = \frac{1}{6\pi} + \frac{3}{256\pi^4}}$$

(verified to 30 digits). The Cabibbo base is $\lambda_Y = \sqrt{\varphi_0^{\rm ret}(1-\varphi_0^{\rm ret})}$, the hierarchy engine of the whole fermion spectrum. [I]

3 The residue and length matrices as spectral objects

With the residue matrix R and the full word-length matrix $L = R + 6W$ (W the rank-one winding),

$$R = \begin{pmatrix} 1 & 3 & 0 \\ 1 & 5 & 2 \\ 2 & 5 & 3 \end{pmatrix}, \qquad L = \begin{pmatrix} 7 & 3 & 0 \\ 7 & 5 & 2 \\ 8 & 5 & 3 \end{pmatrix}, \qquad W = \begin{pmatrix} 1 & 0 & 0 \\ 1 & 0 & 0 \\ 1 & 0 & 0 \end{pmatrix},$$

the spectral invariants are pure compiler numbers (all verified):

Invariant of R	**value**	**Invariant of L**	**value**
$\operatorname{tr} R$	$9 = N_{\rm fam}^2$	$\operatorname{tr} L$	15
$\det R$	$8 = h(D_5)$	$\det L$	$20 = \lvert R(A_4)\rvert = 2\mathcal{A}_\Lambda$
$\chi_R(t)$	$t^3 - 9t^2 + 10t - 8$	$\chi_L(t)$	$t^3 - 15t^2 + 40t - 20$
$\mathrm{PrinMin}_2(R)$	$(2,3,5)$	$\mathrm{PrinMin}_2(L)$	$(5,14,21)$
$\lVert R\rVert_F^2$	$78 = \dim E_6$	—	—

The three principal 2×2 minors of R are $(2,3,5)$ — sheet, family, carrier — with product $2\cdot3\cdot5 = 30 = h(E_8)$ and sum $10 = \mathcal{A}_\Lambda$. The wrong down-branch $\{1,3,4\}$ fails exactly these invariants.

Bilinear anchor forms. With $\mathbf{1} = (1,1,1)^\top$ and $a = (1,1,2)^\top$ (all verified):

R	$\mathbf{1}$	a	L	$\mathbf{1}$	a
$\mathbf{1}^\top$	$22 = \sum R$	$27 = 3^3$	$\mathbf{1}^\top$	$40 = \lvert R(D_5)\rvert$	$45 = \dim D_5$
$a^\top$	$32 = 2^5$	$40 = \lvert R(D_5)\rvert$	$a^\top$	$56 = \dim \mathbf{56}_{E_7}$	$64 = \dim S^+ \cdot \lvert\mu_4\rvert$

The corner value $a^\top L\,a = 64 = (16,4)$ reads exactly one spinor-glue sheet of the E_8 branching. And the rank-one winding update shows up in the inverses:

$$\boxed{R^{-1}\mathbf{1} = \tfrac{1}{4}(1,1,-1)^\top, \qquad L^{-1}\mathbf{1} = \tfrac{1}{10}(1,1,-1)^\top}$$

i.e. the residue response reads the glue denominator $\lvert\mu_4\rvert = 4$, the full-length response reads the decuple denominator $\mathcal{A}_\Lambda = 10$. [I]

Anchor-plane Plücker coordinates: scattered hits become oriented areas [I]

The strongest anti-numerology tool is to read *orientation invariants* of the anchor plane $\langle \mathbf{1}, a\rangle$ rather than loose integers. For $M \in \{R, Q, K, L\}$ the left block $B_M = (\mathbf{1}^\top M;\, a^\top M)$

(a 2×3) and the right block $C_M = (M\mathbf{1}\;\; Ma)$ (a 3×2) have the 2×2 Plücker minors

	$\mathrm{Pl}(\cdot)$ (left)	$\mathrm{Pl}_R(\cdot)$ (right)
R	$(-6, 2, 14)$	$(8, 12, 4) = (h(D_5), \lvert R(A_3)\rvert, \lvert\mu_4\rvert)$
Q	$(3, 6, 3)$	$(0, 4, 5) = (0, \lvert\mu_4\rvert, g_{\mathrm{car}})$
K	$(-1, 6, 4) = (-N_\Phi, \lvert R^+(A_3)\rvert, \lvert\mu_4\rvert)$	$(12, 12, -2) = (\lvert R(A_3)\rvert, \lvert R(A_3)\rvert, -\lvert\mathbb{Z}_2\rvert)$
L	$(6, 26, 14)$	$(20, 30, 10) = (2\mathcal{A}_\Lambda, h(E_8), \mathcal{A}_\Lambda)$

Three load-bearing flavor numbers are now single oriented areas, not coincidences:

$$\boxed{\|\,\mathrm{Pl}(K)\|_1 = 11, \qquad \|\,\mathrm{Pl}_R(K)\|_1 = \mathrm{Pl}(L)_{13} = 26, \qquad \sum \mathrm{Pl}_R(L) = 60 = D_{\mathrm{start}}}$$

(11 is the mass-power anchor-plane norm; 26 the c_t/c_b denominator; 60 the E_8 cascade start). The left-null-space response complements the right one above: the column sums of $\det L \cdot L^{-1}$ and $\det R \cdot R^{-1}$ are $(-5, 1, 6)$ and $(1, -5, 6)$, the same magnitudes $\{N_\Phi, g_{\mathrm{car}}, |R^+(A_3)|\} = \{1, 5, 6\}$. Finally the sector rows themselves are anchor-shifted: $K_e - K_u = (1, 1, 2) = a$ and $L_e - L_u = (1, 2, 3) = \mathrm{Exp}(A_3)$ — leptons are up-quarks displaced by the family roots. All entries are exact [I]; this is an audit layer, not a closure (the physical quark amplitudes stay [P]). [`verification/v37_plucker_anchor.py`]

4 The secondary branching atlas

Each maximal subalgebra of E_8 mirrors a TFPT module. The dimension identities are exact [I]; the readings are audit-level [A]. The seven *slices* of $\dim E_8 = 248$ are shown below: the same total cut seven ways, each cut a different TFPT module.

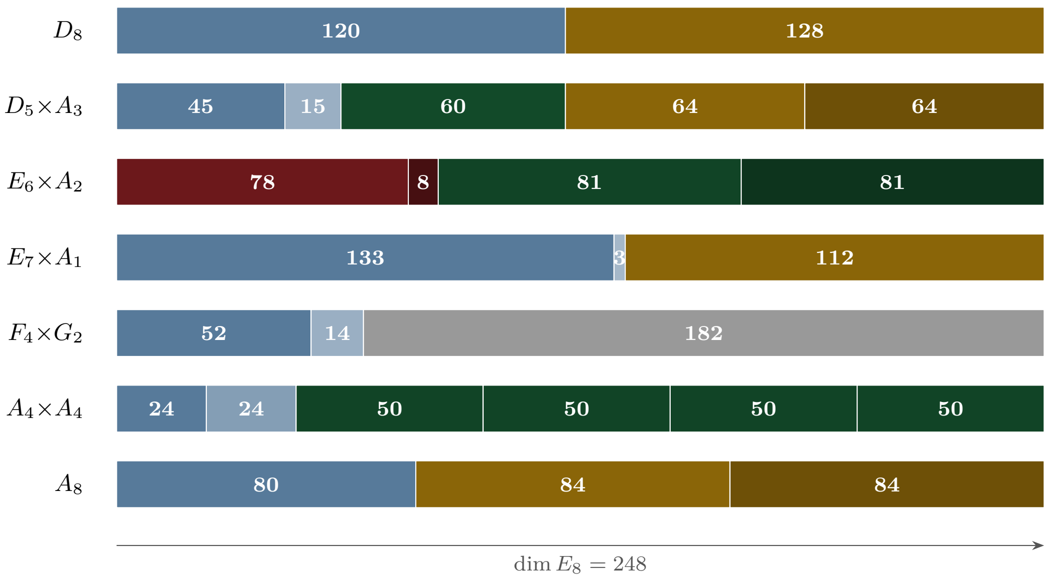

Figure 1: The E_8 slice atlas: seven maximal-subalgebra cuts of 248. Gold = spinor-glue (64=(16, 4), 112=7·16, 84=7·12, 128); red = flavor-residue ($78 {=} \|R\|_F^2$, $8 {=} \det R$); green = family/action blocks; blue = Lie adjoints.

Subalgebra	$248 =$	**TFPT reading**
D_8	$120 + 128$	family-transition adjoint + two spinor-glue sheets ($128{=}2\dim S^+\lvert\mu_4\rvert$)
$D_5 \times A_3$	$45 + 15 + 60 + 64 + 64$	**main split**: $\dim D_5$, $\dim A_3$, $\mathcal{A}_\Lambda\lvert R^+(A_3)\rvert$, two $(16,4)$ glue sheets
$E_6 \times A_2$	$78 + 8 + 81 + 81$	**flavor residue**: $\|R\|_F^2$, $\det R{=}\dim A_2$, cubic family blocks
$E_7 \times A_1$	$133 + 3 + 112$	**scalaron**: $112{=}7 \cdot 16$ (exponent × one family)
$F_4 \times G_2$	$52 + 14 + 182$	**occupancy/gauge**: $\lvert R(F_4)\rvert{=}48{=}\Omega_{\mathrm{adm}}$, $\lvert R(G_2)\rvert{=}12{=}\dim \mathfrak{g}_{\mathrm{SM}}$
$A_4 \times A_4$	$24 + 24 + 4 \cdot 50$	**action ladder**: $24{+}24{=}48{=}\Omega_{\mathrm{adm}}$; off-diagonal $(5,10)$ reps $\mathcal{A}_H, \mathcal{A}_\Lambda$
A_8	$80 + 84 + 84$	rank-decuple + scalaron: $80{=}8 \cdot 10$, $84{=}7 \cdot 12$

Theorem 1 ($E_6 \times A_2$ flavor-residue shadow [A])**.** The principal flavor branching reads the residue matrix:

$$\|R\|_F^2 = 78 = \dim E_6, \qquad \det R = 8 = \dim A_2 = h(D_5), \qquad \mathbf{1}^\top R\, a = 27,$$

and the full $E_6 \times A_2$ dimension count is the residue identity

$$\boxed{248 = \|R\|_F^2 + \det R + 2\,(\mathbf{1}^\top R\, a)\, N_{\mathrm{fam}} = 78 + 8 + 2 \cdot 27 \cdot 3.}$$

Here E_6 reads the Frobenius norm of the residue matrix, A_2 the pure three-family symmetry, and the **27** the cubic family block. (Strongest new E_8 flavor fingerprint.)

Theorem 2 ($E_7 \times A_1$ scalaron shadow [A])**.** With the scalaron exponent $7 = \Omega_{\mathrm{adm}} - 10b_1 = -\deg E + \mathrm{rk}\, E$,

$$248 = 133 + 3 + 112, \qquad 112 = 7 \cdot \dim S^+ = 7 \cdot 16, \qquad |R(E_7)| = 126 = 7 \cdot 18,$$

where $18 = N_{\mathrm{fam}} |R^+(A_3)|$. The scalaron exponent 7 (the seam power $M_{\mathrm{scal}}^2/\bar{M}_{\mathrm{Pl}}^2 = c_3^7$) is thus exactly the E_7 off-block coefficient. *Cross-check:* the centred E_8 exponents give $\sum_{m \in \mathrm{Exp}(E_8)} |m - 15| = 56 = \dim \mathbf{56}_{E_7}$ (exponents $\{1, 7, 11, 13, 17, 19, 23, 29\}$, the primitive residues mod 30).

Proposition 1 ($A_4 \times A_4$ action-ladder and $F_4 \times G_2$ occupancy [A])**.** $248 = 24 + 24 + 4(5 \cdot 10)$ with $24 + 24 = 48 = \Omega_{\mathrm{adm}}$ and off-diagonal reps $(\mathbf{5}, \mathbf{10}) = (\mathcal{A}_H, \mathcal{A}_\Lambda)$ — the action ladder appears inside E_8 (an audit, *not* an $SU(5)$ GUT). And $F_4 \times G_2$: $|R(F_4)| = 48 = \Omega_{\mathrm{adm}}$, $|R(G_2)| = 12 = \dim \mathfrak{g}_{\mathrm{SM}}$, $\mathrm{rank}\, F_4 = 4 = |\mu_4|$, $\mathrm{rank}\, G_2 = 2 = |\mathbb{Z}_2|$.

Flavor-matrix and prime-table fingerprints (audit-level, machine-verified)

Five exact linear-algebra readouts of R, L and the prime atoms, kept here under the no-free-pattern rule.

- **Row norms ladder.** $\|R_u\|^2, \|R_d\|^2, \|R_e\|^2 = (10, 30, 38) = (\mathcal{A}_\Lambda,\ h(E_8),\ 2 \cdot 19)$ with $10{+}30{+}38 = 78 = \dim E_6$ $(19 \in \mathrm{Exp}(E_8))$.
- **Residue inventory.** The entries of R on C_6 are $m = (1, 2, 2, 2, 0, 2)$: the *only* gap is $4 = |\mu_4| = \sum a$. Moments $\sum R = 22$, $\sum R^2 = 78 = \dim E_6$ — so $\|R\|_F^2 = 78$ is the second moment of the C_6 inventory with the μ_4 hole.
- **First-column norm → $E_6{\times}A_2$ off-blocks.** $L_{\cdot 1} = (7, 7, 8)$, $\|L_{\cdot 1}\|^2 = 162 = 81 + 81$ — the two **81** off-blocks of the $E_6{\times}A_2$ slice are exactly the wound first-generation column.
- **Lepton exponent polynomial.** The φ_0^{ret}-ladder lepton powers $(k_e, k_\mu, k_\tau) = (5, 3, 2)$ are the Coxeter primes $30 = 2 \cdot 3 \cdot 5$ in hierarchy order: $(t-5)(t-3)(t-2) =$

$t^3{-}10t^2{+}31t{-}30$ ($e_1 = \mathcal{A}_\Lambda$, $e_3 = h(E_8)$). The lepton c-ratios obey $\frac{(16/7)(7/6)}{4/3} = 2 = |\mathbb{Z}_2|$ and $\frac{(16/7)(7/6)}{(4/3)^2} = \frac{3}{2}$ (the near-Koide factor).

- **The $2, 3, 5$ multiplication table** (the shortest integer-landscape compression):

$$(2, 3, 5) \xrightarrow{\times} (6, 10, 15, 30) \to (8{=}3{+}5,\ 16,\ 240{=}8 \cdot 30,\ 248{=}240{+}8).$$

5 Group-theoretic audits (optional)

Weyl completion index

Over the main stabiliser,

$$\frac{|W(E_8)|}{|W(D_5)|\,|W(A_3)|} = \frac{696729600}{1920 \cdot 24} = 15120 = 63 \cdot 240 = (2^6 - 1)\,|R(E_8)|.$$

A group-theory machine-check, not physics: it quantifies the Weyl symmetry created beyond the $D_5 \times A_3$ stabiliser. [I] (arithmetic) / [A] (reading)

E_8 theta series as a higher-shell raster

$\Theta_{E_8} = E_4 = 1 + 240 \sum_{n \geq 1} \sigma_3(n) q^n$ has shell degeneracies $240,\ 2160,\ 6720,\ 17520, \ldots$ The first is $|R(E_8)| = 240$; the second is $2160 = 9 \cdot 240 = N_{\text{fam}}^2 |R(E_8)|$. If later one-loop or CMB shell degeneracies touch these multiplicities, the theta series is the right place to type them — as an audit raster, not a claim. [A]

6 Further audit lemmas (machine-verified this round)

Five more exact identities tie the flavor matrices and the E_8 exponents to the compiler. All checked with exact integer arithmetic.

(i) Full-length Frobenius norm is three E_6 blocks

$$\|L\|_F^2 = 234 = N_{\text{fam}}\,\|R\|_F^2 = 3 \dim E_6, \qquad \|L\|_F^2 = \|R\|_F^2{+}12\langle R, W\rangle{+}36\|W\|_F^2 = 78{+}48{+}108,$$

with $\langle R, W\rangle = 4 = |\mu_4|$ and $\|W\|_F^2 = 3 = N_{\text{fam}}$. The full word-length load is exactly the residue E_6 norm + the A_3-winding over μ_4 + the family-winding square. [I]

(ii) A_3-winding determinant lemma

The rank-one winding $L = R + 6\,\mathbf{1}e_1^\top$ and the matrix-determinant lemma with $e_1^\top R^{-1}\mathbf{1} = \frac{1}{4}$ give

$$\det L = \det R\Big(1 + \tfrac{|R^+(A_3)|}{|\mu_4|}\Big) = 8\Big(1 + \tfrac{6}{4}\Big) = 8 \cdot \tfrac{5}{2} = 20 = 2\mathcal{A}_\Lambda.$$

So $h(D_5) = 8 \xrightarrow{A_3^+/\mu_4} 20 = |R(A_4)|$: the winding injects the carrier factor 5 into the cokernel ($\operatorname{coker} R \simeq \mathbb{Z}_8$, $\operatorname{coker} L \simeq \mathbb{Z}_{20}$). [I]

(iii) The scalaron 7 is the E_8 exponent variance per decuple

Centring the exponents at $h/2 = 15$,

$$\sum_m (m-15)^2 = 560, \quad \mathrm{Var} = \tfrac{560}{8} = 70 = 7\cdot\mathcal{A}_\Lambda = (\text{scalaron } 7)\cdot(\text{decuple } 10).$$

The halved positive offsets $\{1,2,4,7\}$ satisfy $\sum = 14 = 2\cdot 7$, $\sum{}^2 = 70$, $\prod = 56 = \dim \mathbf{56}_{E_7}$ — a micro-code carrying the scalaron, the decuple and the E_7 minuscule at once. Also $\sum_{\text{pairs}} m(30-m) = 620 = \frac{h\dim E_8}{12} = 2\cdot 10\cdot 31$. [I]

(iv) Exponents + degrees give the 120+128 split

The invariant degrees $d_i = m_i + 1 = \{2,8,12,14,18,20,24,30\}$ obey

$$\sum_i m_i = 120, \qquad \sum_i d_i = 128, \qquad 248 = 120+128, \qquad \prod_i d_i = |W(E_8)|,$$

matching the D_8 branching $248 = 120 + 128$; the degrees are a TFPT dictionary $(2,8,12,14,18,20,24,30 = \text{sheet},\ h(D_5),\ |R(A_3)|,\ \dim G_2,\ p_4,\ \det L,\ \dim A_4,\ h(E_8))$. [I]

(v) Root shell = local $E_7 \times A_1$

Around any E_8 root the inner-product multiplicities are $\{+2{:}1,\ +1{:}56,\ 0{:}126,\ -1{:}56,\ -2{:}1\}$, i.e. the local $248 = (133,1)+(1,3)+(56,2)$. A distinguished root (seam/Higgs direction) leaves an orthogonal E_7 and an active coupling shell 56+56. [I]

7 What this changes, and what stays open

E_8 is one container whose maximal subalgebras each project a different TFPT module:

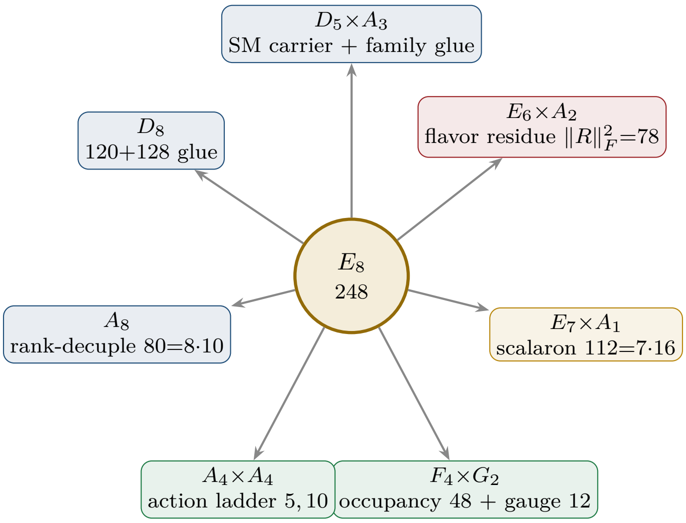

Figure 2: E_8 as an audit container: each maximal subalgebra projects one TFPT module. The discipline rule — every load-bearing number appears in ≥ 1 projection — makes the structure falsifiable.

Net

The atlas turns isolated fingerprints (e.g. $\|R\|_F^2 = 78 = \dim E_6$, the scalaron 7) into *branching projections* of one container. The discipline rule — "every load-bearing number must appear in ≥ 1 projection" — makes the structure falsifiable: a number in no projection is suspect. This is a *program*, not a proof of new physics: all readings are [A]. The two genuinely load-bearing, derived results remain the μ_4 glue ($E_8 = (D_5 \oplus A_3) + \mu_4$, [L] [`verification/v1_e8_glue.py`]) and the spectral selector of the flavor matrix ($\det R = h(D_5)$, $\mathrm{PrinMin}_2 = (2,3,5)$, [I] [`verification/v4_flavor_matrix.py`]).

The single remaining flavor-geometry step (H2) is sharpened in the parabolic note to: *the D_4-equivariant stable parabolic structure on $\mathcal{O}(-2) \oplus \mathcal{O}(-1)^2$ realises the unique branch with* $\det R = 8$ *and* $\mathrm{PrinMin}_2(R) = (2,3,5)$ — no longer "find the matrix", but "prove this one geometric realisation".

Part II

The E_8 cascade bridge

The connection in one sentence

The old E_8 cascade (Paper V1.06: the nilpotent-orbit chain $D_n = 60 - 2n$, $n = 0 \ldots 26$) and the new compiler/atlas are *the same even-integer* E_8 *spine*, read three ways: the cascade *orders* the numbers $\{8, 10, \ldots, 60\}$ as a scale ladder, the atlas *reads* them as E_8 subalgebra branchings, and the compiler *generates* them from the $(1, 1, 2)$ anchor. The cascade *starts* at $D_{\text{start}} = 60 = p_2 p_3$ and *ends* at $8 = h(D_5) = \det R$ — the flavor selector.

8 The old cascade, recalled

Paper V1.06 builds a deterministic scale ladder from the nilpotent orbits of E_8. With centralizer dimension $D = 248 - \dim \mathcal{O}$ and the Hasse chain restricted to $\Delta D = 2$, a unique maximal chain runs from $A_4{+}A_1$ ($D_0 = 60$) down to the regular orbit ($D_{26} = 8$):

$$\boxed{D_n = 60 - 2n, \quad n = 0, \ldots, 26; \qquad \varphi_{n+1} = \varphi_n e^{-\gamma(n)}, \quad \gamma(n) = \lambda \left[\ln D_n - \ln D_{n+1}\right]}$$

The cascade arithmetic (endpoints, exponent rungs, IR-tail product, variance) is machine-checked. [`verification/v5_e8_cascade.py`] with the single structural normalisation $\lambda = \gamma(0)/(\ln 248 - \ln 60) = 0.58770$, fixed by a curvature-extremum on the chain (no fit). The E-windows of the two-loop flow anchor it: $\alpha_3 = 1/(8\pi) = c_3$ (the E_8 window) and $\alpha_3 = 1/(6\pi)$ (the E_6 window, near φ_0^{ret}).

9 The spine: every rung is a current number

Reading $D = 60 - 2n$ against today's compiler atoms, the secondary-branching atlas and the $(1, 1, 2)$ anchor power sums $p_n = 2 + 2^n$, **25 of the 27 rungs are load-bearing numbers** (machine-checked):

n	D	**current meaning**	n	D	**current meaning**
0	60	$D_{\text{start}} = p_2 p_3 = \mathcal{A}_\Lambda \, \lvert R^+(A_3)\rvert$ **(start)**	13	34	p_5 anchor power $= 2 \cdot 17$
1	58	$2 \cdot 29$ (E_8 exponent)	14	32	$2^5 = 2 \dim S^+$
2	56	$\dim \mathbf{56}_{E_7} = a^\top L \mathbf{1} = \sum \lvert m{-}15\rvert$	15	30	$h(E_8) = 2{\cdot}3{\cdot}5$ **(Coxeter rung)**
3	54	$2 \cdot 27$, $27 = \mathbf{1}^\top R a = 3^3$	16	28	$4 \cdot 7$ (half of 56)
4	52	$\dim F_4 = \lvert R(D_5)\rvert {+} \lvert R(A_3)\rvert$	17	26	$2 \cdot 13$ (old hadronic block)
5	50	$A_4{\times}A_4$ block $= 5{\cdot}10$	18	24	$\dim A_4$; $24{+}24{=}48$
6	48	$\Omega_{\text{adm}} = \lvert R(F_4)\rvert = 2 p_1 p_2$	19	22	$\sum R = \mathbf{1}^\top R \mathbf{1}$ (residue sum)
7	46	$2 \cdot 23$ (E_8 exponent)	20	20	$\det L = \lvert R(A_4)\rvert = 2\mathcal{A}_\Lambda$
8	44	$4 \cdot 11$ (E_8 exponent)	21	18	$p_4 = N_{\text{fam}} \lvert R^+(A_3)\rvert$
9	42	$6 \cdot 7 = \lvert R^+(A_3)\rvert \cdot 7$	22	16	$\dim S^+$ (one generation)
10	40	$\lvert R(D_5)\rvert = \sum L = \mathbf{1}^\top L \mathbf{1}$	23	14	$\dim G_2 = \sum L_d = 8{+}6$
11	38	$2 \cdot 19$ (E_8 exponent)	24	12	$\lvert R(A_3)\rvert = \dim \mathfrak{g}_{\text{SM}}$
12	36	6^2 (old EW block)	25	10	$\mathcal{A}_\Lambda = \binom{5}{2}$ (old cosmo block)
			26	8	$h(D_5) = \det R = 2p_1$ **(end: flavor selector)**

Three structural facts jump out (all [I] arithmetic):

- **Endpoints are the compiler boundary.** The cascade *starts* at $60 = p_2 p_3 = D_{\text{start}}$ and *ends* at $8 = h(D_5) = \det R$ — the flavor spectral selector is literally the deepest orbit of E_8.
- **Coxeter rung is $h(E_8)$.** At $n = 15$, $D = 30 = h(E_8)$ — this is the *Coxeter rung*, not the arithmetic centre; the median node is $n = 13$, $D = 34 = p_5(a)$ (below). It matches the exponent-centering audit $\sum_m \lvert m - 15\rvert = 56 = \dim \mathbf{56}_{E_7}$ (which is itself rung $n = 2$).

- **Doubled E_8 exponents are rungs.** $\{58, 46, 38, 34, 26, 22, 14\} = 2 \times \{29, 23, 19, 17, 13, 11, 7\}$ are exactly the cascade rungs at $n = 1, 7, 11, 13, 17, 19, 23$; three of them coincide with anchor numbers ($34 = p_5$, $22 = \sum R$, $14 = \dim G_2$). Only $n = 8$ ($D = 44 = 4 \cdot 11$) is not independently labelled.

Cascade arithmetic — endpoints, centre and length (verified)

The 27 rungs encode E_8 in their endpoints, and their centre is the anchor:

$$\tfrac{D_{\text{start}}\,D_{\text{end}}}{2} = \tfrac{60\cdot 8}{2} = 240 = |R(E_8)|, \qquad 240 + D_{\text{end}} = 248 = \dim E_8, \qquad D_{\text{start}} = 2\,h(E_8) = 60;$$
$$\text{arithmetic median} = \tfrac{60+8}{2} = 34 = p_5(a) = h(E_8) + |\mu_4|, \qquad \textstyle\sum_{n=0}^{26}(D_n - 30) = 27 \cdot 4 = 27|\mu_4|;$$
$$\text{width} = D_{\text{start}} - D_{\text{end}} = 52 = \dim F_4, \qquad 60 \to 30 = h(E_8), \quad 30 \to 8 = \textstyle\sum R = 22;$$
$$\#\text{nodes} = 27 = 3^3 = \mathbf{1}^\top R\,a \ (E_6 \text{ cubic block}), \qquad \#\text{steps} = 26 = \dim \mathbf{26}_{F_4}.$$

So the cascade is an E_8 Coxeter ladder centred on $p_5(a)$ with a per-node $|\mu_4|$ glue shift; its length (27 nodes, 26 steps) carries the E_6 cube and the F_4 fundamental. [I]

The IR tail $n = 24 \ldots 30$: the compiler atoms

Continuing $D = 60 - 2n$ past the last orbit ($D{=}8$, $n{=}26$) is no longer a nilpotent-orbit chain but the *algebraic compiler tail* of the Coxeter clock:

$$D_{24\ldots30} = \underbrace{12}_{|R(A_3)|=\dim \mathfrak{g}_{\text{SM}}}, \underbrace{10}_{\mathcal{A}_\Lambda}, \underbrace{8}_{h(D_5)}, \underbrace{6}_{|R^+(A_3)|}, \underbrace{4}_{|\mu_4|}, \underbrace{2}_{\text{sheet}}, \underbrace{0}_{\text{closure}}.$$

This is exactly what the old "$n = 0 \ldots 30$ clock" was scanning: $n{=}0 \ldots 26$ the E_8 orbit spine, $n{=}27 \ldots 30$ the compiler atoms. (Mark the tail as the *compiler tail*, not orbits.) [I]

Three further cascade lemmas (all machine-verified)

The clock carries more E_8 data than the endpoints:

- **Primitive rungs sum to the root count.** Evaluating $D_n = 60 - 2n$ at the *primitive* indices $n \in (\mathbb{Z}/30)^\times = \{1, 7, 11, 13, 17, 19, 23, 29\} = \operatorname{Exp}(E_8)$ gives $\{58, 46, 38, 34, 26, 22, 14, 2\}$ with

  $$\sum_{n \in \operatorname{Exp}(E_8)} (60 - 2n) = 240 = |R(E_8)|,$$

  and since $30 - n$ is primitive when n is, the clock indexes its own E_8 exponents.
- **The IR tail product is the main Weyl stabiliser.** The nonzero tail $\{12, 10, 8, 6, 4, 2\}$ has

  $$\prod = 46080 = |W(D_5)|\,|W(A_3)| = 1920 \cdot 24, \qquad \sum = 42 = |R^+(A_3)| \cdot 7.$$
- **Cascade variance $=$ E_6 flavour norm times the scalaron-gauge block.** The 27 rungs $60, \ldots, 8$ have mean $34 = p_5(a)$ and

  $$\sum_{n=0}^{26} (D_n - 34)^2 = 6552 = 78 \cdot 84 = \dim E_6 \cdot (7 \dim \mathfrak{g}_{\text{SM}}),$$

  tying the cascade simultaneously to the $E_6 {\times} A_2$ flavour shadow and the A_8 scalaron-gauge shadow. [I] (audit level)

10 Three views of one spine

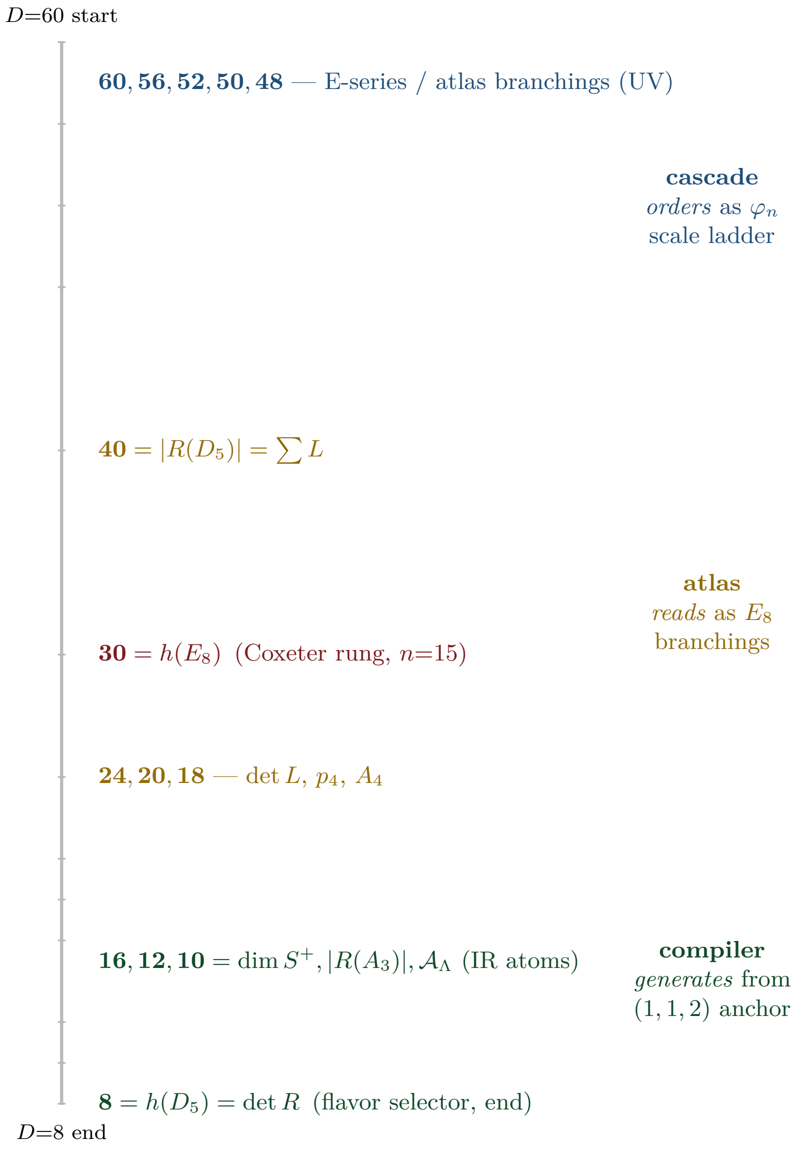

The unification

The same numbers $\{8, 10, \dots, 60\}$ are: (a) the *scale ladder* φ_n of the old cascade (orbit dimensions, ordered by $e^{-\gamma}$); (b) the *branching blocks* of the secondary-atlas ($\dim F_4{=}52$, $\dim \mathbf{56}_{E_7}$, $5{\cdot}10$, $\dots$); (c) the *anchor power sums and matrix invariants* of the compiler ($p_n = 2 + 2^n$, $\det R{=}8$, $\sum L{=}40$). The cascade is the *UV→IR ordering*; the compiler is its *IR endpoint*, where the small numbers $8, 10, 12, 16$ are the flavor/family/carrier data.

11 Direct φ_0^{ret} relations: old vs current

The old cascade's "fundamental relations near $n = 0$" connect directly to the current note — one is *identical*, two are early forms of current refinements:

Quantity	old (V1.06)	current	relation
Ω_b	$\varphi_0^{\mathrm{ret}}(1-2c_3) = 0.04894$	$(1-\frac{1}{4\pi})\varphi_0^{\mathrm{ret}} = 0.04894$	**identical** ($2c_3 = \frac{1}{4\pi}$) [I]
V_{us}/V_{ud}	$\sqrt{\varphi_0^{\mathrm{ret}}} = 0.2306$	$\lambda_C = \sqrt{\varphi_0^{\mathrm{ret}}(1-\varphi_0^{\mathrm{ret}})} = 0.2244$	refined (the $(1-\varphi_0^{\mathrm{ret}})$ factor) [I]
r (tensor)	$(\varphi_0^{\mathrm{ret}})^2 = 0.00283$	$12/N_\star^2 = 0.00397$	two branches, same order [P]

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

So $\Omega_b = \varphi_0^{\mathrm{ret}}(1-2c_3)$ was already the exact current formula in V1.06; the Cabibbo ratio gained its $(1-\varphi_0^{\mathrm{ret}})$ factor; and the old direct $r = (\varphi_0^{\mathrm{ret}})^2$ and the current Starobinsky $r = 12/N_\star^2$ agree in order (~ 0.003) but are distinct predictions — a clean, testable discriminant, stated honestly.

12 Sector inventory: have all the old-cascade sectors been computed?

The old cascade also carried *calibrated absolute scales*: per block, one unit ζ_B was fixed on a single reference and every other scale in the block followed fit-free from the ratio law $\varphi_m/\varphi_n = (D_m/D_n)^\lambda$ (Paper V1.06, App. C). These block assignments are the place where the proton mass, the EW masses and the CMB temperature lived. The honest question is: *has the new compiler re-derived each of these, or only some?* The full inventory, with the current status spelled out:

Old block (rung)	Old observable & calibration	Current-model status	Mark
near $n{=}0$	$\Omega_b = \varphi_0^{\mathrm{ret}}(1{-}2c_3)$, $V_{us}/V_{ud} = \sqrt{\varphi_0^{\mathrm{ret}}}$, $r = (\varphi_0^{\mathrm{ret}})^2$	**covered** — dimensionless, derived from $\{c_3, \varphi_0^{\mathrm{ret}}\}$; see the φ_0^{ret}-table above. Ω_b is *identical.*	[N]
EW, $n{=}12$	$v_H = \zeta_{\mathrm{EW}} M_{\mathrm{Pl}}\varphi_{12}$, $M_W{=}\frac{1}{2}g_2 v_H$, M_Z	**scale covered, absolute scheme-level** — the EW scale enters the current note via the action-grammar exponent and the two-loop flow; M_W, M_Z are dimensionful and remain RG/scheme outputs (one ζ_{EW}), *not* new compiler powers.	[P]
hadronic, $n{=}15, 17$	$\mathbf{m_p}{=}\zeta_p M_{\mathrm{Pl}}\varphi_{15}$, $m_b{=}\zeta_b M_{\mathrm{Pl}}\varphi_{15}$, $m_u{=}\zeta_u M_{\mathrm{Pl}}\varphi_{17}$	**ratios closed, absolute scale an anchor.** The quark mass *ratios* sit on the φ_0^{ret}-word-length ladder and are closed (integer Plücker readouts, v49/v71); the *absolute* values need the per-sector $\Lambda/U_f^\star$ normalisation, an anchor ([A]). The **proton mass** is QCD-confinement, $m_p \sim \Lambda_{\mathrm{QCD}}$ by dimensional transmutation — a *cross-sector* ratio, listed as a genuine frontier item (m_p/m_e, tfpt_4).	[A]
cosmo, $n{=}25$	$T_{\gamma 0}{=}\zeta_\gamma M_{\mathrm{Pl}}\varphi_{25}$, $T_\nu{=}(\frac{4}{11})^{1/3}T_{\gamma 0}$, ρ_Λ	**partly.** ρ_Λ/Λ is *typed* by the action grammar and the seam (Λ-typing, tfpt_2); $T_\nu/T_{\gamma 0}$ is standard QFT (not cascade). The **absolute CMB temperature** $T_{\gamma 0}$ is a downstream readout of the reheating pipeline (Paper 6/7), conditional — not re-derived as a compiler number.	[P]

Direct answer: which sectors are closed, which are not

Closed (derived from $\{c_3, g_{\rm car}, \varphi_0^{\rm ret}\}$): the dimensionless near-n=0 anchors (Ω_b exactly, Cabibbo, tensor branch) and all the *dimension/rung* structure of the spine. **Scale-level only (one ζ_B each, scheme/RG):** the EW masses M_W, M_Z and the cosmological absolute scales $T_{\gamma 0}$ — exactly as in the current SM/cosmology notes; these were *never* pure predictions, not even in V1.06 (each needed a block unit). **Genuinely open:** the **proton mass** m_p (and m_p/m_e) is QCD-confinement across two sectors and is *not* claimed as a compiler power — it stays a frontier item ([A], tfpt_4); the quark mass *ratios* are closed (integer Plücker, v49/v71) and only the *absolute* quark scale m_b, m_u needs the per-sector $\Lambda/U_f^\star$ normalisation, an anchor ([A]). So: yes, the *spine* is fully audited and the *dimensionless* sectors are derived; the *dimensionful* hadronic/cosmo absolute scales are honestly still scheme-level, with the proton mass the one explicitly open cross-sector number — nothing is silently claimed.

13 What is new, what is old, what stays open

New structures found via the bridge

1. The flavor selector $\det R = h(D_5) = 8$ is the *terminal orbit* of the E_8 cascade; the compiler is the cascade's IR endpoint. [I]
2. The atlas branching blocks $(52, 56, 60)$ are the cascade's *UV head* ($n = 0, 2, 4$); the secondary atlas and the orbit cascade are the same spine read top-down vs. as subalgebras. [I]
3. The anchor power sums sit on doubled-exponent rungs: $p_5 = 34$ (n=13), $\sum R = 22$ (n=19), $\dim G_2 = 14$ (n=23). [I]
4. The cascade Coxeter rung $D = 30 = h(E_8)$ (n=15) ties to the exponent-centering audit $\sum |m - 15| = 56$. [I]

Honest separation

The *dimension/rung coincidences* are exact arithmetic [I] — they show the old and new constructions live on one E_8 spine. The old cascade's *scale assignments* φ_n (EW at n=12, hadronic at n=15, 17, cosmo at n=25) are a *different* mechanism from the current $\varphi_0^{\rm ret}$-power mass ladder: the cascade orders *orders of magnitude* (which sector sits where) via $e^{-\gamma}$, while the $\varphi_0^{\rm ret}$-ladder gives *values within* a sector. Both are real, complementary orderings; neither is forced onto the other. The per-sector ζ-calibrated absolute scales (proton mass, EW masses, CMB temperature) remain scheme-level — audited one by one in §12, with the proton mass the single explicitly open cross-sector number.

Bottom line. The old E$_8$ cascade was an early, correct intuition: the physically relevant numbers are exactly the even-integer E_8 orbit spine $\{8, \ldots, 60\}$. The current work *explains why* — they are the $(D_5 \times A_3) + \mu_4$ compiler atoms and their E_8 branching blocks — and identifies the spine's two ends as the compiler boundary ($D_{\rm start} = 60$, $\det R = 8$). Old and new are the same structure at different resolution.

Part III

The self-consistency loop

The question, and the honest answer

The theory rests on two inputs, $c_3 = \frac{1}{8\pi}$ and $g_{\text{car}} = 5$. The old theory had a Möbius seam and the seed was the *start* of the E_8 cascade. Run it *backwards*: can the cascade/E_8 *return* c_3 and g_{car}, closing a self-consistent loop? **Answer:** the *discrete* content is a closed, overdetermined fixed point — $g_{\text{car}} = 5$ and the "8" in c_3 are forced by the E_8 closure. The *only* irreducible primitive left is the geometric π (the Möbius/boundary Gauss–Bonnet). So "two axioms" collapses to "*one* continuous primitive π + *one* discrete self-consistency fixed point."

14 The loop for g_{car} — an overdetermined fixed point

Two *independent* structural conditions, both pure Lie-theory, force $g_{\text{car}} = 5$ once E_8 is the closure object and A_3 is the four-puncture family geometry:

$g_{\text{car}} = 5$ is forced twice over

(A) rank fill.
The glue $D_{g_{\text{car}}} \oplus A_3 \hookrightarrow E_8$ must fill the rank: $\operatorname{rank} D_{g_{\text{car}}} + \operatorname{rank} A_3 = \operatorname{rank} E_8$, i.e. $g_{\text{car}} + 3 = 8 \Rightarrow g_{\text{car}} = 5$.

(B) Coxeter match.
The carrier Coxeter number must equal the E_8 rank, $h(D_{g_{\text{car}}}) = 2g_{\text{car}} - 2 = 8 \Rightarrow g_{\text{car}} = 5$.

Both give $g_{\text{car}} = 5$, and the loop *closes*: $g_{\text{car}} = 5 \Rightarrow D_5 \oplus A_3 \xrightarrow{\mu_4} E_8 \Rightarrow h(E_8) = 30 = 2\cdot 3\cdot 5 \Rightarrow$ largest prime $= 5 = g_{\text{car}}$. [I]

Reverse glue selection — D_5 *and* A_3 from unimodular closure

The closure even *selects the two factors*, not just the rank. The discriminant of $A_{\mu-1}$ is $\mathbb{Z}_\mu$ and of D_g (odd g) is $\mathbb{Z}_4$; a common cyclic glue forces $\mu = 4$, hence $A_{\mu-1} = A_3$. The even-glue norm $q(D_g) + q(A_3) = \frac{g}{4} + \frac{3}{4} = 2$ then forces $g = 5$, hence D_5. A second, independent confirmation of $\mu = 4$ is the *mirror identity* $\dim D_{\mu+1} + \dim A_{\mu-1} = \dim V_{D_{\mu+1}} \cdot |R^+(A_{\mu-1})|$, i.e. $3\mu(\mu+1) = \mu(\mu+1)(\mu-1) \Rightarrow \mu - 1 = 3 \Rightarrow \mu = 4$ (the $45 + 15 = 10 \cdot 6$ identity). So

$$\boxed{\text{common } \mathbb{Z}_4 \text{ glue } + \text{ root norm } 2 \implies D_5 + A_3}$$

— $g_{\text{car}} = 5$ is recovered from the closure, not just posited. [I]

Familyful E_8-closure fixed-point theorem

Among rank-filling pairs $D_g \oplus A_m$ $(g + m = 8)$ *admitting an integer glue index* $\sqrt{\det D_g \det A_m} = \sqrt{4(m+1)}$, *exactly two close to an even unimodular lattice:* $D_8 \oplus A_0$ *(no family) and* $D_5 \oplus A_3$ *(familyful). The unique familyful solution is* $D_5 \oplus A_3$*, with lattice* $\cong E_8$ *and* $g_{\text{car}} = 5$.

D_g+A_m	$g+m$	glue idx^2=4(m+1)	integer?	status
D_8+A_0	8	4	2	valid, *no family*
D_7+A_1	8	8	—	no integer glue
D_6+A_2	8	12	—	no integer glue
$\mathbf{D_5+A_3}$	8	16	**4**=$\lvert\mu_4\rvert$	**valid, familyful,** E_8
D_4+A_4	8	20	—	no integer glue

Equivalently, with $q(A_{\mu-1}) = \frac{\mu-1}{\mu}$ and $g = 9-\mu$ the norm closure $q(D_g) + q(A_{\mu-1}) = 2$ becomes $\mu^2 - 5\mu + 4 = 0$, whose nontrivial root is $\mu = 4$ (A_3, $g = 5$). So $g_{\rm car} = 5$ is forced *three* ways: rank-fill, Coxeter-match, and integer-glue/norm closure. [L] [`verification/v6_bootstrap.py`]

Bootstrap classification: $D_5 \oplus A_3$ is the *unique* familyful E_8 glue

E_8 is the unique even unimodular rank-8 lattice. Its two-factor root-lattice gluings $L_1 \oplus L_2$ (rank $= 8$) that admit a *cyclic* glue are exactly those whose discriminant groups are isomorphic and cyclic:

$$D_5 \oplus A_3\ (\mathbb{Z}_4), \quad E_6 \oplus A_2\ (\mathbb{Z}_3), \quad E_7 \oplus A_1\ (\mathbb{Z}_2), \quad A_4 \oplus A_4\ (\mathbb{Z}_5).$$

Imposing the two TFPT data — one factor with a 16-dimensional half-spinor ($\dim S^+_{D_g} = 2^{g-1} = 16 \Leftrightarrow g = 5$) and one factor giving exactly $N_{\rm fam} = 3$ families (rank $A_m = m = 3$) — selects

$$\boxed{D_5 \oplus A_3 \text{ (uniquely), glue group } \mathbb{Z}_4 = |\mu_4|}\,.$$

$E_6 \oplus A_2$ gives only 2 families, $E_7 \oplus A_1$ only 1; D_8 has non-cyclic glue $(\mathbb{Z}_2)^2$ and A_8 is a single rank-8 factor. So the carrier+family decomposition is forced. [`verification/v15_bootstrap_classification.py`]

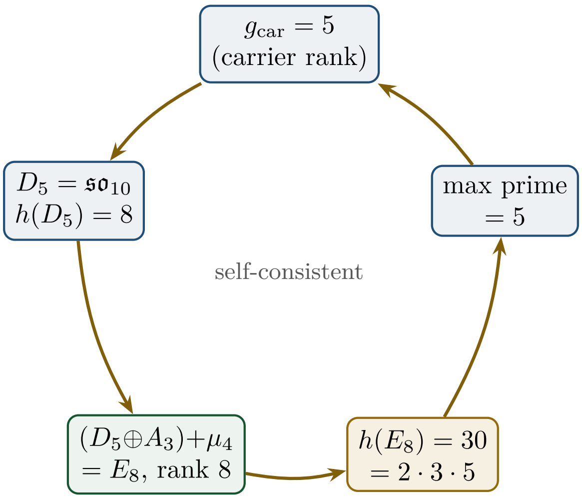

Figure 3: The closed loop: $g_{\rm car} = 5$ builds E_8, whose Coxeter number $30 = 2\cdot 3\cdot 5$ returns the carrier prime 5. The value $g_{\rm car} = 5$ is the unique fixed point (overdetermined by rank-fill *and* Coxeter-match).

The carrier split $3+2$ and the one-generation count $\dim S^+ = 2^{g_{\rm car}-1} = 16$ are then automatic, as is the glue norm $q(D_5) + q(A_3) = \frac{5}{4} + \frac{3}{4} = 2$. The family rank is rank $A_3 = 3 = N_{\rm fam}$ and the glue index is $h(A_3) = 4 = |\mu_4|$ — the Coxeter numbers of the three players are $(h(A_3), h(D_5), h(E_8)) = (4, 8, 30) = (|\mu_4|,\ \det R,\ \text{compiler})$.

15 The loop for c_3 — the “8” is the E_8 rank

The boundary normalisation is $c_3 = 1/(8\pi)$. Its integer is exactly the carrier Coxeter number, which the loop above equals to the E_8 rank:

$$\boxed{c_3 = \frac{1}{h(D_5)\,\pi} = \frac{1}{(2g_{\rm car}-2)\,\pi} = \frac{1}{8\pi}, \qquad 8 = h(D_5) = \operatorname{rank} E_8 = \det R = \varphi(30)}$$

where $\varphi(30) = 8$ is Euler’s totient (the number of primitive residues mod 30 = the E_8 exponents) — so the “8” has *five* concordant readings. The tree seed $\varphi_{\rm tree} = 1/(6\pi)$ has $6 = |R^+(A_3)| = h(E_8)/g_{\rm car}$, the three Gauss–Bonnet boundary cycles ($6\pi = 3 \times 2\pi$). So the *discrete* parts of both c_3 and $\varphi_{\rm tree}$ are E_8/carrier invariants; only the factor π is genuinely continuous.

Honest caveat on 8π

8π is *also* the standard Einstein normalisation ($G_{\mu\nu} = 8\pi G\,T_{\mu\nu}$). The loop *identifies* this gravitational 8 with $\operatorname{rank} E_8 = h(D_5)$ — a non-coincidence *if* E_8 is the closure, but one should be clear that 8π has an independent GR reading. The loop promotes the two 8’s to the same number; it does not derive π.

16 The Möbius origin of the irreducible π

What the loop does *not* remove is π. In the old theory this is exactly the Möbius datum: the orientable double cover of the Möbius fibre is a cylinder with two boundaries plus one $\mathbb{Z}_2$ seam Γ; Gauss–Bonnet gives $\oint k\,ds = 2\pi + 2\pi + 2\pi = 6\pi$ (hence $\varphi_{\rm tree} = 1/(6\pi)$), and the $\mathbb{Z}_2$ identification *is* the sheet “2” of $30 = 2 \cdot 3 \cdot 5$. So the Möbius seam supplies precisely the two things the discrete loop cannot: the continuous π and the sheet parity $\mathbb{Z}_2$.

Net: two axioms → one primitive + one fixed point

$$\{c_3, g_{\rm car}\} \;\longrightarrow\; \underbrace{\pi\ \text{(Möbius/Gauss–Bonnet)}}_{\text{one continuous primitive}} \;+\; \underbrace{E_8\ \text{closure}}_{\text{one discrete fixed point}}\;.$$

The discrete data $\{2, 3, 4, 5, 6, 8, 30\}$ are *all* self-consistent outputs of the E_8 closure ($g_{\rm car} = 5$, $h(D_5) = 8$, $h(A_3) = 4 = |\mu_4|$, $h(E_8) = 30 = 2{\cdot}3{\cdot}5$, $|R^+(A_3)| = 6$). The only irreducible input on the *dimensionless* axis is π. [I] (discrete loop) / [A] (π is primitive)

On the *dimensionful* axis there is exactly *one* more irreducible: a single mass scale. But even that is no longer a *free* constant — fixing the physical Λ branch pins $\bar{M}_{\rm Pl}$, so G_N is a Λ-metrology *output* (`origin_theory`, Λ-typing); the one genuinely irreducible scale is the induced-gravity (Sakharov) anchor. So the complete honest reduction is “π + the E_8-closure fixed point + *one* dimensionful scale”.

One sharpening of the “8”. Both π-coefficients are 2π times a compiler integer: $\varphi_{\rm tree} = 1/(6\pi)$ with $6 = 2N_{\rm fam}$ (three Gauss–Bonnet boundary circles) and $c_3 = 1/(8\pi)$ with $8 = 2|\mu_4|$ (seam winding), so $\varphi_{\rm tree} = (|\mu_4|/N_{\rm fam})c_3 = \frac{4}{3}c_3$. The *same* integer 8 is the Hawking/Einstein coefficient ($G_{\mu\nu} = 8\pi T_{\mu\nu}$, Appendix H): the seam normaliser, the E_8 rank and the universal horizon factor are one number, overdetermined ($8 = 2|\mu_4| = \operatorname{rank} E_8 = h(D_5) = \varphi(30) = \det R$). [I] [`verification/v54_seam_horizon_keystones.py`]

A cyclic element *inside* the hull: the order-30 Coxeter rotation ([`verification/v55_coxeter_cycle.py`])

The compiler hull carries a *literal* cyclic symmetry. Built from the E_8 Cartan matrix, the Coxeter element $c = s_1 \cdots s_8$ has **order exactly** $h(E_8) = 30 = |\mathbb{Z}_2| \cdot N_{\text{fam}} \cdot g_{\text{car}} = 2{\cdot}3{\cdot}5$ (verified by direct matrix powers), and its eigenvalues are the primitive 30th roots $e^{2\pi i\, m/30}$ with m the E_8 *exponents* $\{1, 7, 11, 13, 17, 19, 23, 29\}$ — which are precisely the $\varphi(30) = 8$ *totatives* of 30. Hence

$$\boxed{\operatorname{rank} E_8 = 8 = \varphi(30) = \#\{\text{exponents}\} = \#\{\text{live phases of the order-30 cycle}\}}$$

So the 8 in $c_3 = \frac{1}{8\pi}$ is the count of live phases of a primitive order-30 cyclic rotation of the hull, whose period $30 = |\mathbb{Z}_2| \cdot N_{\text{fam}} \cdot g_{\text{car}}$ is the product of the three core integers. The cyclic structure is *present in the mathematics*, not interposed.

The compiler atoms *are* the E_8 Casimir degrees ([`verification/v66_e8_casimir_degrees.py`])

The load-bearing integers are not an ad-hoc list: they are the fundamental *invariant degrees* of E_8 (exponents $+1$),

$$\deg(E_8) = \{2,\ 8,\ 12,\ 14,\ 18,\ 20,\ 24,\ 30\},$$

with the clean (non-trivial) readings $2{=}|\mathbb{Z}_2|$, $8{=}\operatorname{rank} E_8$, $14{=}\dim G_2$ (the $[K,Q]$ commutator G_2, §Paper 2), $20{=}\det L$, $24{=}|W(A_3)|{=}4!$, $30{=}h(E_8)$ (and the fittable $12{=}|R(A_3)|$, $18{=}p_4(a){=}2{+}2^4$). Two exact sums close it:

$$\sum \deg(E_8) = 128 = 2^7 \quad \text{(the de Sitter prefactor } 128c_3^4\text{; } 7\text{=scalaron exp),}$$
$$\sum \exp(E_8) = 120 = |R^+(E_8)|.$$

Since E_8 is the audit hull, its invariant degrees being the theory's integers is *structurally expected* — an organizing simplification that confirms the E_8 choice rather than a coincidence.

Two-level foundation, and the parameter audit

Local foundation: c_3, g_{car} are the two operational inputs (Papers 1–2). **Global bootstrap foundation:** g_{car} and the integer in c_3 are E_8-closure fixed points; only π is primitive. The audit:

former input	status now	recovered by
$g_{\text{car}} = 5$	discrete fixed point	rank-fill, Coxeter-match, integer-glue/norm, $h(E_8)$
8 in c_3	discrete fixed point	$h(D_5)$, rank E_8, det R, $\varphi(30)$, $2\lvert\mu_4\rvert$ (seam winding)
6 in φ_{tree}	discrete fixed point	$\lvert R^+(A_3)\rvert = h(E_8)/g_{\text{car}}$
π	geometric primitive	Möbius / Gauss–Bonnet (*not* derived)
φ_0^{ret}	derived readout	$\frac{4}{3}c_3 + 48c_3^4$

Correct one-liner: *TFPT has no freely tunable fundamental numbers left; the only remaining continuous origin is π from the Möbius boundary geometry.* (Not “zero axioms”.) [I]/[A]

17 Honest status: bootstrap, not creation from nothing

- **What is genuinely shown [I]:** $g_{\text{car}} = 5$ is an *overdetermined* fixed point of the E_8 closure (two independent Lie-theory conditions); the integer “8” in c_3 is the carrier Coxeter = E_8 rank; the discrete data closes a loop. The earlier landscape scan confirms $g_{\text{car}} = 5$ is the *unique* value at which all hierarchies/family/gauge close — so the fixed point is unique, not one of many. **This uniqueness is now machine-formalised [F] in Lean 4** (Theorem A, `lean4-carrier-rigidity`): the division-free Pascal condition $2^g = g^2 + g + 2$ has the *unique* solution $g = 5$ (`carrier_rank_pascal_unique`; growth lemma for $g \geq 6$; $N_{\text{fam}} = (2^4-1)/5 = 3$, $g+N_{\text{fam}} = 8$), audited to use only the standard kernel axioms.
- **What this is *not*:** a derivation from zero inputs. A self-consistency loop is consistent *at* a value; uniqueness is what makes it predictive (and that is supplied by the landscape scan). The continuous π is not produced by the loop — it is the Möbius/boundary geometry, the one irreducible primitive.
- **Conceptual gain:** the theory’s foundation tightens from “two free axioms” to “*the* E_8-closure fixed point + the geometric π.” That is the Möbius self-consistency you intuited: the seed is the start of the cascade, and the cascade returns the seed’s discrete content.

Part IV

Open items and the E_8-slice scan

18 The honest open-items list

This is the complete current ledger of what is *not* fully closed, graded by how hard the gap is.

#	Open item	Grade	Hardening path / what closes it
1. SM core — finite computations, no new physics			
1a	absolute quark c-scale (ratios closed)	[A]	quark *ratios* are integer Plücker readouts on the derived selector stratum (v49/v69/v71); only the *absolute* $U_f^\star$ normalisation remains — an anchor
1b	θ_{12} solar angle	[N]	**texture closed (below):** explicit $\mu\tau$-symmetric texture verified; only $\varepsilon{=}\frac{3}{4}\varphi_0^{\mathrm{ret}}{\approx}c_3$ (0.23%) stays conditional
2. Scheme layer — standard RG, no new input			
2a	$m_W, m_Z, m_H, \sin^2\theta_W, \alpha_s$	[P]	RG threshold matching R_{SM}; m_H on the closed UV quartic $1/16\pi^2$
2b	hadron pole spectrum (singlet)	[P]	singlet algebra $\mathcal{A}_{\mathrm{sing}}$ explicit; dynamics already closed (gap $6\log\frac{3}{2}$)
3. Conditional — named analytic hypotheses			
3a	gravity (Hodge/R^2 branch)	[P]	APS-Fredholm + truncation written out; full spectral UV completion remains
3b	strong CP / QFT closure	[P]	reflection positivity (RP) is the one remaining input
3c	Λ branch label (rec vs car)	[P]	one discrete choice; typed via $(8\pi)^2$, value not fixed
4. Axioms — not eliminable, only hardenable			
4	P1 (c_3), P2 ($g_{\mathrm{car}}{=}5$)	[A]	more Lean; P2 algebra already machine-verified; remain inputs
5. Frontier — honestly not on a clean ladder (see frontier doc)			
5	η_B, m_p/m_e, Koide, dark matter, full QG	[A]	genuine handles only; not forced (scan results below)
6. Cosmology late-time pipeline — not simplified by the compiler			
6	C_ℓ transfer, $a_{\ell m}$ seam, baryogenesis	[P]	Paper-7 machinery; $a_{\ell m}$ "good CMB world, not yet this CMB world" (SMICA test)

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

Net. Items 1 are finite arithmetic; 2 are standard RG; 3 hinge on named hypotheses (RP, the Λ branch, the UV completion); 4 are the two irreducible axioms; 5–6 are honestly open. The inflation sector (the only cosmology piece the compiler *does* sharpen) is now a parameter-free prediction ($M{=}c_3^{7/2}\bar{M}_{\mathrm{Pl}}$, document 2).

Item 1b closed: the explicit solar-angle Majorana texture [N]

The last open SM angle is now realised by an *explicit* texture, not a fitted number. Take the $\mu\tau$-symmetric Majorana matrix (tri-bimaximal at $\varepsilon{=}0$)

$$M_\nu(\varepsilon) = \begin{pmatrix} -\eta(\varepsilon) & 1 & 1 \\ 1 & 1 & 0 \\ 1 & 0 & 1 \end{pmatrix}, \qquad \eta(\varepsilon) = \frac{2\sqrt{2}}{\tan 2\theta_{12}} - 1, \quad \sin^2\theta_{12} = \tfrac{1}{3}(1-2\varepsilon). \tag{1}$$

$\mu\tau$-symmetry forces $\theta_{23} = 45°$ and $\theta_{13} = 0$ *exactly*; diagonalisation at the TFPT solar deviation $\varepsilon = \frac{3}{4}\varphi_0^{\text{ret}}$ gives

$$\boxed{\sin^2\theta_{12} = \tfrac{1}{3} - \tfrac{\varphi_0^{\text{ret}}}{2} = 0.306747, \quad \theta_{23} = 45^\circ, \quad \theta_{13} = 0} \quad \text{[N]} \tag{2}$$

Honest residual ([`verification/v9_neutrino_texture.py`]). The texture and the angle are exact. The seam identification is $\varepsilon = \frac{3}{4}\varphi_0^{\text{ret}} = \frac{3}{4}\left(\frac{1}{6\pi} + \frac{3}{256\pi^4}\right) = c_3 + \frac{9}{1024\pi^4}$: the *leading geometric term is exactly* $c_3 = \frac{1}{8\pi}$, and the 0.23% residual is purely the φ_0^{ret} seed tail (document 2). So $\varepsilon = c_3$ holds at leading order; the full $\varepsilon = \frac{3}{4}\varphi_0^{\text{ret}}$ is what enters the texture. The reactor angle θ_{13} ($\sin^2\theta_{13} = e^{-5/6}\varphi_0^{\text{ret}}$, document 2) enters through a separate $\mu\tau$-breaking term and is not part of this minimal solar texture. **Scope:** this realises the TFPT solar shift in the $\mu\tau$-symmetric *limit*; it is *not* the full PMNS matrix. The θ_{23}=45° value is the symmetric-limit centre, *not* an octant prediction — current global fits (NuFIT 6.0) leave the octant unresolved — and the reactor angle plus the combined matrix require the separately stated $\mu\tau$-breaking channel.

19 The E_8-slice scan: results (honest hit / miss)

All seven maximal slices of 248 were scanned for the unassigned quantities. The rule was a meaningful (low-complexity) compiler/slice expression matching to $< 0.5\%$; matches that need a free fitted factor are reported as *misses*.

Slice / target	**Result**	**Verdict**
$A_4\times A_4 = SU(5)^2$	$\sin^2\theta_W\vert_{\text{GUT}} = 3/8$ sits in the $SU(5)$ hypercharge embedding; the 12 coset states are the proton-decay X, Y bosons	[I] (known)
$F_4\times G_2$ / Koide	Koide $Q = 2/3 = \lvert\mathbb{Z}_2\rvert/N_{\text{fam}}$ (sheet over families); $J_3(\mathbb{O})$ has dim 27=26+1, 3×3 octonionic “generations” on the diagonal	[P] structural
$D_8 = SO(16)$	$128 = (16,4)\oplus(\overline{16},\overline{4})$ — *same* spinor as $D_5\times A_3$; ν_R (seesaw) lives in the 16; **no new dark-matter state** (only the 4th μ_4 direction)	[I]
$A_8 = SU(9)$, 9=3×3	$\Lambda^3(3,3) = (10,1)\oplus(8,8)\oplus(1,10)$: the $(8,8)$=64 dominates — naive family×colour **fails**	miss
$A_8 \supset SU(5)\times SU(4)$	$\Lambda^3(5+4)$ contains the SM **10**: a genuine chiral-family-unification *lead*, but not TFPT-forced	lead
$E_6\times A_2$ / Jarlskog	cross term $2(\mathbf{1}^\top Ra)N_{\text{fam}} = 162 = 248-78-8$ exact; $J \sim \lambda_Y^6$ only to right *order*, no clean hit	weak
E_8 theta shells	root shells $240, 2160, 6720, \ldots$; $31 = 2^{g_{\text{car}}}-1$ is *not* a clean shell ratio	inconcl.
$m_p/m_e = 1836.15$	no clean E_8/compiler expression (a QCD-confinement ratio); apparent hits need a free factor	**miss**
$\eta_B \approx 6.1\times10^{-10}$	sits *between* $c_3^6 = 4\times10^{-9}$ and $c_3^7 = 1.6\times10^{-10}$; right order, no integer power	order only

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

Genuine new readings the scan produced

- **Koide** = $\lvert\mathbb{Z}_2\rvert/N_{\text{fam}}$. The empirical Koide value 2/3 is exactly the sheet over the family count — a clean (if simple) compiler reading, promoting Koide from “near 2/3” to “democratic $\lvert\mathbb{Z}_2\rvert/N_{\text{fam}}$.”
- **E_8 as the optimal 8-bit code.** E_8 is the optimal 8D lattice (kissing number 240). In byte framing $256-248 = 8 = \operatorname{rank} E_8$ and $256-240 = 16 = \dim S^+$ — a clean information-theoretic reading of the two load-bearing E_8 numbers.
- **Dark matter = no new state.** The $SO(16)$ scan shows the 128 is the *same* matter

spinor; there is no spare E_8 singlet to be a WIMP. Any DM candidate must be the determinant-line axion (frontier doc), not a new E_8 rep — a useful *negative.*

- $SU(9) \supset SU(5){\times}SU(4)$ **chiral lead.** Λ^3 puts the SM **10** in the antisymmetric; the cleanest unexploited route to a family-unification reading, flagged for future work.

Honest negatives

m_p/m_e has *no* clean E_8/compiler derivation (it is a confinement ratio — the apparent “hits” all need a free fitted factor); η_B matches only at the order-of-magnitude level ($\sim c_3^{6.x}$); the Λ rec/car branch is *not* selected by the theta series; the Jarlskog invariant is right only to order λ_Y^6. These are reported as open, not forced.

20 Following the $SU(9)$ lead: cross-slice consistency

The one scan lead worth following is $A_8 = SU(9) \supset SU(5) \times SU(4)$. Splitting $9 = (5,1) \oplus (1,4)$ the E_8 branching $248 = 80 \oplus 84 \oplus \overline{84}$ decomposes *exactly* as:

piece	$SU(5) \times SU(4)$ content	reading
80 (adj)	$(24,1){+}(1,15){+}(5,\bar{4}){+}(\bar{5},4){+}(1,1)$	$SU(5)_{\text{GUT}} + SU(4)_{\text{fam}}$ gauge
$84 = \Lambda^3 9$	$(10,1){+}(10,4){+}(5,6){+}(1,4)$	the SM **10** as family-singlet + family-**4**
$\overline{84}$	$(\overline{10},1){+}(\overline{10},\bar{4}){+}(\bar{5},6){+}(1,\bar{4})$	the conjugate matter

(Dimension checks: 24+15+20+20+1 = 80; 10+40+30+4 = 84; total 248.)

The genuine content: A_8 and $D_5{\times}A_3$ agree on the family [I]

The $SU(4)$ here is *the same* A_3 family group as in the construction slice $D_5 \times A_3$ — the four punctures μ_4 of $\mathbb{P}^1 \setminus \mu_4$. The $SU(5)$ is the Georgi–Glashow GUT sitting inside the carrier $SO(10) = D_5$ ($SU(5) \subset SO(10)$). So the two maximal slices are *consistent*: both read the family as $SU(4){=}A_3$, and A_8 additionally exposes the $SU(5)$ matter content — the SM **10** appears as $(10,1) \oplus (10,4)$, i.e. a family-singlet **10** plus a family-**4**, where the **4** of A_3 is exactly TFPT's “$4 = N_{\text{fam}} + 1$” (three families + the μ_4 direction).

Honest limit: chirality is still the index, not E_8

$84 \oplus \overline{84}$ is *vectorlike* ($10 \oplus \overline{10}$ etc.): E_8 being non-chiral, the slice cannot by itself produce three *chiral* generations. The chirality is the boundary index $\mathrm{Ind}(D^+_{\text{fam}}) = 3$ — the *same* input TFPT already uses. So $SU(9)$ is a consistent re-reading (matter in $SU(5)$, family in A_3), not an independent derivation of $N_{\text{fam}} = 3$. Each $SU(5)$ pair ($\mathbf{10} + \bar{\mathbf{5}}$) is anomaly-free, so the embedding is clean. [P]

Appendix: computational verification

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

- `v1_e8_glue.py` — the μ_4 glue $E_8{=}(D_5{\oplus}A_3){+}\mu_4$ (load-bearing).
- `v4_flavor_matrix.py` — the spectral selector $\det R{=}h(D_5){=}8$, minors $\{2,3,5\}$.

- `v5_e8_cascade.py` — the cascade $D_n{=}60{-}2n$: endpoints, exponent rungs, IR tail, variance.
- `v6_bootstrap.py` — the self-consistency loop: $\mu^2{-}5\mu{+}4{=}0$, $g_{\mathrm{car}}{=}5$ three ways, $8{=}\,\mathrm{rank}\,E_8$.

The E_8-slice/atlas readings are explicitly **[A]** (audit-level), not part of the pass/fail suite. Run `python verification/run_all.py`; full map in `verification/README.md`.