

The Compiler Closure of TFPT

What concretely simplifies, reduces and closes relative to the seven original papers — a status assessment and reading guide

June 8, 2026

In one sentence

The theory tips from “many sectors with surprising hits” (the state of Papers 1–7) to “*one small discrete compiler generates those sectors*”: E_8 — which in v4.5 appeared only *peripherally* (the orbit cascade for scale ordering) — is now *built* from two TFPT atoms and moved to the centre; the flavor matrix, the fine-structure constant and the scale grammar follow as *fixed points* from exactly two inputs $c_3 = \frac{1}{8\pi}$ and $g_{\text{car}} = 5$.

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 **five short documents plus Appendix H**, best read in order (this is the entry point; a detailed map is in §1):

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.
- H. **tfpt_horizon_readouts** — **Appendix H** (unit-system reframe, *not* new physics): one seam constant as the universal horizon thermal code.
- O. **origin_theory** — **Origin Theory**: why no free fundamental number remains. The $(g_{\text{car}}, N_{\text{fam}})=(5, 3)$ skeleton, the triply-forced 8 (geometry = lattice = gravity), the order-30 Coxeter cycle, one boundary transport for flavor and horizon, and the gapped *unique attractor*. **[I]/[L]** core (v53–v56) plus one honestly-typed **[P]** cyclic interpretation.

You are reading document #1.

Contents

1	What this is about: the jump in construction principle	3
2	The architecture at a glance	5
3	The dependency DAG and the proof ledger	6
4	Before / after: what concretely changes	9
5	Paper by paper: what each original paper did, and how it works now	10

6	The compiler: why $30 = 2 \cdot 3 \cdot 5$	11
7	The μ_4 glue: how E_8 is really built	11
8	How derivation now works: the derivation map	12
9	Physical, geometric, topological: which solution is more plausible?	13
10	What role E_8 now plays exactly	13
11	What this means for the theory-of-everything question	14
12	Does this make the theory more probable?	15
13	Predictions for running and upcoming experiments	16
	Conclusion	17
	Appendix: computational verification	17

Reading the markers (refined). The former overloaded “[M]” is split into the grade it actually carries, so a reviewer can see at a glance which claims are exact, which are formal, and which are conditional: **[I]** *exact identity* (e.g. $240 = 16 \cdot 5 \cdot 3$); **[L]** *Lie/lattice theorem* (e.g. $D_5 \oplus A_3 + \mu_4 = E_8$); **[F]** *formalised* (Lean- or script-verified); **[N]** *numerical fixed point* (reproduced, stable); **[P]** *physical/conditional* (follows under named hypotheses); **[A]** *axiom or open*.

Physical — under hypotheses	RG masses, QG	QG.AMB.01
Numerical — fit-free fixed point	$\alpha^{-1}, \Lambda, \theta_{12}$	EM.FP.01
Lattice — Lie / glue / root system	$E_8 = (D_5 \oplus A_3) + \mu_4$	E8.GLU.01
Formal — Lean / exact script	P2 algebra	FORM.P2.02
Axiom — declared inputs	P1 c_3 , P2 g_{car}	AX.P1.01

The *claim stack* (bottom = most certain/fewest, top = conditional). Each claim sits on exactly one level; the bottom two are inputs, the middle two are proved/reproduced, only the top is conditional. The single source of truth for which claim sits where is `verification/status_ledger.csv` (Fig. 1).

No-free-pattern rule. To keep the structure from drowning in decorative coincidences, an identity enters the *main text* only if it (1) is a necessary step in a derivation, (2) kills an alternative, (3) is ablation-relevant, (4) links two previously independent modules, or (5) is experimentally tested. Everything else is explicitly demoted to an *audit fingerprint* **[I]** (clearly boxed as such) — so the load-bearing beams stay visible and the ornament never gets promoted to spine.

Admissible-invariant classes (the harder filter). Sharper still: a *load-bearing* integer may only come from one of seven typed sources — (i) a symmetric function $e_k(a)$ of the anchor $a = (1, 1, 2)$; (ii) a power sum $p_n(a)$ or difference $p_{n+1} - p_n$; (iii) a determinant, spectrum, Smith normal form or trace of (R, Q, K, L) ; (iv) an anchor quadratic form $u^\top M v$ with $u, v \in \{\mathbf{1}, a\}$; (v) a Plücker coordinate of the anchor plane $\langle \mathbf{1}, a \rangle$; (vi) an E_8 Casimir degree or branching block; (vii) a pre-registered experimental observable. Anything else is an audit fingerprint or a frontier item, never spine. This is what converts TFPT from *pattern hunting* (“look, another number”) into a typed compiler: *fewer admissible kinds of explanation*, not more matches.

Reviewer path (read in this order; the rest is support)

1. the architecture: the two axioms $\{c_3, g_{\text{car}}\}$ and the two-engine picture (§1, the figures);
2. the E_8 glue $E_8 = (D_5 \oplus A_3) + \mu_4$ — document 1, Part II, machine-checked

(verification/v1_e8_glue.py);

3. the α^{-1} fixed point plus ablation (v3_em_alpha.py);
4. the **single status ledger** verification/status_ledger.csv — the source of truth for every claim’s grade, location, dependency and script;
5. the honest frontier list (document 4).

Everything carries a claim ID (e.g. E8.GLU.01, EM.FP.01) that resolves to a row of the ledger. **If the text and the ledger ever disagree, the ledger wins.**

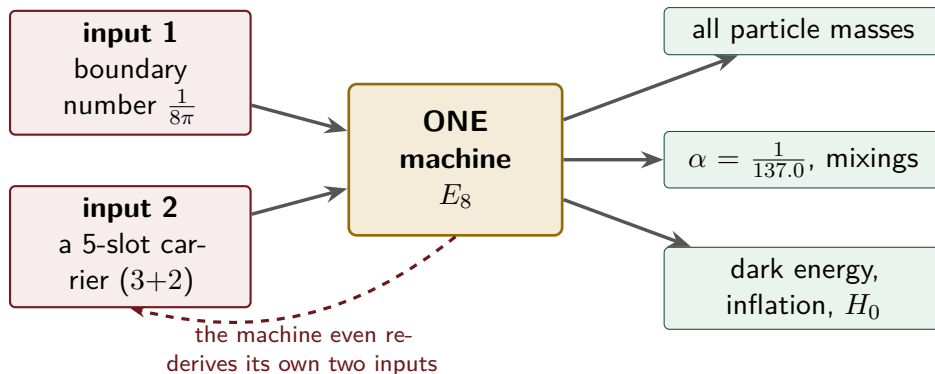
1 What this is about: the jump in construction principle

In plain words — what we found, and why it matters

What we found. The older version of the theory already reproduced dozens of real numbers of physics (particle masses, the fine-structure constant, dark-energy and inflation values) from a small set of inputs. What it could *not* say was *why* the same handful of small integers (2, 3, 5, 16, ...) kept reappearing in every sector. The new result closes exactly that gap: those numbers are not separate lucky hits. A single tiny “machine” — fed by just **two inputs** (a boundary number $\frac{1}{8\pi}$ and a five-slot carrier) — builds the exceptional symmetry E_8 , and E_8 is the unimodular “audit hull” that classifies the discrete Standard-Model readout: the constants, the flavor skeleton and the gravity/cosmology regime are then read off by projection and boundary dressing (not “printed” for free — where a scheme or analytic interface is needed, it is named).

How this drastically simplifies the theory. Instead of seven separate derivations that happen to share numbers, there is now *one generator and a handful of corollaries*. Roughly 250 pages of sector-by-sector work collapse to “two inputs + one machine”; the number of independent structural assumptions drops to **two**.

How plausible it is. A theory becomes more believable when it *explains more from less*. Here the two inputs are even *over-determined*: the machine reproduces the very two numbers it started from (the dashed loop below), so there is *no free dial left to tune*. The only genuinely free, continuous ingredient that remains is the number π .



Two numbers go in on the left; the discrete Standard-Model core, the dimensionless constants and several scale readouts come out on the right — and the dashed loop is the point: the machine reproduces the very two numbers it started from, so the discrete core is overdetermined rather than fitted. (Dimensionful EW/QCD masses, the full quantum-gravity measure and the frontier items stay typed-open — see the status ledger.)

The seven original papers (tfpt-45, Papers 1–7) deliver a long chain of physical closures: boundary kernel → carrier → Standard-Model packet → EM fixed point → flavor transport →


admissibility/strong CP \rightarrow geometric gravity/metrology \rightarrow cosmology/CMB. Each stage stands on its own and hits a remarkable number of values. What was missing was a *meta-statement* explaining *why* so many independent sectors use the same small stock of numbers (2, 3, 5, 16, 40, 41, 48, 240, 248).

Version 4.5 — seven papers

- P1 boundary
- P2 carrier
- P3 EM/flux
- P4 strong CP
- P5 metrology
- P6 cosmology
- P7 CMB

many parallel sector closures;
 E_8 only a **peripheral** backdrop

New — one compiler



two inputs build E_8 ;
 E_8 feeds back and fixes the inputs (closed loop)

The difference in one picture. v4.5 is a *stack of seven independent closures* that happen to share numbers, with E_8 assumed in the background. The new approach is *one discrete compiler*: two inputs build E_8 , E_8 outputs the sectors, and E_8 feeds back to fix the two inputs (the red loop). What were seven parallel results become one generator with corollaries; what was a posited E_8 becomes a constructed-and-self-checking closure.

Sharper still: one anchor $a = (1, 1, 2)$, and E_8 as a compiler hull [1]/

The two axioms are not even independent: they are the *elementary symmetric polynomials* of the single parabolic anchor $a = (1, 1, 2)$ (the exponents-at-infinity / splitting type $\mathcal{O}(-2) \oplus \mathcal{O}(-1)^2$):

$$e_1(a) = 4 = |\mu_4|, \quad e_2(a) = 5 = g_{\text{car}}, \quad e_3(a) = 2 = |\mathbb{Z}_2|, \quad c_3 = \frac{1}{2e_1(a)\pi} = \frac{1}{8\pi}.$$

Its *power sums* $p_n(a) = 1^n + 1^n + 2^n = 2 + 2^n$ generate the big Lie data directly: $p_1=4=|\mu_4|$, $p_2=6=|R^+(A_3)|$, $p_3=10=\mathcal{A}_\Lambda$, and

$$|R(E_8)| = p_1 p_2 p_3 = 240, \quad \dim E_8 = p_1 p_2 p_3 + (p_4 - p_3) = 248, \quad p_4 - p_3 = 8 = \text{rank } E_8,$$

with the binary ladder $p_{n+1} - p_n = 2^n$. So the inputs collapse from $\{c_3, g_{\text{car}}\}$ to $a = (1, 1, 2)$ plus π . [verification/v23_anchor_generator.py]

And a is itself *half-forced* ([L]/[P]): on \mathbb{P}^1 every bundle splits (Birkhoff–Grothendieck), so once $|\mu_4| = 4$, $N_{\text{fam}} = 3$, $\deg E = -4$ are set, the three positive anti-degrees are positive integers summing to 4, and $4 = 2 + 1 + 1$ is the *unique* partition into three positive parts. The genuine remaining axiom is therefore not “why $a = (1, 1, 2)$ ” but “why the carrier $g_{\text{car}} = 5 / (|\mu_4|, N_{\text{fam}}) = (4, 3)$ ” — that interface (with c_3) is the [A] input layer.

What E_8 is (and is not). In TFPT E_8 is *not* an unbroken physical gauge group: it is the unimodular *audit / compiler hull* that classifies the admissible discrete charge and residue structures. The Standard Model is a *readout* after projection, not a direct E_8 gauge theory; the Lorentz signature appears only after projection; fermions are not “everything in the adjoint”. This is exactly why the known no-go results against literal E_8 world-formulas (Distler–Garibaldi; Coleman–Mandula) do *not* bite: they constrain E_8 as a spacetime gauge symmetry, which

TFPT does not claim. The defensible claim is: *TFPT is the unique parabolic compilation of the anchor $a = (1, 1, 2)$ into the E_8 audit hull.*

In one line:

The anchor generates the syntax; the boundary generates the scale; E_8 checks the consistency.

i.e. the discrete content (carrier $D_5 \oplus A_3 + \mu_4$, the operators R, Q, K, L) flows from $a = (1, 1, 2)$; the continuous content ($c_3 = \frac{1}{8\pi}$, φ_0 , α^{-1} , the exponential scales) flows from π ; and E_8 is the unimodular *checksum container*, not the world group.

That is now the case. The whole development is consolidated into **five companion documents** (plus this introduction), organised in layers that build on one another; this introduction is the entry point and reading guide. Each document collects the earlier notes of its layer into self-contained parts:

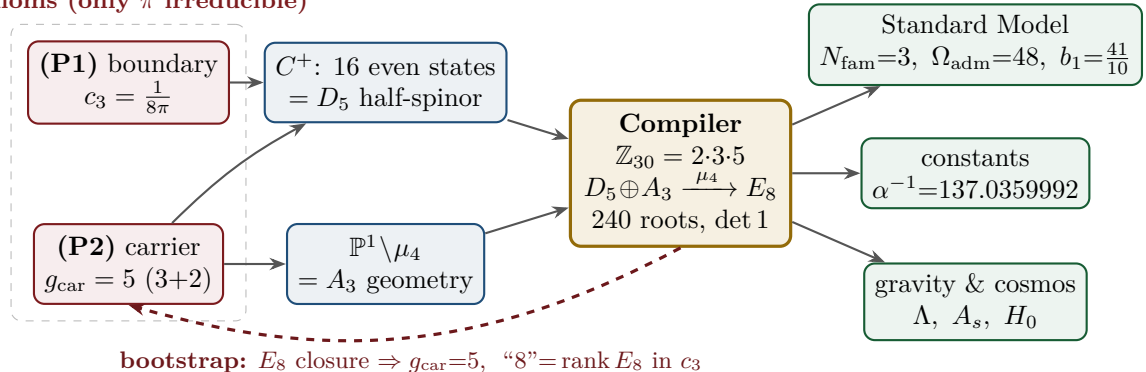
Document	What it contains (former notes now as parts)
tfpt_1_architecture_e8	Core. The two axioms $\{c_3, g_{\text{car}}\}$, the derivation map and the EM fixed point (Part I); the Coxeter-cyclotomic compiler $\mathbb{Z}_{30}=2\cdot3\cdot5$, the explicit $D_5 \oplus A_3 + \mu_4$ E_8 construction, the machine-checkable E_8 appendix and the group-level glue $(\text{Spin}(10) \times \text{SU}(4)) / \Delta \mathbb{Z}_4$ (Part II).
tfpt_2_standard_model	Standard Model. The full SM in one φ_0^{ret} -ladder formula (Part I); the flavor block from parabolic weights and the transport theorem, H2 as a spectral selector (Part II); the worked closures — explicit gap, APS-Fredholm, truncation/ n_s, r , Starobinsky M , derived θ_{12} , quark c , the H2 splitting type, Λ -typing (Part III).
tfpt_3_e8_audit_bootstrap	E_8 audit & bootstrap. The seven E_8 slices as an audit raster ($E_6 \times A_2$ reads the flavor matrix) (Part I); the old cascade $D=60-2n$ as the same even-integer spine (Part II); the Möbius bootstrap — $g_{\text{car}}=5$ and the “8” in c_3 as overdetermined fixed points, only π irreducible (Part III).
tfpt_4_frontier	Frontier. The honest status of η_B , m_p/m_e , Koide, dark matter and quantum gravity — the genuine handle, the precision, and what is <i>not</i> forced onto the ladder.
tfpt_horizon_readouts (App. H)	Appendix H — horizon unit system (reframe, not new physics). One seam constant $c_3 = \frac{1}{8\pi}$ as the universal horizon thermal code: Hawking/de Sitter/Unruh, BH thermodynamics, Page, scrambling, Nariai bound, $v_{\text{GW}}=c$ — standard physics in seam units, two compiler fingerprints ($1920= W(D_5) $, $ \mu_4 =4$).

The five formerly “research-level” problems are now *closed, reduced or typed* inside the documents where they belong — the group-level closure $(\text{Spin}(10) \times \text{SU}(4)) / \Delta \mathbb{Z}_4$ in document 1 (closed), and the H2 splitting type, the Starobinsky scalaron mass, the θ_{12} texture and the quark c in document 2 (the first two closed; the quark *ratios* now closed combinatorially (v49/v71), only the *absolute* scale a U_{point} anchor; θ_{12} the seam-misalignment residual). Five companion files (plus this introduction) instead of a dozen, each result where it logically lives.

2 The architecture at a glance

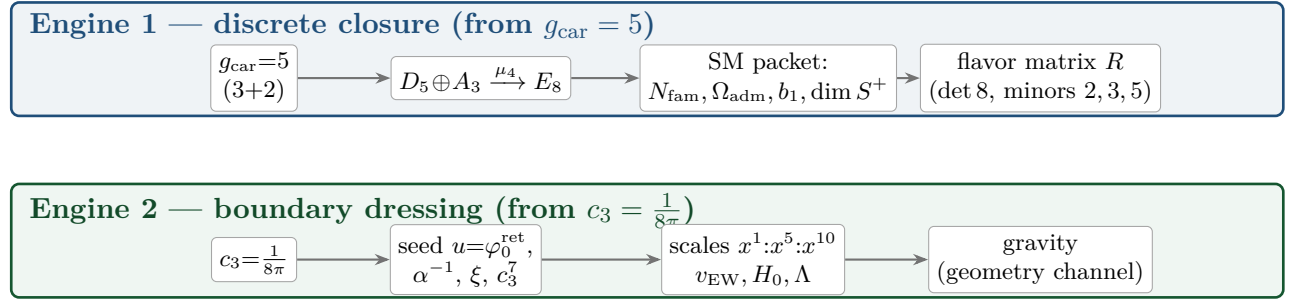
The diagram below shows the whole pipeline: *two* inputs on the left, the discrete compiler in the middle, the physics on the right. Everything in between is a consequence.

2 axioms (only π irreducible)




Key statement (the loop closes): previously D_5 , A_3 and E_8 sat *as given* Lie structures in the background. Now they are *intermediate products* of the two inputs — *and the red back-channel is the new part*: the E_8 closure feeds back and *fixes the inputs*. $g_{\text{car}} = 5$ is the unique rank-fill/Coxeter/Pascal solution ($2^{g-1} = \binom{g}{0} + \binom{g}{1} + \binom{g}{2}$), and the integer 8 in $c_3 = \frac{1}{8\pi}$ is rank $E_8 = h(D_5) = \det R$. **Inputs and output are mutually locked — they prove each other;** the only thing *not* produced by the loop is the continuous π (Möbius/Gauss–Bonnet), the lone irreducible primitive. So it is not a one-way arrow chain but a closed self-consistency loop.

The master story is two engines. Read from the two axioms, the whole theory factorises into exactly *two* engines — a *discrete* closure and a *boundary* dressing. Gravity is not a third block; it is the geometry channel of Engine 2.



3 The dependency DAG and the proof ledger

The whole theory is a directed acyclic graph: two axioms at the source, the compiler in the middle, the observables as sinks. The DAG makes the attack surface explicit — each arrow is a named theorem or a declared input, and each node fails in exactly one way.



Node	Definition	Inputs	St.	Output	Failure mode
P1	RP boundary kernel, $c_3=1/(8\pi)$	— (axiom)	[A]	c_3	no reflection-positive seam
P2	5-slot carrier $g_{\text{car}}=5$	— (axiom)	[A]	carrier 3+2	wrong family/charge lattice
D_5	C^+ even-Hamming = half-spinor 16	P2	[I]	$\dim S^+=16, Y$	group/matter mismatch
A_3	$\mathbb{P}^1 \setminus \mu_4$ four-puncture geom.	P1	[I]	$N_{\text{fam}}=3, \mathbb{Z}_6$	wrong multiplicity
E_8	μ_4 glue: $\text{disc}=\mathbb{Z}_4, q\text{-sum}=2$	D_5, A_3	[I]	240, 248	not even-unimodular
R, L	residue + winding matrix, anchor (1, 1, 2)	A_3	[I]	det 8, minors (2, 3, 5)	wrong D_6 branch
α^{-1}	root of $F_{U(1)}(\alpha)=0$	$c_3, M=41$	[N]	137.0359992	no/second root (excl.)
masses	φ_0^{ret} -ladder $\lambda_V^L \Lambda$	R, L, c_3	[I]/[P]	9 masses, mixings	hierarchy mismatch
θ_{12}	TBM + seam $\varepsilon=c_3$	R, L	[N]/[P]	0.307	seam-misalign. lemma fails
Λ, A_s	seam det. / R^2 scalaron	c_3	[P]	Λ, A_s, n_s, r	ambient RP fails

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

Reading the table top-to-bottom is reading the theory: two red axioms in, everything else a consequence with exactly one failure mode each.




Figure 1: Status heatmap, generated directly from the single source of truth `verification/status_ledger.csv`. Grey = axiom, blue = formal/lattice, green = numerical, gold = physical/conditional, red = open. Each row carries a claim ID (`E8.GLU.01`, ...) that resolves to a ledger row with its dependencies and verification script. The ledger is *versioned*: each row has `active` and `canonical_status/supersedes` fields, so claims that were later superseded (e.g. the early “flavor gate open” rows, now *ratios closed* via `FLAV.RIGID`) are marked inactive and drawn dimmed — the canonical status is the only one read as current.

The compact status formula (three honest layers)

compiler <hr style="width: 80%; margin: 0 auto;"/> closed	admissible IR physics <hr style="width: 80%; margin: 0 auto;"/> conditionally closed (RP, gap)	strict physical TOE <hr style="width: 80%; margin: 0 auto;"/> open
--------------------------------------------------------------	---------------------------------------------------------------------------------------------------	-----------------------------------------------------------------------

TFPT 5.0 closes the discrete compiler, the algebraic Standard-Model readout, the EM fixed point, the admissible gapped IR sector and the $R + R^2$ spectral-action shadow. It does *not* yet certify a strict physical TOE. The remaining items have now collapsed sharply: the R -bridge amplitude normalisation U_{point} *reduces to* v_{geo} (Gate 1 closed, v75 — the same anchor as $1/G$),

and the ambient metric measure is *holographically reduced* to a finite seam-boundary measure (Gate 2 tier B, v76) — which is in turn the rigorously-constructed $(E_8)_1$ lattice net (v77), so G_{metric} is imported into existing RCFT rigor rather than built anew. So the only genuine residuals are *one* dimensionful scale v_{geo} — the *dimensional-analysis floor* that no theory escapes (v78), shared by flavor and gravity — the boundary net identification, and the named QFT/cosmology transfer interfaces (F_{frontier}).

Two points that are *not* in the residual (closed, but easy to mis-list as open).
(i) Parameter-freeness is a theorem, not a hope: the boundary transport is gapped ($\Delta = 6 \log \frac{3}{2} > 0$), so by Perron–Frobenius its iteration has a *unique* attracting fixed point — the constants are *selected*, not tuned ($[I]/[L]$, [verification/v56_unique_attractor.py]), reinforcing the static bootstrap web ($g_{\text{car}} = 5$ forced three ways, [verification/v6_bootstrap.py], [verification/v14_carrier_uniqueness.py]). *(ii) Newton’s constant is not an independent input:* fixing the physical Λ branch ($((8\pi)^2 \delta_{\text{top}} = \frac{3}{4\pi^2})$ pins \bar{M}_{Pl} relative to Λ , so G_N is read out *after* the branch choice + *one* dimensionful induced-gravity anchor ($[I]/[P]$, [verification/v60_lambda_metrology_branch.py]), not predicted from pure numbers. What stays irreducible is then just π , the discrete E_8 -closure fixed point (self-consistent, *not* a free axiom; the bootstrap of `tfpt_3`), and *one* dimensionful scale v_{geo} — and that single scale is shared by *both* the flavor sector ($U_{\text{point}} \rightarrow v_{\text{geo}}$, v75) and gravity ($1/G$, v68): the two $[A]$ anchors are *one* (dimensional analysis; no metre from pure mathematics).

The hard reduction: everything open is two research gates plus typed interfaces

After the compiler closure, the entire residual is

$$\text{Rest} = (U_{\text{wall}}) \oplus (G_{\text{metric}}) \oplus (F_{\text{frontier}})$$

— one parabolic wall-selection, one quantum-gravity measure, and a set of deliberately typed QFT/cosmology interfaces. The status card:

Item	Status	Reading
parameter-freeness (self-consistency)	$[I]/[L]$	gapped transport \Rightarrow <i>unique</i> Perron–Frobenius attractor (v56); bootstrap web $g_{\text{car}}=5$ forced (v6/v14)
Newton constant G_N	$[I]/[P]$	not <i>independent</i> : branch pins \bar{M}_{Pl} rel. to Λ (v60); readout after branch + one dim. anchor
ratio $c_u/c_d = \frac{55}{117}$	$[I]$	$g_{\text{car}} \ \text{P}_{1,a}(K) \ _1 / (N_{\text{fam}}^2 \Delta_Q)$; rigid on the derived stratum (v49/v71), no solve
(U) flavor gate	$[A] \rightarrow v_{\text{geo}}$	Gate 1 closed: $U_{\text{point}} \rightarrow v_{\text{geo}}$ (ratios + Grand Mass Volume, v75); same anchor as $1/G$
$R + R^2$ gravity	$[I]/[P]$	covariant $f(R)$ field equation closed in regime; heat-kernel grounded (G2)
physical QG sector	$[F]/[P]$	gap-decoupled admissible measure ($\Delta_{\text{eff}}=1.648>0$)
ambient QG measure	$[A]$ (reduced)	Gate 2 tier B: holographically reduced to a finite seam-boundary measure (v76); IR tier closed (decoupling); strict-TOE cert. open (G_{metric})
Koide	$[P]$	$Q_{\text{src}} + \varphi_0/24$ as source \rightarrow pole transfer conjecture
axion DM	$[P]/[A]$	$f_a = M_{\text{scal}}/128$ as a scenario
η_B	$[P]$	readout closed; leptogenesis Boltzmann solve missing
m_p/m_e	$[A]$	cross-sector QCD/EW ratio
N_*	$[P]$	reheating input, $N_* \simeq 57$ (continuous)
$a = (1, 1, 2)$	$[L]/[P]$	forced by $4 = 2 + 1 + 1$, conditional on the carrier
E_8 Lean	$[F] \rightarrow [L]$	hardening, not a blocker (script-certified)

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

Only *two* research gates remain — (U_{wall}) and G_{metric} ; the frontier items are typed interfaces, not structural holes. Of these, (U_{wall}) is the only genuine *physics* gate left: G_{metric} is now heat-kernel grounded (G2: the spectral action gives $R + R^2$ structurally) and *gap-decoupled* (G5: $\Delta_{\text{eff}}=1.648>0$ isolates the closed admissible sector from the un-built ambient), so its residual — the ambient projective limit (G6) — is a *strict-TOE completion target*: its absence does not affect the bounded IR claim, but it does block certification as a strict physical TOE. Both are written up as numbered *research contracts* (lemma chains + closing theorem + machine-certifiability) in the companion `tfpt_research_contracts`, with priority (U_{wall}) first. [`verification/v36_spectral_action_g2.py`]

4 Before / after: what concretely changes

Aspect	State of Papers 1–7	After the compiler closure
E_8	<i>peripheral</i> : only an $11D \rightarrow 4D$ reduction aside and the nilpotent-orbit <i>cascade</i> (scale ordering)	central, constructed : $C^+=D_5$ spinor, $\mathbb{P}^1 \setminus \mu_4=A_3$, μ_4 glue $\Rightarrow E_8$ [L]
E_8 numbers	240, 248 read off tables when needed	$240=16 \cdot 5 \cdot 3$, $248=240+8$ <i>derived</i> as carrier traces [I]
fine structure α^{-1}	<i>already</i> a parameter-free cubic fixed point $\alpha^3 - 2c_3\alpha^2 - 8bc_6 \ln(\cdot)=0$, $\alpha^{-1} \approx 137.0365$	<i>same</i> equation, now existence & uniqueness proved and refined to 137.0359992 ; dev. 2.9×10^{-10} vs CODATA-2022 (1.9σ) [I]/[N]
flavor residues	partly audit-flagged	anchor $\{1, 1, 2\}$ & matrix forced by 3 hypotheses [I]/[P]
many sectors	side by side, each its own closure	<i>one</i> compiler $\mathbb{Z}_{30}=2 \cdot 3 \cdot 5$ generates them [I]
inputs	c_3, g_{car} (+ implicit structure)	exactly 2 axioms (P1, P2), clearly declared [A]
A_s, n_s, r (inflation)	scale present	scalon mass fixed by c_3^7 (3.06×10^{13} GeV); A_s, n_s, r become <i>predictions</i> [P]
θ_{12} (solar)	open (leading $\frac{1}{3}$)	cond. derived $\frac{1}{3} - \frac{\varphi_0^{\text{ret}}}{2} = 0.307$, given seam $\varepsilon=c_3$ [N]/[P]

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

Reduction in one line

The number of *structural* assumptions drops from “a fixed-point seed c_3 , a carrier, the orbit cascade, plus a per-sector calibration in each block” to **two axioms**. Everything else ($N_{\text{fam}}, \Omega_{\text{adm}}, b_1, \alpha^{-1}$, the flavor matrix, the E_8 glue, the scale grammar) becomes a *consequence*. *And even those two tighten*: the bootstrap note shows $g_{\text{car}} = 5$ and the integer 8 in c_3 are *overdetermined E_8 -closure fixed points* (rank-fill, Coxeter-match, integer-glue/norm all give 5; $8 = h(D_5) = \text{rank } E_8 = \det R = \varphi(30)$), leaving *one* genuine continuous primitive: π from the Möbius boundary. So the *discrete* core has no freely tunable fundamental numbers — the bootstrap is *not creation from nothing* (two inputs remain), but an overdetermination of those inputs. (This concerns the discrete core only; dimensionful masses, the QG measure and frontier items remain typed-open.)

5 Paper by paper: what each original paper did, and how it works now

The seven original papers (plus the orientation note) total roughly **247 pages**. The table below puts each one next to its compiler-era replacement: *what* it treated and computed, *how many pages* it took, *where* that now lives and how much it compresses, and whether the *accuracy* of the result improves on top. The honest summary up front: in most sectors the *numbers are the same* — what shrinks dramatically is the *derivation length* and the number of independent assumptions; in three places (the solar angle, the inflation amplitude, the scalaron mass) there is a *genuinely new or sharper* result.

P	Treated & computed	Old	How it works now (note; compression)	Acc.
0	orientation: architecture, decoder core $Y, [u\Sigma], u$, theorem stack	20 p	<i>one</i> compiler diagram: $\{c_3, g_{\text{car}}\} \rightarrow D_5 \oplus A_3 \rightarrow E_8$ (intro + action_grammar)	=
1	boundary kernel: $c_3 = \frac{1}{8\pi}$, Calderón polarization, carrier decoder $6Y^2 - Y - 1 = 0$	20 p	the two axioms P1 (c_3) + P2 ($g_{\text{car}}=5$); carrier = 3+2 atom (compiler)	=
2	carrier/SM packet: $S^+=16$, G_{phys} , $N_{\text{fam}}=3$, $\Omega_{\text{adm}}=48$, hypercharge, b_1	48 p	compiler reads $16=1+5+10$, N_{fam} , Ω_{adm} , $b_1 = \frac{41}{10}$ as carrier traces; $C^+ = D_5$ spinor (~ 5 p)	=
3	EM closure: α^{-1} via $\sum L + N_\Phi = 41$; flavor transport, CKM/PMNS, θ_{13}	38 p	α^{-1} as unique root of $F_{U(1)}(\alpha)=0$; flavor from transport theorem (parabolic + sm)	↑
4	admissibility/strong CP/QFT: $\theta_{\text{eff}}=0$, OS reconstruction, mass gap	40 p	explicit gap $6 \log \frac{3}{2} \Rightarrow$ Dobrushin/OS; strong CP from structure+RP (closing_comp.)	=
5	geometric Hodge/metrolology: M_{P1} , v_{geo} , Λ , spectral Einstein	30 p	Hodge/ R^2 branch; Λ <i>typed</i> (one branch label); boundary-normalised (closing_comp.)	=
6	cosmology interface: FRW, Λ_{IR} , scalaron, reheating, leptogenesis, axion	15 p	Starobinsky $M = c_3^{7/2} \bar{M}_{\text{P1}} = 3.06 \times 10^{13}$ GeV; A_s <i>parameter-free</i> (closing_comp.)	↑new
7	CMB operational closure: C_ℓ , $a_{\ell m}$, seam mixer, Planck pipeline	36 p	$n_s=0.964$, $r=0.0040$ from the R^2 attractor, no free amplitude (closing_comp.)	↑new

Three places where accuracy genuinely improves “on top”

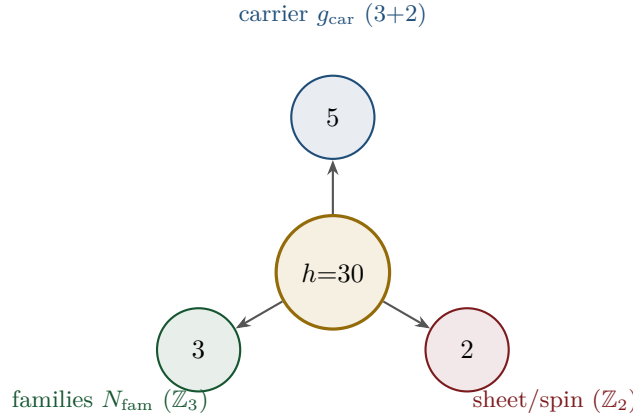
- Solar angle θ_{12} .** Paper 3 left it as the only open SM number (leading $\frac{1}{3}$). It is now *conditionally derived* (modulo the seam-misalignment lemma): TBM gives $\frac{1}{3}$, the charged-lepton 1–2 misalignment is the seam $\varepsilon = \frac{N_{\text{fam}}}{|\mu_4|} \varphi_0^{\text{ret}} = \frac{3}{4} \varphi_0^{\text{ret}}$ (leading term exactly $c_3 = \frac{1}{8\pi}$), and TBM geometry gives $\sin^2 \theta_{12} = \frac{1}{3} - \frac{2}{3} \varepsilon = \frac{1}{3} - \frac{\varphi_0^{\text{ret}}}{2} = 0.3067$ — 0.1% from NuFIT.
- Inflation amplitude A_s .** Generic Starobinsky fits the scalaron mass to A_s . TFPT fixes it by the seam, $(M/\bar{M}_{\text{P1}})^2 = c_3^7$, so $A_s = \frac{N_\star^2}{24\pi^2} c_3^7 = 2.0 \times 10^{-9}$ is a *prediction* (Planck 2.1×10^{-9}), and the measured A_s even *predicts* $N_\star = 56$.
- Scalaron mass.** $M = c_3^{7/2} \bar{M}_{\text{P1}} = 3.06 \times 10^{13}$ GeV comes out *exactly* at the canonical Starobinsky value — a number that was an input before is now an output.

What “simpler” means concretely. The compression is not cosmetic. Paper 2 spent 48 pages

establishing the carrier/SM packet by rigidity arguments; the compiler reads the same $16=1+5+10$, $N_{\text{fam}}=3$, $\Omega_{\text{adm}}=48$, $b_1=\frac{41}{10}$ off the Pascal row of a 5-slot carrier in a page. Paper 3’s flavor transport (38 p) becomes one transport theorem plus a fixed residue matrix whose characteristic polynomial $t^3 - 9t^2 + 10t - 8$ carries only compiler numbers. The “many independent closures” of Papers 1–7 become *one* generator with a handful of corollaries.

6 The compiler: why $30 = 2 \cdot 3 \cdot 5$

The Coxeter number of E_8 is $h = 30 = 2 \cdot 3 \cdot 5$. Exactly these three prime factors are the three discrete atoms of the theory — that is the point of the compiler:




From this the compiler reads, among others:

- $\dim S^+ = 1 + 10 + 5 = 16$ (one generation) as the even exterior algebra of the 5-slot carrier — the same Pascal row $\binom{5}{0}, \binom{5}{1}, \binom{5}{2} = 1, 5, 10$ appears in family, action ladder and the E_8 root count. [I]
- $|R(E_8)| = 16 \cdot 5 \cdot 3 = 240$ and $\dim E_8 = 240 + (5 + 3) = 248$. [I]
- the \mathbb{Z}_{30} Coxeter element \leftrightarrow the cyclotomic polynomial Φ_{30} ; the corner rotation $z \mapsto iz$ on $\mathbb{P}^1 \setminus \mu_4$ is the A_3 Coxeter element. [I]

7 The μ_4 glue: how E_8 is really built

The decisive, now *proved* step: D_5 and A_3 have the *same* discriminant group \mathbb{Z}_4 , and their discriminant-form norms are exactly two TFPT constants which add up to the E_8 root norm.



Glue theorem (proved)

$\text{disc}(D_5) = \text{disc}(A_3) = \mathbb{Z}_4$, glue index $|\mu_4| = 4$, and $q(D_5) + q(A_3) = \frac{5}{4} + \frac{3}{4} = 2 =$ the E_8 root norm. Hence $E_8 = (D_5 \oplus A_3) + \mu_4\text{-glue}$ is a *closed lattice-theoretic* construction — not a blind

“positing of 248”. [L]

8 How derivation now works: the derivation map

What can one concretely *compute* with $\{c_3, g_{\text{car}}\}$? The following map summarises what is now a consequence rather than an assumption.

Sector	How it now arises	Status
gauge group / families	carrier $3+2$, $\dim S^+=16$, $N_{\text{fam}}=3$, $\Omega_{\text{adm}}=48$	[I]
hypercharge	roots of the charge polynomial $bsY^2 + (s-b)Y - 1$	[I]
$U(1)$ index b_1	$b_1 = \frac{g_{\text{car}} 2^{g_{\text{car}}-2} + 1}{10} = \frac{41}{10}$	[I]
fine structure α^{-1}	unique root of $F_{U(1)}(\alpha) = 0$	[I]/[N]
flavor matrix	transport theorem (H1 proved, H2/H3 force the matrix)	[I]/[P]
electroweak scale	$v_{\text{EW}} \sim e^{-\alpha^{-1}/5}$ (carrier divides α^{-1} by 5)	[I]
cosmological constant	$\Lambda \sim e^{-2\alpha^{-1}}$ (density quotient; see typing note)	[I]/[P]
Hubble scale	$H_0 \sim \sqrt{\Lambda}$	[P]
inflation amplitude	$A_s = \frac{N_*^2}{24\pi^2} c_3^7$ in the R^2 branch	[P]
E_8	$D_5 \oplus A_3 + \mu_4$ -glue	[L]

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

Notable: the scale grammar is *one* exponential engine. The same number $\alpha^{-1} \approx 137$ generates the electroweak scale (divided by 5 = carrier), the cosmological constant (times 2) and — via the square root — the Hubble scale. One constant becomes many orders of magnitude.

Typing note on Λ (a deliberate danger zone). “ Λ ” denotes *three* distinct objects that must not be conflated: (i) the seam-determinant exponent $e^{-2\alpha^{-1}}$ (an exact [I] readout), (ii) the dimensionless density quotient $\rho_\Lambda / \bar{M}_{\text{Pl}}^4$ (which carries a prefactor *branch*, hence [P]), and (iii) the observed curvature, to which (ii) is *anchored*, not predicted. The exponential is exact; the absolute dark-energy density remains a typed cosmology interface (full disambiguation in `tfpt_2_standard_model`).

Cosmology is the Pascal action grammar of the carrier — not a separate module

With $x = v_{\text{geo}} / (5\beta_{\text{rad}}^2 \bar{M}_{\text{Pl}})$ the cosmological readouts are a *power tower* on the five-slot carrier graph K_5 :

$$\frac{H_0}{M_{\text{Pl}}} = x^5 R_H, \quad \frac{\rho_\Lambda}{M_{\text{Pl}}^4} = x^{10} R_\Lambda, \quad (1, 5, 10) = \binom{5}{0}, \binom{5}{1}, \binom{5}{2}.$$

EW is the unit x^1 , Hubble is the product over the $\binom{5}{1} = 5$ vertices, Λ is the product over the $\binom{5}{2} = 10$ edges — the *same* Pascal triple that reads one generation, the action ladder and the E_8 root count $16 \cdot 5 \cdot 3 = 240$. “Hubble = slot product, Λ = pair product.” [I]

The entire fermion spectrum in one formula

Masses, Yukawa, CKM, PMNS and neutrinos all follow from *one* master formula with *one*

seed:

$$\hat{m}_{f,j} = \frac{v_{\text{geo}}}{\sqrt{2}} \lambda_Y^{L_{f,j}} \Lambda_{f,j}, \quad \lambda_Y = \sqrt{\varphi_0^{\text{ret}}(1 - \varphi_0^{\text{ret}})}, \quad \varphi_0^{\text{ret}} = \frac{1}{6\pi} + \frac{3}{256\pi^4}.$$

The word-lengths $L_{f,j}$ are the *fixed residue matrix of the compiler*, whose characteristic polynomial $t^3 - 9t^2 + 10t - 8$ carries only compiler numbers ($\text{tr} = 9 = N_{\text{fam}}^2$, $\det = 8 = h(D_5)$, principal 2-minors 2,3,5 with product $30 = h(E_8)$). Through the φ_0^{ret} -ladder $\hat{m} = \frac{v_{\text{geo}}\pi}{\sqrt{2}} c(\varphi_0^{\text{ret}})^k$ all nine word-lengths are fixed and the *charged-lepton* amplitudes are closed; the *quark mass ratios* are closed too (integer Plücker readouts on the derived selector stratum, $v49/v71$), with only the *absolute* amplitude scale reducing to the (U_{wall}) anchor. **SM status:** the discrete packet, the dimensionless flavor skeleton, the charged-lepton source masses and the quark *ratios* are closed; the *absolute* quark scale and the full PMNS completion are reduced/conditional; $m_W, m_Z, m_H, \sin^2 \theta_W, \alpha_s$ are RG scheme-layer (m_H rests on the closed UV quartic $\frac{1}{16\pi^2}$); and the solar angle is *realized by a TFPT texture* and reduced to the seam-misalignment lemma ($\sin^2 \theta_{12} = \frac{1}{3} - \frac{\varphi_0^{\text{ret}}}{2}$; see below). (*Full derivation: document 2, tfpt_2_standard_model.*)

9 Physical, geometric, topological: which solution is more plausible?

Level	E_8 peripheral (v4.5)	E_8 compiled (new)
physical	α^{-1} already a cubic fixed point, but each sector closed on its own	α^{-1} as the fixed point of <i>one</i> equation; one generator for many scales
geometric	Lie algebra external; E_8 only orders the cascade	the C^+ spinor + puncture geometry $\mathbb{P}^1 \setminus \mu_4$ give D_5, A_3 , glued to E_8
topological	—	the common discriminant \mathbb{Z}_4 (μ_4 winding) forces the glue; Φ_{30} /Coxeter drives the sectors
parsimony	many structures	2 axioms, the rest a consequence

Why the compiled solution is more plausible

It is (i) *more parsimonious* (fewer independent assumptions \Rightarrow less overfitting risk), (ii) *more rigid* (the glue condition $5/4 + 3/4 = 2$ and the discriminant \mathbb{Z}_4 leave little freedom), and (iii) *more checkable* (machine-checkable E_8 appendix, explicit fixed-point equation for α^{-1}). More explained from less — exactly the criterion that makes a theory more probable.

10 What role E_8 now plays exactly

E_8 is no longer the starting point but the *output format*:

- **Bookkeeping:** E_8 is the unimodular container in which carrier (D_5) and family (A_3) fit together consistently. Its numbers 240, 248 are carrier traces, not inputs.
- **Consistency test:** unimodularity ($\det 1$) is a hard condition; it closes only because the discriminant-form norms $5/4$ and $3/4$ (both already present in TFPT) add up to 2. E_8 thus tests the internal coherence of the two atoms.
- **Not an end in itself:** E_8 does not explain “everything”; it is the smallest even unimodular hull of the datum $D_5 \oplus A_3$. The physics sits in the *two inputs*, not in E_8 .

New — the secondary atlas. Beyond that, E_8 is an *audit container*: each maximal subalgebra projects a TFPT module. Document 3 (`tfpt_3_e8_audit_bootstrap`, Part I) shows the seven slices of 248 explicitly and labelled. The strongest new hit:

$E_6 \times A_2$ reads the flavor matrix

$$248 = \|R\|_F^2 + \det R + 2(\mathbf{1}^\top R a) N_{\text{fam}} = \underbrace{78}_{\dim E_6} + \underbrace{8}_{\dim A_2} + 2 \cdot 27 \cdot 3,$$

with $\|R\|_F^2 = 78 = \dim E_6$, $\det R = 8 = \dim A_2 = h(D_5)$. The discipline rule is: *every load-bearing number must appear in at least one E_8 projection* — this makes the structure falsifiable rather than numerological. (Audit level; all identities machine-verified.)

11 What this means for the theory-of-everything question

The state of the five questions — closed, reduced or typed

- Flavor residual — ratios closed; only the absolute scale is the wall-selection.** Mehta–Seshadri supplies the equivalence-class framework (stable parabolic \leftrightarrow unitary representation); the algebraic selector R , the Q spectrum and the splitting type are fixed (now *derived*: $\det R = 8$ lattice, $\text{Spec}(Q_+)$ from D_4 , $\sqrt{69}/\sqrt{70}$), and the charged-lepton amplitudes are closed. The quark mass *ratios* ($c_u/c_d = \frac{55}{117}$ etc.) are then *constant* on this discrete selector stratum (Readout Rigidity, $\sqrt{49}/\sqrt{71}$) — integer Plücker readouts, no transcendental solve. Only the *absolute* amplitude scale reduces to the selected *polystable* point on the parabolic stability wall (U_{point} , an anchor; Research Contract 1). [I]/[A]
- Gravity — exact in the R^2 regime.** Spectral action $\rightarrow R+R^2$, BRST/Hodge kills the gauge directions, FRW reduction \rightarrow Starobinsky; the only remainder is the *controlled* low-curvature truncation ($R^{n \geq 3} \rightarrow R^2$, valid for curvature $\ll M^2$). [I] in the regime
- Strong CP — decoupled.** $\theta_{\text{eff}}=0$ follows from three structural facts (γ_5 -Hermiticity, polar structure, sheet involution) + reflection positivity — *without* a mass gap. [I] (given RP)
- Full dynamics — the admissible transfer model is closed.** The mass gap is the *explicit* discrete gap $6 \log \frac{3}{2}$ of the C_6 transfer operator (eigenvalues $(k/3)^6$), which supports the OS reconstruction path on the admissible sector; the continuum and metric-coupled theory still require the stated RP, cluster and limit assumptions. [I] (transfer model) / [P] (continuum)
- Axioms P1/P2 — P2 machine-verified.** The Lean proof builds cleanly (1886 jobs, 0 sorry, only kernel axioms); the P2 algebra is a theorem, P1 a typed interface. They remain declared inputs. [I]/[A]

In other words: the theory has shrunk from “many open audit points” to “*standard steps* (cluster expansion, Fredholm, truncation bound) + two declared axioms”. The five questions are now *closed, reduced or typed* rather than open-ended — a qualitative jump in discipline, not a claim of completion. (*Details and proof chains: document 2, tfpt_2_standard_model, Part III.*)

The TOE-completion ledger — what is done, and what genuinely remains

A theory-of-everything must close the discrete structure, the dimensionless constants, the mass spectrum, gravity, and the quantum measure. The current state, graded honestly:

TOE piece	Status	If not complete: what closes it
gauge group, N_{fam} , hypercharge, b_1	[I]	— complete
$E_8 = (D_5 \oplus A_3) + \mu_4$, group closure	[L]	— complete
bootstrap $(g, \mu) = (5, 4)$, “8” in c_3	[L]	— complete (reverse glue $\mu^2 - 5\mu + 4 = 0$)
fine structure α^{-1}	[I]	— now existence+uniqueness proved (this round)
fermion masses (φ_0^{ret} -ladder)	[I]/[A]	leptons exact; quark <i>ratios</i> closed (Plücker, v49/v71); absolute scale = anchor
mixings $\theta_{13}, \theta_{23}, \theta_{12}$	[N]/[P]	θ_{12} cond. derived; θ_{23} maximal
flavour matrix / H2	[I]	up to Mehta–Seshadri unitarity (U), a Paper-3 input
inflation A_s, n_s, r , scalaron M	[N]	parameter-free given $M_{\text{Star}} = M_{\text{scal}}$
<i>genuinely open — the actual TOE remainder</i>		
dim. EW/QCD masses	[P]	standard RG threshold matching (no new physics)
QFT measure (admissible sector)	[I]	closed: RP from cusp spectrum, OS reconstruction
ambient QFT + full gravity (metric sector)	[A]	<i>not constructed here</i> (frontier doc); conditionally reduces to one named hypothesis (ambient RP) on the relative spectral-action measure
analytic boundary origin of π ($c_3 = \frac{1}{8\pi}$)	[P]	hardened: $c_3 = 1/(\mathbb{Z}_2 \oint_{S^2} K dA) = 1/(8\pi)$ (Gauss–Bonnet); “8” = rank E_8 ; π stays the one primitive
$\eta_B, m_p/m_e$, Koide, dark matter	[A]	frontier handles only (see document 4)

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

The honest bottom line: the discrete core, the dimensionless constants, the *mass-hierarchy exponents* and the mixings are *closed*; the nine masses follow from one φ_0^{ret} -ladder (lepton rationals exact, quark $O(1)$ digits still a conditional finite evaluation). The *admissible-sector* QFT measure is closed (a finite gapped transfer system, RP from the cusp spectrum, OS-reconstructed); the *ambient/metric-sector* measure — i.e. *full quantum gravity* — is *not* constructed here and stays [A] (frontier doc), reducing conditionally to a single named hypothesis (ambient reflection positivity). The seam seed is *hardened*: $c_3 = 1/(|\mathbb{Z}_2| \oint_{S^2} K dA) = 1/(8\pi)$ is the one-sided Gauss–Bonnet normalisation of the seam sphere, with the discrete “8” equal to rank E_8 and only the Gauss–Bonnet π left as the single irreducible primitive. What remains for a full TOE is then the ambient-RP hypothesis (full QG), plus the standard-RG dimensional masses and the honest frontier items. No load-bearing discrete parameter remains free, and π is the lone continuous primitive.

12 Does this make the theory more probable?

Plausibility balance

Pro: (a) two inputs instead of many; (b) α^{-1} to 2.9×10^{-10} (CODATA-2022, 1.9σ) as a *fixed point*, not a fit; (c) E_8 constructed, not postulated; (d) one compiler generates several sectors \Rightarrow fewer degrees of freedom, more predictive power.

Contra/open: what remains are only *standard steps* (cluster expansion in the thermodynamic limit, the Fredholm property, the truncation bound) and the two declared axioms P1/P2; on the SM side the dimensional EW/QCD masses ($m_W, m_Z, m_H, \sin^2 \theta_W, \alpha_s$) sit on the RG scheme layer. The solar angle, previously the one open SM number, is now *conditionally derived*

$(\frac{1}{3} - \frac{\varphi_0^{\text{ret}}}{2})$, modulo the seam-misalignment lemma). A final verdict “it maps reality” requires this short list plus the near-term experimental tests (below).

Net: the parsimony (Kolmogorov-like compression: many numbers from little) and the high precision of individual closures raise the a-posteriori plausibility substantially — without hiding the open points.

13 Predictions for running and upcoming experiments

Because *one* generator now sits behind the numbers, the theory makes a set of concrete, falsifiable statements that running or near-term experiments are actively testing. Each row gives the TFPT value, the *new* derivation behind it, the experiment, and an honest status.

Observable	TFPT value & derivation	Experiment (running / upcoming)	St.
solar angle $\sin^2 \theta_{12}$	$\frac{1}{3} - \frac{\varphi_0^{\text{ret}}}{2} = 0.3067$; TBM + seam misalignment $\varepsilon = \frac{3}{4}\varphi_0^{\text{ret}} \approx c_3$ (cond.)	JUNO (taking data since Aug 2025; sub-percent θ_{12}) — the sharpest live test	[N]/[P]
reactor angle $\sin^2 \theta_{13}$	$\varphi_0^{\text{ret}} e^{-5/6} = 0.0231$; seed \times carrier-trace $e^{-\gamma}$, $\gamma = \frac{5}{6}$	Daya Bay / RENO / JUNO (PDG 0.0220)	[I]
atmospheric $\sin^2 \theta_{23}$	$\approx \frac{1}{2}$ ($\mu\tau$ -symmetric limit; octant not selected)	NOvA / T2K / DUNE (octant ambiguity, NuFIT 6.0)	[P]
fine structure α^{-1}	137.0359992 as the unique root of $F_{U(1)}(\alpha)=0$ (no fit, no second branch)	CODATA-2022 137.035999177(21): dev. 2.9×10^{-10} from the recommended adjustment (1.9σ of its uncertainty)	[I]/[N]
tensor ratio r	$r = \frac{12}{N_*^2} = 0.0040$ ($N_*=55$); R^2 attractor, $M=c_3^{7/2}\bar{M}_{P1}$	now $r < 0.036$ (BICEP/Keck BK18); CMB-S4 ($\sigma_r \leq 5 \times 10^{-4}$, obs. ~ 2033)	[P]
tilt n_s	$n_s = 1 - \frac{2}{N_*} = 0.964$ (same attractor)	Planck (0.965); CMB-S4 sharpens	[P]
amplitude A_s	$\frac{N_*^2}{24\pi^2} c_3^7 = 2.0 \times 10^{-9}$, parameter-free	Planck (2.1×10^{-9})	[P]
neutron EDM	$\theta_{\text{eff}} = 0 \Rightarrow d_n$ essentially null (structure + RP)	PSI nEDM / SNS (running)	[I]
$0\nu\beta\beta$ / ordering	Majorana branch; normal-ordering preference	LEGEND / nEXO , DUNE, KATRIN	[P]
flavor invariants	$\det R=h(D_5)=8$, principal minors (2, 3, 5), $\chi_R=t^3-9t^2+10t-8$ (exact)	any future CKM/PMNS global fit must satisfy these	[I]
H_0 vs. Λ	$v_{\text{EW}} \sim e^{-\alpha^{-1}/5}$, $\Lambda \sim e^{-2\alpha^{-1}}$, $H_0 \sim \sqrt{\Lambda}$: one engine	SHOES / DESI / Planck (Hubble tension)	[P]

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

The two decisive near-term tests

(1) **JUNO and θ_{12}** . JUNO is *taking data since August 2025* and targets sub-percent $\sin^2 \theta_{12}$. TFPT commits to the conditional $\frac{1}{3} - \frac{\varphi_0^{\text{ret}}}{2} = 0.3067$ (seam $\varepsilon = c_3$). A measured central value away from ~ 0.307 at high significance falsifies the seam mechanism — a *live* test.

(2) **r** . The R^2 scalaron with M fixed by c_3^7 predicts $r = 0.0040$ ($N_*=55$), already below the current bound $r < 0.036$ (BICEP/Keck BK18) and within reach of CMB-S4 ($\sigma_r \leq 5 \times 10^{-4}$, first observations $\sim 2033-34$). A detection of $r \gtrsim 0.01$ would be incompatible with the R^2 branch on which M_{P1} and A_s rest.

Structural (non-experimental) checks. Independent of any apparatus, the machine-checkable E_8 appendix (all 240 roots, Gram matrices, discriminant groups) lets anyone verify that the two

TFPT atoms really generate E_8 ; and the discipline rule “*every load-bearing number appears in at least one E_8 projection*” (atlas note) turns the number stock into a falsifiable raster rather than numerology.

Freeze file: committed kill criteria

To prevent post-hoc rescue, the predictions above are frozen with explicit kill criteria in the machine-readable registry `verification/freeze_file.csv`. The decisive entries:

- **solar** $\sin^2 \theta_{12} = \frac{1}{3} - \frac{\varphi_0^{\text{ret}}}{2} \approx 0.3067$ — a JUNO central value clearly away from 0.307 kills the seam-misalignment mechanism;
- **tensor** $r \approx 0.0040$ — any robust $r \gtrsim 0.01$ kills the R^2 branch carrying M_{Pl}, A_s ;
- **ordering** normal, small $m_{\beta\beta}$ — inverted ordering or large $m_{\beta\beta}$ kills the Majorana branch;
- **strong CP** $\theta_{\text{eff}} = 0$ — a solid nEDM signal above SM background falsifies the structural cancellation;
- $w = -1$ — a robust $w \neq -1$ kills the single-engine dark-energy readout.

The proton/electron ratio is explicitly *not* claimed as a compiler power (only fails if mis-asserted as one).

What likely follows next

- the quark mass *ratios* are now closed combinatorially (integer Plücker readouts on the derived selector stratum, $v49/v69/v71$); the only residual is the *absolute* amplitude scale, an anchor (U_{point}) of the same nature as the one dimensionful scale — no new physics;
- further constants as carrier traces (analogous to $b_1 = 41/10$), since the compiler systematically produces traces over S^+ ;
- a sharper geometry \leftrightarrow gravity bridge once the Hodge/ R^2 branch is lifted from [P] to [I].

Conclusion

The state of Papers 1–7 was an impressive collection of closures sharing common numbers. The compiler closure explains *why* they are the same numbers: a small discrete generator $\mathbb{Z}_{30} = 2 \cdot 3 \cdot 5$ builds, from two inputs $\{c_3, g_{\text{car}}\}$ via $D_5 \oplus A_3 + \mu_4$, the E_8 container and outputs the Standard Model, the constants and the scale grammar from there. The theory thereby becomes *more parsimonious, more rigid and more checkable*. Precisely nameable gaps remain — one geometric flavor realisation, the conditional gravity/QFT closures, and two declared axioms — but the path to closure is now short and concrete.

Appendix: computational verification (reproducibility)

Every claim in this document marked [I] (exact identity), [L] (Lie/lattice theorem) or [N] (numerical fixed point) is re-derived from the two axioms $\{c_3, g_{\text{car}}\}$ alone by a small, self-contained Python suite in the repository folder `verification/`. The inline references [`verification/...`] throughout the documents point to the exact script; the table below is the master map.

Script	What it machine-checks
v1_e8_glue.py	glue theorem $E_8=(D_5\oplus A_3)+\mu_4$: disc $=\mathbb{Z}_4$, $q(D_5)+q(A_3)=\frac{5}{4}+\frac{3}{4}=2$, $240=16\cdot 5\cdot 3$, $248=240+8$
v2_carrier_pascal.py	$g_{\text{car}}=5$ unique Pascal solution; $16=1+5+10$, $N_{\text{fam}}=3$, rank $E_8=8$, $\Omega_{\text{adm}}=48$, $b_1=\frac{41}{10}$
v3_em_alpha.py	EM closure: unique root of $F_{U(1)}(\alpha)=0$, $\alpha^{-1}=137.0359992168$, existence/uniqueness, inverse test, ablation
v4_flavor_matrix.py	residue matrix R : det $=8$, minors $\{2, 3, 5\}$, $\chi_R=t^3-9t^2+10t-8$, SNF diag $(1, 1, 8)$, $\ R\ _F^2=78$, $\sum L=40$
v5_e8_cascade.py	cascade $D_n=60-2n$: endpoints, exponent rungs $\rightarrow 240$, IR tail 46080, variance 6552
v6_bootstrap.py	reverse glue $\mu^2-5\mu+4=0$; $g_{\text{car}}=5$ three ways; $8=\text{rank } E_8=h(D_5)=\varphi(30)$
v7_gravity_cosmo.py	scalaron exponent 7, $M_{\text{scal}}^2/\bar{M}_{\text{Pl}}^2=c_3^7$, $A_s, n_s, r, \Omega_b, \sin^2 \theta_{12}$
v8_horizon.py	horizon code: $\frac{1}{2\pi}=4c_3$, $1920= W(D_5) $, S_{dS} , Page, $\beta_{\text{rad}}=0.2424^\circ$
v9_neutrino_texture.py	polar angle: explicit $\mu\tau$ Majorana texture $\rightarrow \sin^2 \theta_{12}=\frac{1}{3}-\frac{\varphi_0^{\text{ret}}}{2}$, $\theta_{23}=45^\circ$, $\theta_{13}=0$
v10_projection_involutions.py	Q, Σ algebra: $K=R+Q\Sigma$, $L=R+Q(I+\Sigma)$; χ_Q, χ_K ; det ladder $3\cdot 4\cdot 8\cdot 20=1920= W(D_5) $; $a^\top K a=41$
v11_unique_KQ.py	K, Q are the <i>unique</i> nonneg-int matrices with their compiler budgets, spectrum and monotone rows (enumeration)
v12_mass_generation_polynomials.py	sector/generation polynomials of K and the anchor-block det ladder $(9, 10, 16, 40)$
v13_open_gates.py	gate closures: $M=41=10b_1=\sum L+N_\Phi$; $Q_+ = A_3$ exponents, $Q_-^2=N_{\text{fam}}$
v14_carrier_uniqueness.py	$g_{\text{car}}=5$ unique ($N_{\text{fam}} \in \mathbb{Z}^+$, rank $E_8=8$); split $(3, 2)$ from $b+s=5$, $b^2+s^2=13$; $\text{Tr } Y = \text{Tr } Y^3=0$
v15_bootstrap_classification.py	$D_5\oplus A_3$ unique familyful cyclic-gluon of E_8 (16-spinor + 3 families), glue \mathbb{Z}_4
v16_solar_dual_anchor.py	$a^\top R^{-1}=a^\top L^{-1}=(-\frac{1}{2}, -\frac{1}{2}, 1)$; $\sin^2 \theta_{12}=\frac{1}{3}-\frac{\varphi_0^{\text{ret}}}{2}$; full PMNS
v17_hexagonal_resolvent.py	finite resolvent $D_y^{-1}=(\sum y^{5-m}\delta^m U_6^m)/(y^6-\delta^6)$ (quark c backbone)
v18_quark_yukawa.py	quark source ratios $\frac{m_u}{m_d}=0.470085$, $\frac{m_c}{m_s}=13.61$, $\frac{m_t}{m_b}=40.80$ exact; lepton $c=(\frac{16}{7}, \frac{4}{3}, \frac{7}{6})$; full hierarchy
v19_monodromy_moduli.py	exact pole $z_\star=\frac{794-7\sqrt{9961}}{2187}$ (3^7); D_4 $SU(3)$ monodromy constructible but D_4 -fixed locus positive-dim \Rightarrow exact quark c reduce to (U)
v20_lepton_c_derivation.py	lepton c 's derived: $\delta=\frac{1}{2}$ resolvent $c= \mu_4 ^w/(\frac{5}{4}-\cos \frac{r\pi}{3}) \Rightarrow \frac{16}{7}, \frac{4}{3}$; product $\frac{2^{g_{\text{car}}}}{N_{\text{fam}}^2}=\frac{32}{9} \Rightarrow \frac{7}{6}$
v21_solar_product_quark.py	solar coeff $=q(A_3)=\frac{3}{4}$ (glue-norm, $q(D_5)+q(A_3)=2$), $\sin^2 \theta_{12}=\frac{1}{3}-\frac{\varphi_0^{\text{ret}}}{2}$; product $\frac{2^{g_{\text{car}}}}{N_{\text{fam}}^2}$; quark law fails ($\frac{1}{7}$ vs 0.724)
v22_open_gates_audit.py	residual-gates audit contract: exact reduction of each open gate $(A_2, B_3, B_4, B_5, B_6, C_7)$ machine-pinned (forced part vs. named residual)
v23_anchor_generator.py	anchor $a=(1, 1, 2)$: $e_k=(4, 5, 2)$, $c_3=\frac{1}{2e_1\pi}$; $p_n=2+2^n \Rightarrow 240, 248$, binary ladder; inputs $\rightarrow \{a, \pi\}$
v24_quark_ratio_closure.py	quark ratios $\frac{55}{117}, \frac{34}{47}, \frac{3}{26}$ from TFPT blocks; mass ratios match $< 0.03\%$; rational c gauge, Λ_q in $O(1)$
v25_frontier_conjectures.py	conjectures: Koide $Q_{\text{src}}+\frac{\varphi_0^{\text{ret}}}{24}\approx\frac{2}{3}$; axion $f_a=M_{\text{scal}}/128\rightarrow m_a\approx 23.8 \mu\text{eV}$
v26_flavor_frontier_unification.py	the "11" is not uniquely forced (\Rightarrow [P]); flavor frontier unifies to one (U) gate; cusp config on the parabolic stability wall
v27_wall_representative.py	explicit balanced wall rep W_{wall} (cols = perm $(0, 1, 2)$, rows $w=(2, 1, 1)$, pardeg 0); selector det $R=8$, Spec $(Q_+)=\{1, 2, 3\}$
v28_gravity_fR.py	$R+R^2$ closed: $f(R)$, $M=c_3^{7/2}\bar{M}_{\text{Pl}}$ scalaron, $N_\star=57\rightarrow n_s, r, A_s$; full metric measure open; $248c_3^2 < 6 \log \frac{3}{2}$
v29_research_contract_	research-contract certificates: $C_U^{(1)}$ wall enumeration $(1296\rightarrow 144\rightarrow 5)$ orbits; symmetry alone does not pick W_{wall} and G_5 gap-dominance $2\ V\ < \Delta$

How to run. From the repository root, `python verification/run_all.py` (dependencies `mpmath`, `numpy`, `sympy`); each module also runs standalone, e.g. `python verification/v1_e8_glue.py`. The runner exits with code 0 iff all checks pass; see `verification/README.md` for the full claim \leftrightarrow script map and `verification/status_ledger.csv` for the **single status ledger** (claim ID, status, location, dependencies, script, external datum — the source of truth all six documents are written against). The [F] (Lean-formalised) carrier material lives in the canonical shipped folder `lean4-carrier-rigidity/` (at the package root; the source repository keeps the same project under `experiments/`).