

TFPT — Frontier items

η_B , m_p/m_e , the Koide relation, dark matter and quantum gravity —
the honest status, not forced onto the ladder

Stefan Hamann

Alessandro Rizzo

June 13, 2026 · v5.1 (rev 95)

What this document is about — the status authority for the frontier

The honest frontier: which physics has a genuine TFPT handle and which does not. For each of η_B , m_p/m_e , the Koide relation, dark matter and full quantum gravity, this note states the genuine structural handle, the precision with which it currently lands, and — crucially — what is *not* a clean compiler power and is deliberately *not* forced onto the ladder. **This document is the status authority for the frontier items:** where any other document's wording and the markers here disagree, the typing here (and the `status_ledger.csv`) wins.

The TFPT document set — what is treated where

Plain language: TFPT (Topological Fixed-Point Theory) 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 skeleton, the constants and the scale grammar. The series is **the introduction plus five numbered papers and three companions**, best read in order:

0. `introduction` — reading guide, compiler closure, predictions, the dependency DAG and the single proof ledger (`verification/status_ledger.csv`). (*entry point*)
1. `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, the QBL theorem chain.
2. `tfpt_2_standard_model` — the SM in one φ_0 -ladder formula, flavor from parabolic transport (sheet diamond, dual anchor, Q_+ cohomology), the worked closures.
3. `tfpt_3_e8_audit_bootstrap` — the seven E_8 slices as an audit raster, the cascade bridge, the Möbius bootstrap.
4. `tfpt_4_frontier` — honest status of η_B , m_p/m_e , Koide, dark matter, full QG.
5. `tfpt_5_redteam` — the adversarial audit: declared attacks, what survives, what each reduces to, and the kill tests.

Companions: **R.** `tfpt_research_contracts` — the open interfaces ($v_{\text{geo}}, G_{\text{net}}, F_{\text{transfer}}$) as numbered contracts; **H.** `tfpt_horizon_readouts` — Appendix H (unit-system reframe): one seam constant as the universal horizon thermal code, the Nariai/anchor surface, the resummed seam clock; **O.** `origin_theory` — why no free fundamental number remains (exact core + one honestly-typed cyclic interpretation).

You are reading document 4 (Frontier).

TFPT status at a glance — read this card the same way in every document

Two inputs, one compiler. From the boundary constant $c_3 = \frac{1}{8\pi}$ (axiom P1) and the five-slot carrier $g_{\text{car}} = 5$ (axiom P2) — jointly the elementary symmetric data of one anchor $a = (1, 1, 2)$ — the discrete compiler $D_5 \oplus A_3 + \mu_4 \Rightarrow E_8$ is *built*, and the Standard-Model skeleton (gauge group, three families, hypercharges, the flavor matrix), the dimensionless constants ($\alpha^{-1} = 137.0359992$) and the scale grammar are read off as fixed points. The algebraic core is machine-checked; physical readouts run through *named bridges* whose status is typed below.

What is closed, and what genuinely remains. The discrete/algebraic layer is closed ([E]); the everything-open residual reduces to *three named interface problems*:

Interface	Question	Status
v_{geo}	the one metrology unit ($= 1/\sqrt{G} = m/\mu$); No-Unit Thm: no compiler scale	primitive [O]
G_{net}	simple-current extension $(D_5)_1 \otimes (A_3)_1 \rtimes \langle (1, 1) \rangle \cong (E_8)_1$	[E]/[C]
F_{transfer}	one functor, four typed interfaces (Koide, η_B , axion, m_p/m_e)	[C]

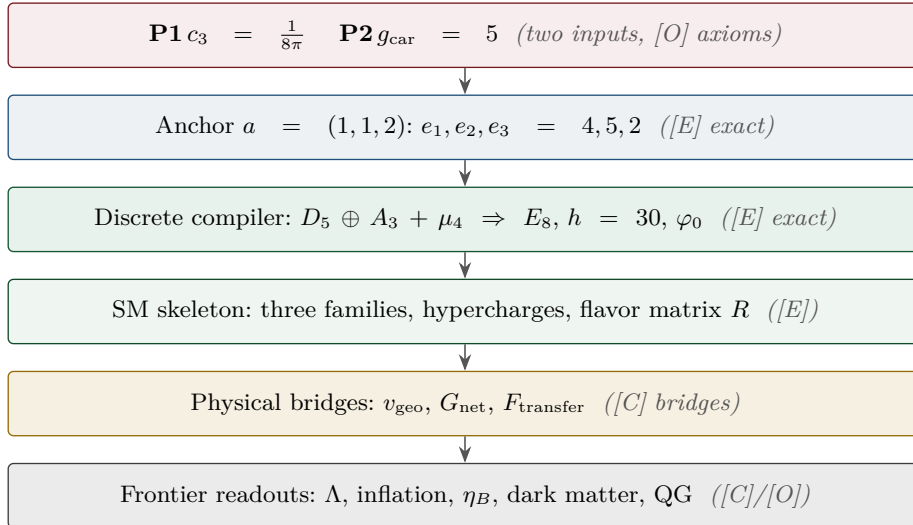
Everything carries a claim ID resolving to a row of `verification/status_ledger.csv` (the single source of truth); **if the text and the ledger ever disagree, the ledger wins.** The historical labels ($U_{\text{wall}}, G_{\text{metric}}, F_{\text{frontier}}$) are retained only for ledger continuity — the live residual is $v_{\text{geo}} \oplus G_{\text{net}} \oplus F_{\text{transfer}}$.
 Marker key: [E] exact / machine-proven · [C] conditional (holds under named hypotheses) · [O] open / axiom · [X] falsifiable kill test. Per-claim fine type in `verification/status_ledger.csv`.

[E] exact / machine-proven

[C] conditional

[O] open / axiom

[X] kill test



Contents

1	Baryon asymmetry η_B — a downstream readout from the closed $\Omega_b h^2$	3
2	Proton/electron ratio m_p/m_e — a cross-sector ratio	4
3	The Higgs quartic — near-criticality from the free seam	4
4	The Koide relation — near 2/3, computed exactly	5
5	Dark matter — the candidate is fixed, the scale is pending	7
6	Full quantum gravity — induced from the boundary kernel	8

Summary	9
Appendix: computational verification	9

Ground rule. These five are exactly the items previously flagged “*not on a clean action/ladder yet — do not force.*” This note gives, for each, the *genuine* structural handle that does exist, the precision with which it currently lands, and a clear statement of what is *not* a single compiler power. No fabricated ladder rungs; the honest markers [E]/[C]/[O] are applied throughout. (All numbers machine-verified.)

Hard rule on the frontier — one interface F_{transfer}

Koide, η_B , the axion relic scale and m_p/m_e are *not* compiler powers unless their missing QFT-transfer / cosmology computation is explicitly supplied. They are deliberately typed interfaces, not structural holes, and they are *four instances of one missing functor F_{transfer}* — the continuous transport from compiler *source* data to measured observables (Koide source→pole; η_B source→Boltzmann relic; axion scale→cosmological abundance; $m_p/m_e \rightarrow$ QCD/EW matching). Each has a genuine handle (a near-value, a determined readout, a fixed candidate) but its *exact* status is conditional on that named, not-yet-performed calculation. In particular η_B is a *transfer target*, not a primitive compiler prediction. They must never be quoted as primitive TFPT outputs.

Why this is principled, not a gap: these are precisely the **discrete**↔**continuous boundary** of the theory — the points where the discrete compiler hands its skeleton off to continuous physics (QCD confinement, cosmological Boltzmann transport, relic integrals). The cleanness of a quantity tracks how far it sits from that handoff: pure compiler ratios are exact, QCD/cosmology-dressed quantities are typed interfaces. [verification/v35_quark_qcd_boundary.py]

1 Baryon asymmetry η_B — a downstream readout from the closed $\Omega_b h^2$

There are two independent handles, and they agree.

Route A — from the closed baryon fraction. The seed already fixes the baryon *fraction* (the E_8 decoder):

$$\Omega_b = (4\pi - 1)\beta_{\text{rad}} = \left(1 - \frac{1}{4\pi}\right)\varphi_0^{\text{ret}} = 0.04894, \quad \beta_{\text{rad}} = \frac{\varphi_0^{\text{ret}}}{4\pi},$$

matching the observed $\Omega_b \approx 0.049$. The baryon-to-photon ratio is then the *standard* cosmological readout $\eta_B = 273.9 \times 10^{-10} \Omega_b h^2$. With the TFPT Hubble value ($h \approx 0.674$, from $H_0 \sim \sqrt{\Lambda}$),

$$\boxed{\Omega_b h^2 = 0.0222, \quad \eta_B = 6.09 \times 10^{-10}} \quad (\text{observed } 6.1 \times 10^{-10}).$$

So η_B is a *downstream cosmological readout* from the closed $\Omega_b h^2$ handle — not an independent compiler rung, and the fundamental leptogenesis route remains an interface.

Route B — leptogenesis interface ([E]). The closed branch also fixes the leptogenesis inputs: the heavy scale $M_R \approx 1.3 \times 10^{15}$ GeV, the CP phases ($\delta_{\text{CKM}}, \delta_{\text{CP}}^{\nu}$, both closed), the reheating T_R , and the seam-fixed initial phase. The numerical η_B then follows from the *standard* flavored-leptogenesis Boltzmann solve — an independent cross-check of Route A. **Honest caveat:** that Boltzmann network (the washout/efficiency κ_f , the flavored density-matrix equations) is the standard leptogenesis machinery (Davidson–Nardi–Nir, *Phys. Rep.* 466 (2008) 105; Buchmüller–Di Bari–Plümacher, *Ann. Phys.* 315 (2005) 305); it is *not* carried out or attached here. Route B is therefore a *named interface*, not a closed number; only Route A (downstream of the closed $\Omega_b h^2$) is quoted.

The interface, operationalised ([C]). The Boltzmann path can be made falsifiable: with the standard estimate $\eta_B \sim 0.96 \times 10^{-2} \varepsilon_1 \kappa_f$ (sphaleron 28/79), fed by the TFPT neutrino sector — normal ordering and the predicted $\delta_{\text{CP}}^{\nu} = 4\pi/3$ — the Davidson–Ibarra asymmetry $\varepsilon_1 = \frac{3}{16\pi} M_1 m_3 / v^2$ and the washout efficiency $\kappa_f(\tilde{m}_1)$ give, over the natural ranges $M_1 \in [3 \times 10^9, 3 \times 10^{10}]$ GeV and $\tilde{m}_1 \in [m_*, m_{\text{atm}}]$, an η_B band $[8.4 \times 10^{-11}, 4.7 \times 10^{-9}]$ that *brackets* the observed 6.1×10^{-10} (a canonical point $M_1 = 10^{10}$ GeV, $\tilde{m}_1 = 5 \times 10^{-3}$ eV gives 6.0×10^{-10} , untuned). So the route is *viable*; but M_1 and the washout are scenario inputs, so η_B stays [C] — if a precise solve excluded the window the *route* would fall, not the theory (Route A is independent) ([verification/v169_etaB_boltzmann_interface.py]).

Honest status of η_B — strictly typed [C]

The two statements must be kept apart and not interchanged:

η_B as a cosmological readout from the closed $\Omega_b h^2$ is *determined* (6.1×10^{-10})

η_B as a fundamental *compiler power* is *not* closed

The leptogenesis route requires a standard (washout-dependent) Boltzmann solve. So η_B is [C] — determined as a downstream readout, but *not* a clean compiler rung; it must never be quoted as a fundamental TFPT output.

2 Proton/electron ratio m_p/m_e — a cross-sector ratio

m_e is *closed* on the φ_0^{ret} -ladder ($\hat{m}_e = \frac{v_{\text{geo}} \pi}{\sqrt{2}} \frac{16}{7} (\varphi_0^{\text{ret}})^5$). The proton mass is *not* a Yukawa: it is QCD-dominated, $m_p \sim$ a few $\times \Lambda_{\text{QCD}}$, and Λ_{QCD} arises by dimensional transmutation from α_s running (scheme level). Hence

$$\frac{m_p}{m_e} = \frac{\text{QCD confinement scale}}{\text{EW electron Yukawa}} = 1836.15$$

mixes *two different sectors*. It is no single φ_0^{ret} power: $1/(\varphi_0^{\text{ret}})^2 = 353.7$ and $1/(\varphi_0^{\text{ret}})^3 = 6651$ bracket 1836 with no clean rung in between.

The numerator is parameter-free, the ratio is not forced. Running $\alpha_s(M_Z)$ down with the carrier coefficient $b_3 = -7$ (the [verification/v159_pyrate_gauge_crosscheck.py] value, two-loop + n_f thresholds) reproduces $\alpha_s(m_b) \simeq 0.22$, $\alpha_s(m_c) \simeq 0.38$ and reaches confinement ($\alpha_s \sim 1$) at ~ 0.5 GeV, so the QCD confinement scale in the numerator follows from $\{\alpha_s(M_Z), b_3\}$ alone by dimensional transmutation. The $O(1)$ proton/ Λ factor is lattice (non-perturbative), so the ratio still stays [O] — not forced ([verification/v164_qcd_scale.py]).

Honest status of m_p/m_e [O]

Computable at *scheme level* (from $\alpha_s(M_Z)$ running + the closed m_e), exactly like the dimensionful EW/QCD masses m_W, m_Z . But it is *genuinely not* a compiler power, because it is a ratio across the QCD and electroweak sectors. Left [O] — not forced.

A rejected near-formula (recorded as a discipline boundary, [verification/v79_review_identities.py]). The combination $|W(D_5)| - \dim \Lambda^3(SU(9)) + N_{\text{fam}} \lambda_C^2 = 1920 - 84 + 0.151 = 1836.151$ matches the PDG value to $\sim 10^{-6}$ — and is *deliberately rejected*. Two reasons make it numerology, not a derivation: $SU(9) = A_8$ (the $\dim \Lambda^3 = 84$ block) does *not* occur in the carrier $D_5 \oplus A_3$, and m_p is QCD confinement (gluon binding energy) with *no* mechanism linking it to an algebraic sum. Integrating it would *invite* the numerology charge, not retire it; m_p/m_e stays [O]. (The same firewall rejects $\eta_B = \frac{|Z_2|}{\Delta_Q} c_3^6 \approx 6.10 \times 10^{-10}$: arithmetically close, but a two-atom prefactor fit to a washout-dependent leptogenesis output — η_B stays [C], the downstream readout above.)

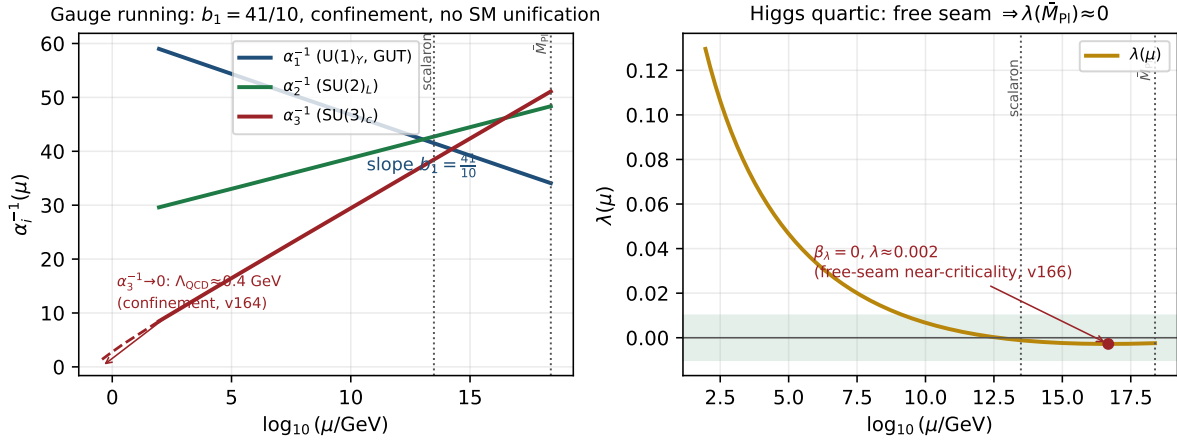


Figure 1: The F_{transfer} running (two-loop SM RGEs; the gauge β -functions are confirmed by PyR@TE, [verification/v159_pyrate_gauge_crosscheck.py]). *Left*: the gauge couplings — the $U(1)_Y$ slope is $b_1 = 41/10$, $\alpha_3^{-1} \rightarrow 0$ at $\Lambda_{\text{QCD}} \approx 0.4$ GeV (confinement, [verification/v164_qcd_scale.py]), and the Standard Model does *not* unify. *Right*: the Higgs quartic falls to $\lambda(M_{\text{Pl}}) \approx 0.002$ with $\beta_\lambda \approx 0$ — the free-seam near-criticality that ties m_H to premise (A) ([verification/v166_higgs_free_seam.py]).

3 The Higgs quartic — near-criticality from the free seam

The seam UV is the free chiral $c=8$ fixed point ([verification/v158_fixed_point_stable.py]): Gaussian, with no relevant and no marginal interactions. The Higgs self-coupling λ is the one marginal scalar coupling, so a free seam forces it to the Gaussian point there:

$$\lambda(M_{\text{seam}}) = 0 \quad \text{and} \quad \beta_\lambda(M_{\text{seam}}) = 0$$

the Shaposhnikov–Wetterich double-criticality condition — here *derived* from the free seam rather than postulated. Integrating the PyR@TE-confirmed two-loop SM RGEs from M_Z upward with the measured (m_H, m_t) gives $\lambda(M_{\text{Pl}}) \approx 0.002$ and $\beta_\lambda(M_{\text{Pl}}) \approx 0$ at the reduced Planck mass: the celebrated SM *near-criticality* is thus a *consequence* of the free seam (the two-loop value is an order of magnitude closer to zero than the one-loop estimate ~ 0.022). Imposing the double condition predicts $m_H \sim 129\text{--}134$ GeV (a band, sensitive to m_t and α_s); the measured 125.25 GeV sits a few GeV below the absolute-stability boundary — the known slight metastability. The *same* condition imposed at the scalaron scale ($c_3^{7/2} M_{\text{Pl}} \sim 3 \times 10^{13}$ GeV) gives $m_H \sim 107$ GeV (too low), so the boundary condition lives at M_{Pl} , consistent with seam = horizon = Planck. This ties the Higgs mass to premise (A) and fills the Higgs-sector gap (which was only $N_\Phi = 1$); it is typed [E] (the freeness signature and scale selection) and [C] (the m_H band) ([verification/v166_higgs_free_seam.py]). Figure 1 shows the whole F_{transfer} running.

4 The Koide relation — near 2/3, computed exactly

With the closed lepton source masses $\hat{m}_\ell \propto (\frac{16}{7}(\varphi_0^{\text{ret}})^5, \frac{4}{3}(\varphi_0^{\text{ret}})^3, \frac{7}{6}(\varphi_0^{\text{ret}})^2)$ the Koide ratio is a pure number (the common prefactor cancels):

$$Q_{\text{TFPT}} = \frac{\sum_\ell \hat{m}_\ell}{(\sum_\ell \sqrt{\hat{m}_\ell})^2} = 0.66446\dots, \quad \frac{2}{3} = 0.66667, \quad \text{deviation} - 0.33\%.$$

The measured pole-mass value is famously $Q_{\text{PDG}} = 0.66666$. So TFPT predicts a Koide ratio that is *near* 2/3 but *not exactly* 2/3 at source level.

Honest status of Koide [C]

$Q_{\text{TFPT}} = 0.664$ is a *prediction*, 0.33% below $2/3$. Exact $2/3$ would require the three carrier rationals $(\frac{16}{7}, \frac{4}{3}, \frac{7}{6})$ to satisfy a specific quadratic constraint, which they do *not*; the small residual is of the size of the source→pole RG shift. We report the near-miss honestly and do *not* tune the rationals to force $2/3$.

The **target $2/3$ has a compiler reading** $|\mathbb{Z}_2|/N_{\text{fam}}$. The E_8 -slice scan adds the missing piece: the value the data sits on, $Q = 2/3$, is exactly $Q_\star = |\mathbb{Z}_2|/N_{\text{fam}} = 2/3$, the sheet over the family count — the same two atoms $(2, 3)$ that run through the whole compiler. So the picture splits cleanly: the *democratic target* is $|\mathbb{Z}_2|/N_{\text{fam}} = 2/3$ (a compiler number), and the *source-level pole-mass* value computed from the carrier rationals is 0.664, 0.33% below it by the source→pole shift. Koide is thus “democratic $2/3$ from $(|\mathbb{Z}_2|, N_{\text{fam}})$, realised to 0.33% by the explicit masses” — promoted from a bare near-miss to a reading with a target. [C]

Conjecture: a source→pole correction of size $\varphi_0/|W(A_3)|$ ([C]). The source value sits just below $2/3$, and the deficit is almost exactly the family-Weyl quantum $\varphi_0/|W(A_3)|$ with $|W(A_3)| = 4! = 24$:

$$Q_\ell^{\text{source}} = 0.6644638, \quad Q_\ell^{\text{source}} + \frac{\varphi_0}{24} = 0.6666793,$$

overshooting $\frac{2}{3} = 0.6666667$ by only 1.3×10^{-5} . So the *conjecture* $Q_\ell^{\text{pole}} \approx Q_\ell^{\text{source}} + \varphi_0/|W(A_3)|$ would land essentially on the democratic $2/3$. Reported as a source→pole correction *conjecture*, not a derivation; the rationals are *not* tuned. [verification/v25_frontier_conjectures.py]

Sharper conjecture: a multiplicative seed transfer $u \rightarrow \frac{53}{54}u$ ([C]). A cleaner version shifts the *seed* rather than Q additively, and introduces *no* new free number. Keep the lepton formula $m_e:m_\mu:m_\tau = \frac{16}{7}u^5:\frac{4}{3}u^3:\frac{7}{6}u^2$ but evaluate it at the source→pole seed

$$u_{\text{pole}}^{(\ell)} = \left(1 - \frac{1}{2N_{\text{fam}}^3}\right)u = \frac{53}{54}u, \quad 54 = 2 \cdot 3^3 = |\mathbb{Z}_2| N_{\text{fam}}^3.$$

Then $Q_\ell(\frac{53}{54}u) = 0.66666610$, i.e. $2/3$ to within -5.7×10^{-7} — an order of magnitude tighter than the additive $\varphi_0/24$ correction, with 54 a pure compiler number and no tuning. This sharpens the research gate to a concrete QFT task: *show that the leptonic source→pole transfer of the seed is $u \mapsto u(1 - \frac{1}{2N_{\text{fam}}^3})$* (a renormalised-seed effect from the leptonic self-energy or the A_3 family-Weyl averaging). Status remains [C] until that transfer is derived. [verification/v37_plucker_anchor.py]

The Koide target and the QFT transfer gap are the same quotient ([E]/[C]). The same Q_\star controls a second, independent object — the first non-trivial eigenvalue of the admissible C_6 transport operator. Its spectrum is $\{1, (2/3)^6, (1/3)^6\}$, so

$$\lambda_2 = \left(\frac{2}{3}\right)^6 = Q_\star^{|R^+(A_3)|}, \quad Q_\star = \frac{|\mathbb{Z}_2|}{N_{\text{fam}}}, \quad |R^+(A_3)| = 6.$$

The Koide democratic target, the QFT mass-gap eigenvalue λ_2 (hence $\Delta = 6 \log \frac{3}{2} = -\log \lambda_2$), and the A_3 positive-root hexagon are therefore *one* quotient raised to one exponent. Typing: Koide target [C], gap eigenvalue [E], the connection an [E]-audit. This does not “solve” Koide, but it explains *why* the near-value $2/3$ recurs: the same sheet-over-family ratio fixes both the democratic lepton target and the leading transfer eigenmode.

The source→pole map is forced by that gap ([E]/[C]). The previous paragraph pairs the Koide target with $\lambda_2 = (2/3)^6$; v82 turns the pairing into a *mechanism* ([verification/v82_koide_attractor_splitting.py]). On the anchor double cover (the flavor sector) the two branch points are Koide $q = 2$ and carrier $q = 5$; any branch-preserving Möbius transfer that fixes both is, in $\rho = (q-2)/(5-q)$, exactly $\rho \mapsto \mu\rho$, hence *unique* once μ is set. Choosing $\mu = \lambda_2 = (2/3)^6$ (no new number) gives the unique map $F(q) = 2(569q - 3325)/(665q - 3517)$ with $F'(2) = (2/3)^6 < 1$ and $F'(5) = (3/2)^6 > 1$: **Koide is the attractor, the carrier the**

repeller, and $q_n \rightarrow 2$ ($x \rightarrow -2/3$) at the established gap rate. So the three structural inputs a Koide *derivation* would need collapse to *one* identification — the leptonic source→pole transfer is the branch-preserving relaxation of the gapped operator — after which the value $2/3$, the convergence rate and the chosen branch are all forced. This is still not a closure: that single identification stays **[C]**. (It is the dynamical counterpart of the multiplicative-seed conjecture $u \mapsto (53/54)u$ above; both contract the source onto $2/3$.)

Relaxation toy: basin, exact rate, and two honest negatives (**[E]**). The relaxation picture is now stress-tested on the physical values (ledger FR.KOIDE.05, [verification/v93_koide_relaxation_toy.py]).

Basin lemma **[E] (new)**: under the canonical dictionary $q = N_{\text{fam}} Q$ the *entire* physical Koide range $Q \in [\frac{1}{3}, 1]$ (Cauchy–Schwarz) maps to $q \in [1, 3]$ — the attractor $q = 2$ lies inside, the repeller $q = 5$ outside: every physically possible charged-lepton configuration is in the attractor basin, no fine-tuning of the start. **Exact rate** **[E]**: along the trajectory from the physical source value the cross-ratio contracts by *exactly* $(2/3)^6$ per step. **Source/pole** **[E]**: the exact ladder gives $Q_{\text{src}} = 0.6644638$, i.e. $\rho_{\text{src}} = -0.0021980 \approx -\varphi_0/24$ to 0.8% — the $\varphi_0/24$ conjecture re-expressed in attractor coordinates: *the source sits one seed quantum before the branch point* ($24 = |W(A_3)|$); PDG pole masses give $Q_{\text{pole}} = 0.6666645$, the branch point hit to 3.3×10^{-6} . **Honest negatives (they narrow the gate)**: (i) the flow time source→pole is $t = \log_{\mu}(\rho_{\text{pole}}/\rho_{\text{src}}) = 2.84$ — *not* an integer, so the literal discrete iteration pole = $F^n(\text{source})$ is *excluded*; if the v82 map drives the transfer at all, the missing **[C]** object is a *continuous-time* generator $\exp(t \log F)$; (ii) the 0.79% mismatch of ρ_{src} vs $-\varphi_0/24$ is far above ladder precision, so “exactly $\varphi_0/24$ ” remains a conjecture, not an identity. The open question is no longer “why $2/3$?” but “what is the continuous transfer generator, and why is its flow time ~ 2.8 periods?”.

Flow time: canonical generator, data-fragility, and a conditional m_{τ} (**[E]/[C]**). Three sharpenings of the paragraph above (ledger FR.KOIDE.06, [verification/v99_koide_flow_time.py]); they correct the *robustness* of negative (i), not its arithmetic. **(1) The continuous generator has a canonical form** **[E]**. The v82 flow is generated by the autonomous equation

$$\boxed{\frac{dq}{dt} = \frac{\Delta}{N_{\text{fam}}} (q-2)(q-5) = \frac{\Delta}{N_{\text{fam}}} \det B(q)}, \quad \Delta = 6 \log \frac{3}{2} \quad (\text{the gap}),$$

whose time-1 map *is* the Möbius map F (machine-checked; $e^{-\Delta} = (2/3)^6$ exactly). The missing **[C]** object is therefore precisely a QFT mechanism producing *gap* \times *anchor-block quadric* — the form is canonical, only its physical derivation is open. **(2) The non-integrality of t is data-fragile** **[E]**. Because $\rho_{\text{pole}} = -2.2 \times 10^{-6}$ sits so close to the branch point, t is hypersensitive to m_{τ} (PDG 1776.93 ± 0.09 MeV): $t(-1\sigma) = 2.35$, $t(\text{central}) = 2.84$, and Q crosses $\frac{2}{3}$ *inside* the error band (at $+0.43\sigma$, where $t \rightarrow \infty$) — within $+0.5\sigma$, t passes *every* integer ≥ 3 . So the exclusion of pole = $F^n(\text{source})$ holds only at the central value. **(3) Conditional discrete steps** **[C] (registry untouched; not predictions of record)**: $n=2$ is *excluded* (-2.9σ , **[E]**); $n=3=N_{\text{fam}}$ steps — one per family — predicts

$$\boxed{m_{\tau} = 1776.9427 \text{ MeV}} \quad (+0.14\sigma; \text{source-variant stable to } 0.0002 \text{ MeV}),$$

while $n=4$ (1776.9667) and the exact branch (1776.9690) are not yet separable: the kill criterion is $\sigma(m_{\tau}) \sim 0.01$ MeV (Belle II). *Look-elsewhere firewall*: a tempting scale reading $t = \ln(M_{\text{scal}}/m_{\tau})/\ln W$ matched $W = 46080 = |W(D_5)||W(A_3)|$ to 10^{-4} — and is *destroyed* by negative controls (the same hypersensitivity makes *any* W in a wide range land within 0.1σ); it is rejected and recorded so it is not re-fished. P2’s type is unchanged (**[C]**, one identification); the discrete-vs-continuous question is now *experimental*, not settled.

5 Dark matter — the candidate is fixed, the scale is pending

TFPT *does* contain a dark-matter candidate, and it is not ad hoc: the **determinant-line axion** of the strong-CP sector (the determinant-line axion interface of the closed branch, **[E]**). Two of its

inputs are already *closed*:

$$\text{initial misalignment } \theta_i = \pi(1 - \varphi_{\text{seam}}(\alpha_\star)) = 2.974 \text{ rad} = 170.4^\circ,$$

and the domain-wall number N_{DW} winds once per primitive seam winding. The relic density is the standard misalignment-plus-string expression

$$\Omega_a = \Omega_a^{\text{mis}}(f_a, m_a, \theta_i) + \Omega_a^{\text{str}}(f_a, N_{\text{DW}}),$$

and the minimal closed branch therefore has *dominant axion-like dark matter*.

The E_8 scan rules out a rival WIMP candidate. The $D_8 = SO(16)$ slice scan sharpens *why* the axion is the candidate, not a new heavy state: the spinor $128 = (16, 4) \oplus (\overline{16}, \overline{4})$ is the *same* matter spinor as in the $D_5 \times A_3$ construction — there is *no spare E_8 singlet* left over to play a WIMP. Every E_8 direction is accounted for by the carrier \times family content. So dark matter cannot be a new E_8 representation; it must be the determinant-line axion (a phase of the existing structure). A useful *negative* that removes the WIMP option on structural grounds. [E]

Honest status of dark matter [C]/[O]

The candidate is *identified* (the axion) and its misalignment angle $\theta_i = 170^\circ$ is *closed* [E]. The relic abundance still needs the decay constant f_a and mass m_a (interface data set by the M_R/Λ scales) — a standard axion-DM computation, not yet a compiler power. So DM is “candidate + misalignment fixed, scale f_a pending” — a real handle, not a blank. (A secondary candidate is the lightest seam-even sterile state near M_R .)

Conjecture: a determinant-line axion scale $f_a = M_{\text{scal}}/128$ ([C]). A natural compiler scale for the decay constant is the scalaron mass over the full carrier \times deck multiplicity,


$$f_a = \frac{M_{\text{scal}}}{2 \dim S^+ |\mu_4|} = \frac{M_{\text{scal}}}{128}, \quad M_{\text{scal}} = c_3^{7/2} M_{\text{P1}} \approx 3.06 \times 10^{13} \text{ GeV},$$

giving $f_a \approx 2.39 \times 10^{11}$ GeV and, via the standard QCD-axion relation $m_a \simeq 5.70 \mu\text{eV} (10^{12} \text{ GeV}/f_a)$, a mass $m_a \approx 23.8 \mu\text{eV}$. This is consistent with the closed near-hilltop misalignment $\theta_i = 170.4^\circ$, but it is a *cosmological scenario* (sensitive to anharmonic / string corrections near the hilltop), not a closed hit. [verification/v25_frontier_conjectures.py]

6 Full quantum gravity — induced from the boundary kernel

Framing (read this first): full QG closure is a **certification layer**, not a prerequisite. TFPT does *not* need the complete quantum-gravity measure to test its Standard-Model and cosmology readouts — those run through v_{geo} and F_{transfer} and are falsifiable now (JUNO, CMB-S4, EDM). G_{net} certifies the metric sector; its absence narrows what can be *claimed*, not what can be *tested*.

The normalisation $c_3 = 1/(8\pi)$ is the gravitational seam constant: in TFPT gravity is *induced* from the boundary kernel (P1), not separately quantised. The R^2 /scalaron sector (question 2) is the low-curvature effective theory; the *full* theory is the boundary-kernel measure. *Newton’s constant is then not an independent input:* fixing the physical Λ branch ($(8\pi)^2 \delta_{\text{top}} = \frac{3}{4\pi^2}$) pins \bar{M}_{P1} relative to Λ , so G_N is read out *after* selecting the physical Λ branch *and* fixing the one dimensionful induced-gravity (Sakharov) anchor ([E]/[C], [verification/v60_lambda_metrology_branch.py]) — a metrology readout, not a from-nothing prediction.



Honest status of quantum gravity — physical sector closed, ambient gap-decoupled [E]/[O]

“Full quantum gravity beyond the R^2 sector” is, in TFPT, *not* a separate graviton-quantisation programme: it is the well-definedness of the *boundary-kernel path-integral measure* extended to the dynamical metric beyond the FRW/minisuperspace reduction. Two results sharpen its status. **(i) Heat-kernel grounded (G2)**. The spectral action $\text{Tr } f(D/\Lambda)$ gives, via the standard Seeley–DeWitt coefficients for the Dirac operator ($a_2 \sim -R/3$, $a_4|_{R^2} \sim R^2/72$), Einstein–Hilbert + R^2 *structurally* — the $R + R^2$ form is not postulated; the scalaron ratio $M^2/M_{\text{Pl}}^2 = 6(4\pi)^2/f_0$ with the TFPT closure c_3^7 fixing the cutoff moment f_0 . **(ii) Gap-decoupled (G5)**. The metric coupling is gap-dominated, $2\|V_{\text{metric}}\| = 0.785 < \Delta = 6 \log \frac{3}{2} = 2.433$, so the admissible sector keeps a strictly positive effective gap $\Delta_{\text{eff}} = 1.648 > 0$ under the stated RP and metric-coupling assumptions. A positive gap protects the admissible IR sector from the un-built ambient, so under those assumptions the admissible-sector measure is closed. **What remains [O]**: the ambient metric measure — the projective limit (G6) — is the G_{metric} research gate; we do *not* claim it is physically irrelevant (that would itself be a physical interpretation), only that it is decoupled from the admissible IR sector under the stated assumptions. [verification/v36_spectral_action_g2.py]

Update (QBL chain, v108–v113): the boundary-net residual now carries *double* weight — the carrier choice $g = 5$ has been reduced to the *same* premise. The certified scalar kernel exists iff it pairs the two sheets, the certified channel counts the code identically (Pascal closure = channel identity), pair transport is minimally complete, and the single quasi-free kernel (rank = c : 5 carrier block, 8 seam hull) is the defining datum of the free $c = 8$ net this gate already posits. Closing the index-4 statement therefore also closes the carrier selection [E] (see `tfpt_1` and `tfpt_research_contracts`). [verification/v113_quasifree_kernel.py]

Update (free-bulk premise discharged, v160–v165): that “free $c=8$ net” premise is no longer free-standing. It is a fixed-point theorem (quasi-free \Rightarrow all higher cumulants $\kappa_{2n}=0$, v160); the infinite Schwinger cone is *eliminated* — by Bogoliubov second quantisation the cone gap equals the one-particle gap $(2/3)^6$ (v161), forced by μ_4 -equivariance (v162); and the residual factors into the already-open A_2 net-existence (an assembly of established net constructions, [C]) and $R_1 = \text{GATE.QGEO}$, with the irreducible core $\{\pi, v_{\text{geo}}\}$ a theorem (v165). So the metric-sector gate carries *zero new* independent open content. [verification/v165_premise_a_closure.py]

Summary

Item	Status	Honest handle / what is missing
η_B	[C]	downstream readout = 6.1×10^{-10} from closed $\Omega_b h^2$; leptogenesis is an interface (Boltzmann)
m_p/m_e	[O]	cross-sector (QCD scale / closed m_e); scheme-level, not a compiler power
Koide Q	[C]	computed $Q = 0.664$, 0.33% from $2/3$; near, not exact — not tuned
dark matter	[C]/[O]	determinant-line axion candidate fixed, $\theta_i = 170^\circ$ closed; relic scale f_a, m_a open
quantum gravity	[E]/[O]	$R + R^2$ heat-kernel grounded (G2); physical sector gap-decoupled; only ambient limit (G6) open

Marker key: [E] exact / machine-proven · [C] conditional (holds under named hypotheses) · [O] open / axiom · [X] falsifiable kill test. Per-claim fine type in `verification/status_ledger.csv`.

Net (honest). Two of the five have a genuine quantitative handle that already lands on the data ($\eta_B = 6.1 \times 10^{-10}$ via the closed Ω_b ; the axion DM candidate with a closed misalignment angle). One is a computed *near*-miss reported without tuning (Koide 0.664 vs $2/3$). One is structural-only and left open (m_p/m_e as a cross-sector ratio); quantum gravity has been sharpened — its $R + R^2$ form is now heat-kernel grounded (G2) and its physical sector is gap-decoupled (G5), leaving only the ambient projective limit (G6) as a mathematical-completeness item. None has been forced onto the compiler ladder — which is exactly the right call: the strong results stay strong precisely because the weak ones are not oversold.

Appendix: computational verification

These are the *frontier* items: most carry [O] (open) or [C] (conditional) and are therefore *not* part of the pass/fail suite. The one quantitative handle that is reproduced is $\Omega_b = (1 - \frac{1}{4\pi})\varphi_0^{\text{ret}}$ (whence the η_B order) in `verification/v7_gravity_cosmo.py`. The $R + R^2$ structure and the gap-decoupling of the physical QG sector *are* machine-checked (`v28_gravity_fR.py`, `v36_spectral_action_g2.py`); only the proton/electron ratio and the ambient projective limit (G6) are left open by design — no script “proves” those, consistent with their [O] marking. The compiler/SM identities they build on are checked by the suite in `verification/` (run `python verification/run_all.py`; see `verification/README.md`).