# TFPT — Research Contracts for the two open gates

($U_{\text{wall}}$): the parabolic flavor wall-selection and ($G_{\text{metric}}$): the full quantum-gravity measure

June 8, 2026

**What this document is about**

After the compiler closure the entire residual is

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

This note turns the two genuine research gates into *contracts*: a numbered chain of lemmas, the single theorem that closes each gate, and — for every step — whether it is machine-certifiable today. ($F_{\text{frontier}}$) (Koide, $\eta_B$, axion relic scale, $m_p/m_e$) is *not* a gate: those are deliberately typed QFT/cosmology interfaces (see `tfpt_4_frontier`). **Priority:** ($U_{\text{wall}}$) *first* (finite, algebraic, falsifiable), then ($G_{\text{metric}}$) (deep analytic programme).

## Research Contract 1 — ($U_{\text{wall}}$)

**Goal.** Prove

$$\nabla_F^\star = \text{Sel}_{D_4,\,\text{cusp},\,\det R=8,\,\text{Spec}(Q_+)=\{1,2,3\}}(\mathcal{M}_{\text{par}}(\mathbb{P}^1, \mu_4, E, \alpha)),$$

with $E = \mathcal{O}(-2) \oplus \mathcal{O}(-1) \oplus \mathcal{O}(-1)$ and $\alpha = \{0, \frac{1}{3}, \frac{2}{3}\}$ at the four points of $\mu_4$. The substrate is standard: degree-zero stable parabolic bundles correspond to unitary representations with prescribed local monodromies (Mehta–Seshadri); tame non-abelian Hodge gives the Higgs / flat-connection side. The TFPT content is to pin the *one* $D_4$-symmetric, rank-3, four-point point that realises the selectors.

**Lemma 1** (U0 — normalisation)**.** *The gate data are equivalent to the $SU(3)$ character variety $\mathcal{X}_{\text{cusp}} = \{(M_1, \dots, M_4) \in C^4_{\text{cusp}} : \; M_1M_2M_3M_4 = \mathbf{1}\}/SU(3)$ with $C_{\text{cusp}} = \text{Ad}_{SU(3)}\,\text{diag}(1, \omega, \omega^2)$, $\omega = e^{2\pi i/3}$, and $\pi_1(\mathbb{P}^1 \setminus \mu_4) = \langle \gamma_1, \dots, \gamma_4 \mid \prod \gamma_i = 1\rangle$. To show: $\mathcal{M}^{D_4}_{\text{par}} \simeq \mathcal{X}^{D_4}_{\text{cusp}}$ as polystable Mehta–Seshadri / tame Hodge realisation.* Machine: the group presentation and $D_4$-action are finite and codable (Sage/Wolfram).

**Lemma 2** (U1 — wall landscape, *corrected*)**.** *Parabolic degree zero forces the line weight-sums onto the wall $(w_1, w_2, w_3) = (2, 1, 1)$. An explicit balanced representative (weights in $\frac{1}{3}$) is*

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

*each column a permutation of $(0, 1, 2)$, row sums $(2, 1, 1)$.* ***Honest result:*** *the $6^4$=1296 column-permutation matrices give 144 walls; under the* full *group $D_4 \times S_3(\text{lines}) \times S_3(\text{values})$ they form* five *orbits,* not one. *So $W_{\text{wall}}$ is* not *singled out by symmetry — the uniqueness must come from the selector (Lemma U4), not from a symmetry quotient. The wall landscape is nonetheless a finite, explicitly enumerated set; $W_{\text{wall}}$ is one of the five. Machine: fully certified $C_U^{(1)}$ (*[`verification/v29_research_contract_certs.py`]*).*

***The splitting type is now algebraic (RH route,*** [`verification/v32_rh_splitting.py`]***).*** *Passing from the monodromy to the Fuchsian* residues *$A_k$ (eigenvalues = weights $\{0, \frac{1}{3}, \frac{2}{3}\}$, $A_k$ =*

$U^{k-1}A_0U^{-(k-1)}$, $U = \mathrm{diag}(1, i, -i)$) *the twisted* $\mathbb{Z}_4$*-average collapses* exactly: $\sum_k U^k A_0 U^{-k} = 4\,\mathrm{diag}(A_0)$, *so the exponents at infinity are* $4\,\mathrm{diag}(A_0)$. *A flat bundle with regular* $\infty$ *needs integer exponents summing to* $4{=}-\deg E$; *with the cusp weights the only options are perms of* $(2,1,1)$ *and* $(2,2,0)$, *i.e.*

$$\boxed{\mathcal{O}(-2)\oplus\mathcal{O}(-1)^2 \iff \mathrm{diag}(A_0) = (\tfrac{1}{2}, \tfrac{1}{4}, \tfrac{1}{4})}$$

*(the sibling* $(2,2,0) \to \mathcal{O}(0)\oplus\mathcal{O}(-2)^2$*). Such a Hermitian* $A_0$ *(spectrum* $\{0, \frac{1}{3}, \frac{2}{3}\}$*, that diagonal) exists by Schur–Horn. So the* splitting *is an algebraic condition on* $\mathrm{diag}(A_0)$. Honest wall: *the global* $\prod M_k = \mathbf{1}$ *is the path-ordered monodromy (*not $\exp(2\pi i \sum A_k)$*), constraining* off-diag$(A_0)$ *via the genuine RH solve; and R-extraction still needs the* $C_6 \leftrightarrow \mathbb{P}^1 \setminus \mu_4$ *bridge. The RH route fixes the* splitting *algebraically but not* $c_u/c_d$.

**Lemma 3** (U2 — the kill switch, *run and hardened* [N]). *The selected polystable wall representative must* not *collapse to the diagonal abelian* $U(1)^3$ *point if it is to give non-trivial* $SU(3)_F$ *transport mixing:* ***(A)*** $\nabla_F^\star$ *polystable but* not *simultaneously diagonal (the desired breakthrough), or* ***(B)*** *diagonal collapse, in which case* $c_u/c_d$ *stays honestly [P].* ***Result: the cheap test is run robustly and lands on A.*** *The kill-switch fires iff the* $D_4$*-symmetric cusp representation is* reducible *(commutant non-scalar). Across* 400 *random cusp-class generators of the* $\mathbb{Z}_4$ *family* $M_k = U^{k-1}MU^{-(k-1)}$ *the commutant is scalar at* every *point (irreducible* 400/400*), the reducible locus is hit with probability* 0 *(measure-zero, codimension* $\geq 1$*), and the mixing strength stays bounded away from zero (*min off-diag $= 0.35 > 0.05$*). So case A is the* generic *behaviour, not an accident of the one explicit point of [verification/v33_explicit_flat_bundle.py]:* ($U_{\mathrm{wall}}$) survives *its own cheapest falsification test. [verification/v38_uwall_killswitch.py]*

**Lemma 4** (U3 — $D_4$ fixed locus as an explicit variety, *positive-dimensional*). *With* $M_k = U^{k-1}AU^{-(k-1)}$ *(*$U = \mathrm{diag}(1, i, -i)$*, the* $\mathbb{Z}_4$ *rotation) and the reflection* $M_1 = VM_1^{-1}V^{-1}$, $M_3 = VM_3^{-1}V^{-1}$, $M_2 = VM_4^{-1}V^{-1}$ *(with* $V = -[[1,0,0],[0,0,1],[0,1,0]]$, $V^2{=}1$, $VUV^{-1}{=}U^{-1}$*), the relations become polynomials in the entries of* $A \in C_{\mathrm{cusp}}$*;* $\mathcal{X}^{D_4}_{\mathrm{cusp}}$ *is cut out as* $\bigcup_k C_k$*, parametrised by trace coordinates* $x = \mathrm{tr}(M_1M_2)$ *(note* $\mathrm{tr}\,A = 1{+}\omega{+}\omega^2 = 0$*).* ***Result (***$C_U^{(2)}$***, [verification/v30_d4_character_variety.py]):*** *adding the reflection to the* $\mathbb{Z}_4$ *locus of v19 does* not *isolate the point — the* full *$D_4$-fixed product locus is still* positive-dimensional *(*$|\,\mathrm{tr}(M_1M_2)|$ *varies continuously over* $\approx [0,3]$*). So even the full* $D_4$ *symmetry cannot select* $\nabla_F^\star$. Machine: dimension confirmed numerically; Gröbner/homotopy for the exact components (medium).

**Lemma 5** (U4 — selectors as equations). *The readouts* $\rho \mapsto \det R(\rho)$, $\rho \mapsto \mathrm{Spec}(Q_+(\rho))$, $\rho \mapsto \Lambda_{f,j}(\rho)$ *are well-defined (cyclotomic/algebraic) on* $\mathcal{X}^{D_4}_{\mathrm{cusp}}$. ***Acceptance criterion:***

$$\#\{\rho \in \mathcal{X}^{D_4}_{\mathrm{cusp}} :\ \det R = 8,\ \mathrm{Spec}(Q_+) = \{1,2,3\}\} = 1$$

*in the polystable quotient. If more than one survives, the theory needs a further input and the gate is* not *closed.* ***The bottleneck (made precise by U3):*** *since the* $D_4$*-fixed variety is positive-dimensional, the* selector *must do all the cutting — and evaluating* $\det R(\rho)$, $\mathrm{Spec}(Q_+(\rho))$ *as functions on* $\mathcal{X}^{D_4}_{\mathrm{cusp}}$ *requires the* parabolic-degree $\leftrightarrow$ residue dictionary $R(\rho)$*, i.e. the H2 equivalence. That dictionary is the single remaining analytic input of the (*$U_{\mathrm{wall}}$*) gate.*

***What*** $R(\rho)$ ***is ([verification/v31_R_dictionary.py]).*** *Two facts pin its nature.* (i) Case A is alive: *every* $D_4$*-fixed cusp monodromy found is* irreducible*, so the wall point carries non-trivial* $SU(3)_F$ *mixing — the U2 kill switch does* not *fire to the diagonal case B.* (ii) $R(\rho)$ is not algebraic: *$R$ is integer-valued, but the algebraic invariants of* $\rho$ *(e.g.* $\mathrm{tr}(M_1M_2)$*) vary* continuously *on the locus; an integer readout cannot be a continuous algebraic function. Hence* $R(\rho)$ *is the discrete invariant of the* non-abelian Hodge / Mehta–Seshadri *map (the parabolic-degree / Hodge-filtration jumps of the harmonic bundle* $E_\rho$*): integer output, but its* computation *is the transcendental harmonic-metric (Hitchin) solve* on the generic Higgs locus. ***Refinement (U4***$''$***, [verification/v40_harmonic_metric.py]):*** *at the* polystable *point TFPT actually selects this pessimism does not apply — see the box below.*

## Progress: the *selector side* of U4 is closed on the explicit point [N]

While $R(\rho)$ as a function on the whole locus needs the Hitchin solve, the two U4 *selectors* are *algebraically* accessible on the explicit flat bundle of [`verification/v33_explicit_flat_bundle.py`] — no PDE. Computed exactly from $A_0$: the exponents at infinity $\operatorname{eig}(\sum_k A_k) = \{2, 1, 1\}$ give the splitting type $E = \mathcal{O}(-2) \oplus \mathcal{O}(-1)^2$, i.e. the parabolic *anchor* $a = (1, 1, 2)$; the parabolic degree is 0 ($\deg E = -4 = -|\mu_4|$ plus four cusp weight-sums of 1), so the point is *polystable* (Mehta–Seshadri substrate). The two selectors then read off the bundle directly:

$$\boxed{\begin{aligned} \det R = 8 = n{\cdot}a \quad &\text{(anchor pairing on the splitting type)}, \\ \operatorname{Spec}(Q_+) = \{1, 2, 3\} = 3\alpha + 1 \quad &(\alpha = \text{cusp weights } \{0, \tfrac{1}{3}, \tfrac{2}{3}\}) \end{aligned}}$$

So $\det R = 8$ is the anchor pairing of the splitting type and $\operatorname{Spec}(Q_+)$ is the affine image ($d$=3 the weight denominator) of the cusp weights — *both selectors are read off the explicit bundle's algebraic data.* **Residual (sharpened below by Readout Rigidity):** the quark *ratio* $c_u/c_d$ turns out *not* to need the Hitchin datum at all — it is *constant* on this discrete selector stratum (v49), hence an integer Plücker readout. The only genuinely transcendental input left is the *absolute* amplitude normalisation ($U_{\text{point}}$), which is an anchor, not a missing PDE. [`verification/v39_uwall_selectors.py`]

## Breakthrough: the harmonic metric is *finite linear algebra*, not a PDE [N]

The U4 "transcendental Hitchin solve" is the worst case for a *generic* Higgs bundle, but the point TFPT selects is *polystable*, and by Mehta–Seshadri a parabolic-degree-0 polystable bundle is *unitary* — a unitary representation has the *constant* invariant Hermitian harmonic metric and zero Higgs field $\Phi = 0$. So at the selected point the harmonic metric is finite linear algebra. On the explicit bundle of [`verification/v33_explicit_flat_bundle.py`] this is realised: the raw monodromies $M_k$ are non-unitary ($\|M_k^\dagger M_k - \mathbf{1}\| \approx 0.83$), but there is a *unique* common invariant Hermitian form $H$ ($\dim\ker\{M_k^\dagger H M_k = H\} = 1$), it is *positive-definite*, and the unitarised holonomy $\widetilde{M}_k = H^{1/2} M_k H^{-1/2}$ is unitary ($\sim 10^{-8}$), lies in the cusp class $\{1, \omega, \omega^2\}$ with $\prod \widetilde{M}_k = \mathbf{1}$, and has *clean* matrix elements $|\widetilde{M}_k|_{ij} \in \{0, \frac{1}{2}, \frac{1}{\sqrt{2}}\}$. **So $C_U^{(3)}$ reduces from a transcendental PDE to: this finite unitarisation (done) + the combinatorial word-length $R$ (closed: H1 + $\det R = 8$).** *Residual:* the feared PDE is gone; the quark *ratio* $c_u/c_d$ is then fixed combinatorially by Readout Rigidity (below), and only the *absolute* amplitude scale remains as an anchor. [`verification/v40_harmonic_metric.py`]

## The final leg-assignment test (honest result) [N]

Does the harmonic-frame holonomy $\{0, \frac{1}{2}, \frac{1}{\sqrt{2}}\}$ feed the $\delta = \frac{1}{2}$ resolvent to give the lepton amplitudes $\Lambda = (0.475, 1.107, 0.917)$? **(i) Ruled out as a literal identity:** $\Lambda_\mu = 1.107 > 1$, while unitary holonomy entries have modulus $\leq 1$, so the amplitudes *cannot* be holonomy matrix elements — they are the non-unitary resolvent Green function (as v19 already noted). **(ii) Suggestive positive:** the harmonic-frame holonomy diagonal modulus is exactly $|\operatorname{diag} \widetilde{M}_0| = (0, \frac{1}{2}, \frac{1}{2})$, and that $\frac{1}{2}$ *equals* the distinguished lepton transport value $\delta = \frac{1}{2}$ that v20 used by hand — so the geometry plausibly *supplies* $\delta$ (an observation at the explicit point; genericity across the locus is future work). **(iii) Clean negative:** no natural holonomy construction (e.g. $(\mathbf{1} - \frac{1}{2}\widetilde{M})^{-1}$) maps the holonomy to the amplitude vector. **Verdict:** the amplitudes are the $\delta = \frac{1}{2}$ resolvent (closed combinatorially via the word-length $R$); the holonomy supplies $\delta$ and the leg selection, not the amplitude magnitude. $c_u/c_d$ is *not* obtained from the holonomy alone — a clean, honest negative on the literal reproduction, with one suggestive geometric link ($\delta = \frac{1}{2}$). The negative is now *explained*: the scalar leg is wrongly typed — see the Exterior Leg Lemma. [`verification/v41_leg_assignment.py`]

**Exterior Leg Lemma — the quark $u/d$ leg is a $\Lambda^2 F$ area, not a scalar [I]/[P]**

The v41 negative is a *category error*, not a dead end. $u$ and $d$ share the hexagon-radial data $(r, w) = (1, 1)$ ($L = 7$ for both), and the $C_6$ resolvent $c(r, w) = |\mu_4|^w / (\frac{5}{4} - \cos\frac{r\pi}{3})$ reads *only* $(r, w)$ — so any *scalar* leg gives $c_u/c_d \equiv 1$. A scalar sees radius, not orientation. The $u/d$ separation lives in the smallest *non-scalar* invariant of the $SU(3)_F$ composition: the exterior square $\Lambda^2 F$, i.e. the oriented anchor-plane area $\mathrm{Pl}_{\mathbf{1},a}(K)$. The "11" is then *generated*, not read:

$$(K, \mathbf{1}, a) \to B_K = \left(\begin{smallmatrix} 13 & 8 & 4 \\ 18 & 11 & 6 \end{smallmatrix}\right) \to \mathrm{Pl}(K) = (-1, 6, 4) = (-N_\Phi, |R^+(A_3)|, |\mu_4|) \to \|\,\mathrm{Pl}(K)\|_1 = 11,$$

$$\boxed{\frac{c_u}{c_d} = \frac{g_{\mathrm{car}}\,\|\,\mathrm{Pl}_{\mathbf{1},a}(K)\|_1}{N_{\mathrm{fam}}^2\,\Delta_Q} = \frac{5 \cdot 11}{9 \cdot 13} = \frac{55}{117}} \quad \left(\text{anchor microcode } = \frac{e_2(a)(p_3(a)+1)}{p_0(a)^2(2p_2(a)+1)}\right).$$

**Typing:** [I] for the Plücker / microcode identities ($\mathrm{Pl}(K) = (-1, 6, 4)$, $\|\cdot\|_1 = 11$, $5 \cdot 11/(9 \cdot 13) = 55/117$); and — by Readout Rigidity (v49, below) — [I] also for the *physical* ratio $\mathcal{C}_{ud}(\rho^\star)$ = this $\Lambda^2 F$ readout, which is *constant* on the derived selector stratum and so does *not* need any point-selection on $\Lambda^2 F$. The quark residual is thereby re-typed from "find the right scalar $y$" to a closed exterior readout for the *ratio*; only the *absolute* amplitude scale ($U_{\mathrm{point}}$) stays as an anchor, no new scalar, no fishing. [`verification/v42_exterior_leg.py`]

**Bridge test ([`verification/v43_exterior_bridge.py`]).** The typing is confirmed: for the $SU(3)$ holonomy the exterior square is the conjugate representation, $\Lambda^2(\widetilde{M}) = \overline{\widetilde{M}}$ exactly — the exterior leg *is* the $\bar{\mathbf{3}}$. But the *continuous* action $\Lambda^2(\widetilde{M}) \cdot (\mathbf{1} \wedge a)$ does *not* equal the *integer* $\mathrm{Pl}(K) = (-1, 6, 4)$: the integer Plücker is the combinatorial $K$ datum. So the bridge $\rho^\star \to \mathrm{Pl}(K)$ is the *discrete* non-abelian-Hodge invariant (the integer $R$; $\det R = 8$ and $\mathrm{Spec}(Q_+) = \{1, 2, 3\}$ already confirmed on the bundle, Lemma U4 box), not a continuous exterior-square computation — exactly the H2 equivalence, now pinned from the $\Lambda^2 F$ side.

**Lie-level grounding — no special math ([`verification/v44_carrier_exterior.py`]).** The discrete invariant has a home already in TFPT: the carrier half-spinor is the *even exterior algebra* of the 5-slot carrier, $16 = \Lambda^{\mathrm{even}}(5) = \binom{5}{0} + \binom{5}{2} + \binom{5}{4} = 1 + 10 + 5$, and under $SU(5) \subset SO(10)$ the up quark sits in $10 = \Lambda^2(5)$ while the down sits in $\bar{5} = \Lambda^4(5)$. So the *exterior degrees* are $\deg(u^c) = 2$, $\deg(d^c) = 4$, with $\deg(u) + \deg(d) = 6 = |R^+(A_3)|$ (the hexagon) and $\deg(d) - \deg(u) = 2 = |\mathbb{Z}_2|$ (the sheet). The "$\Lambda^2$" of the exterior leg is therefore not exotic Hodge data — it is the carrier's own $\Lambda^2$ grading (the up quark *lives* in $\Lambda^2(5)$). This re-types the residual once more: from "discrete non-abelian-Hodge invariant" to "the carrier exterior degree" — a standard [L] fact. *Honest scope:* this grounds the *type*; the Plücker *number* 55/117 stays the family $\mathrm{Pl}(K)$ [I], and the carrier-$\Lambda^2$ $\leftrightarrow$ family-$\Lambda^2 F$ link (via family $\otimes$ carrier $= 15 = 16-1$) plus the normalisation remain [P].

**The "11" is Pascal — a third, simplest origin ([`verification/v45_family_exterior.py`]).** Following that link, $\mathrm{Pl}(K)$ is the *exterior algebra of the family fundamental* $4 = \mu_4$ (the $A_3{=}SU(4)$ fundamental): $\Lambda^k(4) = \binom{4}{k}$ is Pascal row $4 = (1, 4, 6, 4, 1)$ with full sum $2^4 = 16 = \dim S^+$, and $|\,\mathrm{Pl}(K)| = (1, 4, 6) = \{\Lambda^0, \Lambda^1, \Lambda^2\}(4)$ are exactly its distinct exterior powers. Hence

$$\boxed{\|\,\mathrm{Pl}(K)\|_1 = \textstyle\sum_{k=0}^{2} \binom{4}{k} = 11 = 2^4 - \binom{4}{3} - \binom{4}{4} = 16 - g_{\mathrm{car}}}$$

(the cumulative $\Lambda^\bullet(\mu_4)$ up to degree $2 = \deg(u^c)$), and family $\otimes$ carrier $= 15 = \dim \mathfrak{su}(4) = \Lambda^2(5) \oplus \Lambda^4(5) = 10 + 5 = N_{\mathrm{fam}} g_{\mathrm{car}}$ read three ways. So the "11" is generated a *third* way — the cumulative family exterior algebra up to the up-quark's degree — pure Pascal/branching, *no special math.* The [P] physical normalisation is unchanged; the algebraic origin of 55/117 is now triply grounded.

**Lemma 6** (U5 — the "11" worry, *discharged* by Readout Rigidity)**.** *The original worry was that* $c_u/c_d = \mathcal{C}_{ud}(\rho^\star)$ *must be evaluated at the uniquely selected* $\rho^\star$ without *using* 11 *as input, so that until then* 55/117 *would be a mere table-level rational shadow.* ***This is now resolved (v49):*** $\mathcal{C}_{ud}$ *is* constant *on the discrete selector stratum* $\mathcal{S}_{8,Q}$*, so the "11" is* earned *as the rigid Plücker readout* $\|\,\mathrm{Pl}(K)\|_1$ *(v42)* — independently *of point-uniqueness. Point-uniqueness of* $\rho^\star$ *enters only the* absolute *amplitude scale, not the ratio.*

**Lemma 7** (U6 — the four certificates)**.** $C_U^{(1)}$ *wall enumeration (done);* $C_U^{(2)}$ $D_4$*-fixed character solver;* $C_U^{(3)}$ *selector uniqueness (one S-equivalence class);* $C_U^{(4)}$ *amplitude extraction — the* ratios $R, Q, c_u/c_d$ *are read off the (derived) selector stratum (closed, v49/v71); only the* absolute $U_f^\star, c_q$ *normalisation is the anchor.*

**Progress: an explicit valid flat bundle is realised (RH solve) [N]**

A numerical Riemann–Hilbert solve produced an *explicit* Hermitian residue $A_0$ (eigenvalues $\{0, \frac{1}{3}, \frac{2}{3}\}$, $\mathrm{diag} = (\frac{1}{2}, \frac{1}{4}, \frac{1}{4})$) for which the Fuchsian system $A(z) = \sum_k U^{k-1} A_0 U^{-(k-1)}/(z - i^k)$ simultaneously satisfies: the cusp class at each puncture; the splitting $\mathcal{O}(-2) \oplus \mathcal{O}(-1)^2$ (exponents $\sum_k A_k = \mathrm{diag}(2, 1, 1)$); the *path-ordered* monodromy at infinity is trivial, $\|M_\infty - \mathbf{1}\| \sim 10^{-9}$, i.e. $M_1M_2M_3M_4 = \mathbf{1}$; and the representation is *irreducible* — **case A** (non-trivial $SU(3)_F$ mixing, not the diagonal case B). So a valid ($U_{\mathrm{wall}}$) flat bundle *exists and is constructed*; the U2 kill switch lands on A. *Honest scope:* this is one realising point (existence + case A), likely one of a residual family; the unique $\nabla_F^\star$ still needs the $\det R = 8$ selector and $c_u/c_d$ still needs the $C_6 \leftrightarrow \mathbb{P}^1 \setminus \mu_4$ dictionary (H2). [`verification/v33_explicit_flat_bundle.py`]

*H2 bridge probed ([`verification/v34_h2_bridge_attempt.py`]).* The explicit per-puncture monodromies $M_k$ (cusp class, $\prod M_k = \mathbf{1}$ to $10^{-5}$) were computed; but the natural resolvent-dressed diagonal extraction ($|\mathrm{diag}\, M_k| = (0, \frac{1}{2}, \frac{1}{2})$ times the $C_6$ resolvent factor) does *not* reproduce the lepton amplitudes (0.475, 1.107, 0.917). So the precise $\Gamma_{ij}^{\min}$ geodesic-to-word dictionary (how the 6-site hypercharge hexagon combines with the 3-dim family $\rho^\star$) is the genuine remaining analytic input — it is *not* fixed by a natural guess, and we do *not* fish for 55/117.

**Theorem U — the four-way split (sharper than one monster gate) [F]/[A]**

*$D_4$ fixes the admissible chamber, not the physical point* (the $D_4$-fixed locus is positive-dimensional, U3); the discrete selectors $\det R = 8$, $\mathrm{SNF}(R) = (1, 1, 8)$, $\mathrm{Spec}(Q_+) = \{1, 2, 3\}$ do the cutting. The gate splits into four pieces, three of them essentially closed:

- $U_{\mathrm{unitary}}$ [N]: parabolic degree 0 + polystable ⇒ unitary (Mehta–Seshadri) ⇒ $\Phi = 0$; the harmonic metric is the *finite* invariant form $H$ ($M_k^\dagger H M_k = H$), not a Hitchin PDE ([`verification/v40_harmonic_metric.py`]).
- $U_{\mathrm{H2}}$ [N]: $R, Q, K$ are read off the bundle (splitting type $= a$, $\det R = 8$, $\mathrm{Spec}(Q_+)$) ([`verification/v39_uwall_selectors.py`]).
- $U_{\Lambda^2}$ [I]: **Readout Rigidity** — on the discrete stratum $\mathcal{S}_{8,Q} = \{\det R = 8, \mathrm{SNF}(R) = (1, 1, 8), \mathrm{Spec}(Q_+) = \{1, 2, 3\}\}$ the anchor-plane readout is *constant*, $\mathcal{C}_{ud}(\rho) = g_{\mathrm{car}} \|\,\mathrm{Pl}_{\mathbf{1},a}(K)\|_1/(N_{\mathrm{fam}}^2 \Delta_Q) = \frac{55}{117}$, *independent* of the continuous $D_4$ position — so $c_u/c_d$ needs no point uniqueness ([`verification/v49_readout_rigidity.py`], [`verification/v42_exterior_leg.py`]).
- $U_{\mathrm{point}}$ [A]: full uniqueness of $\rho^\star$ — needed only for the *complete* $U_f^\star$ amplitude matrix,

*not* for $c_u/c_d$.

**Headline:** the remaining flavor bridge is not a PDE gate — it is a finite exterior *readout-rigidity* theorem; only $U_{\text{point}}$ stays open.

**The simple closure ([verification/v71_simple_r_bridge.py]).** The selector stratum $\mathcal{S}_{8,Q}$ is now *fully derived*: $\det R = 8 = \operatorname{rank} E_8$ with $\mathrm{SNF}(R) = (1,1,8)$ is the lattice invariant (cf. $\det Q = 3 = N_{\text{fam}}$, v70), and $\mathrm{Spec}(Q_+) = \{1,2,3\}$ is the $D_4$-equivariant result (v69). The four flavor operators are integer lattice maps with compiler-atom determinants $(\det Q, \det K, \det R, \det L) = (3,4,8,20)$, product $1920 = |W(D_5)|$. So by Readout Rigidity the quark *ratios* ($c_u/c_d = g_{\text{car}}\cdot 11/(N_{\text{fam}}^2\cdot 13) = \frac{55}{117}$, etc.) are *pure integer Plücker readouts — no transcendental monodromy solve.* The earlier monodromy machinery was over-engineering for the ratios; the only transcendental piece, $U_{\text{point}}$, is the *absolute* amplitude normalisation — an anchor of the same nature as the one dimensionful scale, not a missing computation.

**Gate 1 complete: $U_{\text{point}} \to v_{\text{geo}}$ ([verification/v75_upoint_to_vgeo.py]).** The absolute amplitudes are not free either. The lepton amplitudes are exact ($c = (\frac{16}{7}, \frac{4}{3}, \frac{7}{6})$, v20), every cross-sector $c$-ratio is closed (Plücker), and each *sector product* is a clean $\varphi_0^{\text{ret}}$-power (Grand Mass Volume, v46: $\det M_{\text{sector}} \sim (\varphi_0^{\text{ret}})^{6,9,10}$; the lepton product is $\frac{16}{7}\cdot\frac{4}{3}\cdot\frac{7}{6} = \frac{32}{9} = 2^{g_{\text{car}}}/N_{\text{fam}}^2$). Since (pairwise ratios) + (product) ⇔ (individuals) is a bijection, the nine charged-fermion amplitudes are *fixed up to one overall scale* $v_{\text{geo}}$. **So $U_{\text{point}}$ is not an independent gate: it reduces to $v_{\text{geo}}$, the *same* one irreducible dimensionful anchor as gravity's $1/G$ (v68) — the two $[A]$ anchors collapse to one.**

## Research Contract 2 — ($G_{\text{metric}}$)

**Goal.** Construct the TFPT metric measure

$$Z[J] = \lim_{\chi\to\infty} \int_{\mathcal{G}_\chi/\text{Diff}} \exp\!\big(-S_{\text{rel},\chi}[g,A,\psi] + J\cdot O\big)\, d\mu_{\text{Cal},\chi},$$

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

the relative spectral action whose low-curvature shadow is the already-closed $R + R^2$ equation.

**Lemma 8** (G0–G1 — configuration space & BRST/Hodge Fredholm)**.** *For $s > 5/2$ a Sobolev space of $\Theta$-symmetric metrics on the doubled seam collar $M_\Sigma$ admits a gauge slice; the gauge-fixed gravitational complex $\Omega^0(TM) \xrightarrow{K} \Gamma(S^2T^*M) \xrightarrow{K^*} \Omega^0(TM)$ with $\Delta_{\text{gh}} = K^*K$ and $\Pi_{\text{phys}} = \Pi_{\text{BRST/Hodge}}\Pi_{TT}$ is elliptic Fredholm; the conformal mode is removed as BRST-exact / non-propagating in the Calderón polarisation (*not *rhetorically).* Machine: principal symbols in xAct/-Cadabra; linear-algebra certificates in Lean (medium).

**Lemma 9** (G2 — finite relative spectral action, *heat-kernel grounded* [N])**.** *$S_{\text{rel},\chi}$ is finite, differentiable, with a local Seeley–DeWitt expansion* $\operatorname{Tr} f(D^2/\chi^2) \sim f_4\chi^4 a_0 + f_2\chi^2 a_2 + f_0 a_4 + \cdots$. *For the Lichnerowicz Dirac operator $D^2 = \nabla^2 + \frac{R}{4}$ (so $E = -\frac{R}{4}$, spinor trace $\operatorname{tr}\mathbf{1} = 4$) the Gilkey coefficients give, per $(4\pi)^{-2}$,*

$$a_2 = \tfrac{1}{6}\operatorname{tr}(6E + R\mathbf{1}) = -\tfrac{R}{3} \quad (\textit{Einstein–Hilbert } \textstyle\int R), \qquad a_4\big|_{R^2} = \tfrac{1}{360}\operatorname{tr}(180E^2 + 60RE + 5R^2\mathbf{1}) = \tfrac{R^2}{72}$$

*(the $R^2$/Starobinsky term), so the spectral action* structurally *produces* $S_{\text{eff}} = \int\sqrt{g}\,\frac{\bar M_{\text{Pl}}^2}{2}(R + \frac{R^2}{6M^2}) + \cdots$ *with the scalaron mass ratio fixed by the cutoff moment, $M^2/\bar M_{\text{Pl}}^2 = 6(4\pi)^2/f_0$. The TFPT closure $M^2/\bar M_{\text{Pl}}^2 = c_3^7$ fixes $f_0 = 6(4\pi)^2/c_3^7$, i.e. $M = c_3^{7/2}\bar M_{\text{Pl}} = 3.06\times 10^{13}$ GeV — the boundary normalisation $c_3$* is *the gravitational scale — and hence the Starobinsky slow-roll forms $n_s \simeq 1 - \frac{2}{N}$, $r \simeq \frac{12}{N^2}$. The $R + R^2$* structure *is convention-independent (standard Gilkey); the precise rational $\frac{1}{72}$ and factor 6 are the standard Dirac conventions used here ([verification/v36_spectral_action_g2.py], [verification/v28_gravity_fR.py]).*

**The simpler resolution: "full QG" for physics is the gap-decoupled admissible sector [F]**

The infinite-dimensional ambient measure (G6) is the summit, but the admissible IR sector does not wait for it. The metric coupling is gap-dominated (G5): $2\|V_{\text{metric}}\| = 0.785 < \Delta = 6\log\frac{3}{2} = 2.433$, 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 dynamically *isolates* the admissible sector from the un-built ambient/deep-UV: under those assumptions the admissible (IR) measure is closed, and the ambient completion is decoupled from it (we do *not* assert it is physically irrelevant — that would itself be a physical interpretation). So the honest reading is

$$\boxed{\begin{array}{c}\text{admissible IR sector} = \text{closed under stated RP;}\\ \text{ambient metric measure (G6)} = G_{\text{metric}}\ \text{gate}\end{array}}$$

G2 supplies the local action ($R + R^2$, heat-kernel); G5 supplies the decoupling of the admissible sector; together they reduce ($G_{\text{metric}}$) to the ambient projective limit (G6). **Precise status (Alessandro's boundary):** the admissible IR sector is closed under the stated RP and gap-dominance assumptions; the ambient metric projective measure remains a *strict TOE completion target*. Its absence does *not* affect the bounded IR claim, but it *does block certification as a strict physical TOE* — so we do not call it "mere mathematical completeness". [`verification/v36_spectral_action_g2.py`]

**Lemma 10** (G3–G4 — reflection positivity, fixed and integrated)**.** *For each $\Theta$-invariant $g$ in the Calderón ball the physical boundary kernel $\mathcal{K}_g = P_{\text{adm}}\Pi_{\text{phys}}\mathcal{N}_g^{TT}\Pi_{\text{phys}}P_{\text{adm}}$ is reflection positive, $\langle\Theta F, F\rangle_{\mathcal{K}_g} \geq 0$ (G3); and $d\mu_{\text{Cal},\chi}$ is $\Theta$-invariant, supported where G3 holds uniformly, so $\int\langle\Theta F, F\rangle_{\mathcal{K}_g}\, d\mu_{\text{Cal},\chi} \geq 0$ (G4). **G4 is the decisive step:** RP after integrating over dynamical metrics, not merely at fixed background.*

**Lemma 11** (G5 — gap dominance [F])**.** $2\|V_{\text{metric}}\|_{\text{rel}} < \Delta = 6\log\frac{3}{2}$*, with* $\|V_{\text{metric}}\|_{\text{rel}} \leq 248c_3^2 = \frac{31}{8\pi^2} = 0.39262$*, hence* $2\|V\| = 0.7852 < 2.4328$ *and* $\Delta_{\text{eff}} \geq \Delta - 2\|V\| > 1.647 > 0$*. The numerical inequality is trivial (*[`verification/v29_research_contract_certs.py`]*); the work is the operator-norm bound (hard).*

**Theorem 1** (Decoupling theorem — physical low-energy claims do *not* need the full $G_{\text{metric}}$)**.** If *the admissible sector is reflection-positive (G4) and the metric coupling is gap-dominated,* $2\|V_{\text{metric}}\| < \Delta = 6\log\frac{3}{2}$ *(G5),* then *the admissible IR measure is stable under the ambient completion: a strictly positive effective gap $\Delta_{\text{eff}} = \Delta - 2\|V\| > 0$ protects the admissible sector from the (un-built) ambient/deep-UV, so every* physical low-energy *read-out (masses, mixings, $\alpha^{-1}$, the $R+R^2$ regime) is independent of the open ambient measure. Symbolically*

$$\boxed{\text{RP} \;+\; 2\|V\| < \Delta \implies \textit{admissible IR measure stable under ambient completion}\,.}$$

*Thus ($G_{\text{metric}}$) is* not *a diffuse "full quantum gravity missing" block: it is the sharply delimited* ambient projective limit *(G6), a strict-TOE certification target whose closure the falsifiable IR physics does not depend on.* ***Typing:*** *the inequality is machine-checked [**F**]; the operator-norm bound on $\|V_{\text{metric}}\|$ is the conditional input.*

**Lemma 12** (G6–G7 — projective limit & Ward identities)**.** *A projective system $\mu_{\chi_n}$ with $\pi_{m,n*}\mu_{\chi_m} = \mu_{\chi_n}$, uniform moment bounds $\sup_n\int\|g\|_{H^s}^p\, d\mu_{\chi_n} < \infty$, tightness and regulator-independence yields $\mu_{\text{QG}} = \lim\mu_\chi$ on $\mathcal{G}/\text{Diff}$ (G6, the hardest step); the limit obeys the BRST Ward identities $\langle s\mathcal{O}\rangle = 0$, whence $\nabla^\mu T_{\mu\nu} = 0$ and the $\eta$ boundary phase cancels the APS anomalies (G7).*

**Gate 2 reduction: G6 is a *boundary* projective limit, not a bulk one ([verification/v76_gmetric_reduction.py])**

The exact decoupling margin is $\Delta_{\text{eff}} = \Delta - 2\|V\| = 6\log\frac{3}{2} - \frac{31}{4\pi^2} = 1.648 > 0$, with $\|V_{\text{metric}}\|_{\text{rel}} \le 248\, c_3^2 = \dim(E_8)\, c_3^2 = \frac{31}{8\pi^2}$ $(31 = 2^{g_{\text{car}}} - 1)$. **Holographic reduction:** because the seam is a *finite causal boundary*, the ambient *bulk* metric measure is reconstructed from a finite-codimension seam (Calderón) *boundary* measure — so *G6* is a *boundary* projective limit. The conditional theorem is then

$$\boxed{\text{RP(seam-Calderón boundary kernel)} + \text{tightness} \implies G_{\text{metric}} \text{ closes}}$$

i.e. the dimension of the open problem drops from bulk QG to a finite seam-boundary measure. **Residual:** the constructive boundary projective limit (constructive-QFT grade) — reduced, not closed. So Gate 2 is now *two clean tiers*: Tier A (IR physics) closed under RP+gap; Tier B (strict TOE) reduced to a boundary measure, no longer a diffuse "full QG missing". [P]/[A]

**And that boundary measure is a *rigorously constructed* net ([verification/v77_e8_conformal_net.py]).** The level-1 central charges $c = \dim\mathfrak{g}/(1 + h^\vee)$ give $c(E_8)_1 = \frac{248}{31} = 8 = \operatorname{rank} E_8$, $c(D_5)_1 = \frac{45}{9} = 5$, $c(A_3)_1 = \frac{15}{5} = 3$, and $c(D_5){+}c(A_3) = 5{+}3 = 8 = c(E_8)$ is a *conformal embedding* $(D_5)_1 \times (A_3)_1 \subset (E_8)_1$ (coset $c = 0$). Now $(E_8)_1$ is the holomorphic $c = 8$ $E_8$-*lattice* VOA / conformal net — one of the most rigorously constructed objects in mathematical QFT (Frenkel–Kac–Segal; Kawahigashi–Longo / Carpi). **So the conditional theorem becomes concrete:** *identify the seam-Calderón boundary measure with the* $(E_8)_1$ *lattice net* (carrier = $(D_5)_1 \times (A_3)_1$ subnet), and the gap $\Delta_{\text{eff}} > 0$ supplies the clustering/tightness — so *G6* is imported into existing RCFT/conformal-net rigor rather than built from scratch. Residual = the net identification + bulk reconstruction. [I] (charges, embedding) + [P] (identification)

**Theorem G — the closure statement [A] (target)**

The relative spectral boundary kernel over the diffeomorphism-quotiented metric sector admits a regulator-independent, reflection-positive projective-limit measure $\mu_{\text{QG}}$. Its Ward identities imply covariant conservation, and its low-curvature expansion is the TFPT $R + R^2$ seam action

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

so the closed covariant equation becomes the controlled IR shadow of the full measure. *Status: open; the deepest item. G5 certified; the regime equation closed [verification/v28_gravity_fR.py].*

## Certifiability and recommended order

| Building block | Machine | Hardness |
|---|---|---|
| $W_{\text{wall}}$ enumeration ($C_U^{(1)}$) | Lean / Python / Wolfram | easy (*done*) |
| $D_4$ fixed-locus polynomial ($C_U^{(2)}$) | Sage / Mathematica | medium |
| selector uniqueness ($C_U^{(3)}$) | interval arithmetic | medium |
| SNF, determinants, spectra | Lean | easy–medium |
| principal-symbol ellipticity (G1) | xAct / Cadabra / Lean | medium |
| gap inequality (G5) | Lean | easy (*done*) |
| operator-norm bound (G5) | analysis + numerics | hard |
| projective limit (G6) | not primarily machine | very hard |
| Ward identities (G7) | BV/BRST formal + checks | hard |

**Recommended order**

$$(U_{\text{wall}}) \;\to\; \text{freeze the frontier status} \;\to\; (G_{\text{metric}})$$

($U_{\text{wall}}$) can visibly upgrade the theory — it hits the last concrete flavor ambiguity, including the "11" — and it is finite and falsifiable (the U2 kill switch can destroy it cheaply). ($G_{\text{metric}}$) is more fundamental but a *programme*, not a single proof: heavy machinery, with the projective limit (G6) the genuine summit.