

## Prime Front — extracted research

Historical source: `tfpt_research_contracts.tex` at revision `ba2f0fc`. Original labels, results and caveats are retained; extraction changes ownership only. Complete original: `provenance/tfpt-source`. The publication assembly omits presentation boxes around excerpts.

## Contents

|                                                            |           |
|------------------------------------------------------------|-----------|
| <b>Externalization Contract: PRIME.RESIDUE.EXTERNAL.01</b> | <b>79</b> |
|------------------------------------------------------------|-----------|

**Hecke from geometry on the glue census (HECKE.GEOM.01).** The  $\mu_4$ -refined  $E_8$  class thetas that underlie the glue Q-system now carry an *in-suite* Hecke package (consolidated from discovery T27/T31/T32; 25 checks,  $\sim 11$  s). (A) Kneser  $p$ -neighbours realise  $\#iso\_lines = \sigma_3(p) \# \mathbb{P}^3(\mathbb{F}_p)$  identically and at  $p = 2, 3, 5, 7$  (135/1120/19656/137600). (B) The marked neighbour-sum is the affine Hecke element  $\nu_p = a \text{Id} + b T_p$  with cusp rule  $a_p = b - \sigma_3$ , recovering  $a_3 = -4$ ,  $a_5 = -2$  ( $p = 7$  consistency  $\lambda_{\text{odd}} = 19840$ ). (C) The census form space has  $\dim_{\mathbb{Q}} V = 7 = 5+2$  as the 2-adic oldform hull of  $E_4$  and  $f_8$ , with recovery  $\pi_{\text{cusp}} = (28 - T_3)/32$  and newform via  $(1 - V_2 U_2)$ . Classical theorems named classical; the claim is the compiler identification (correspondence  $\rightarrow$  operator  $\rightarrow a_p$ ; redundancy =  $\mu_4/2$ -adic). Fence: both new systems stay weight 4 — GL(1)/weight-drop transport remains [O]; no RH claim; no marker moves. Machine-checked: [verification/v535\_hecke\_from\_geometry.py]. [E] (mechanics) / [O] (weight drop)

**Eichler trace layer at the  $E_8$  lattice (HECKE.GEOM.EICHLER.01).** Consolidated promotion of discovery T33/T36/T37/T42 (23 checks,  $\sim 30$  s; builds on HECKE.GEOM.01). Witt  $\lambda_{\text{Eis}} = \sigma_3(\# \mathbb{P}^3 - \sigma_3)$  closed in  $p$ ; Eichler  $\lambda_{\text{geom}} = \lambda_{\text{Eis}} + a_p^2$  two-sided at  $p = 3$  (live type-(i)/(ii):  $352 = 364 - 12$ ) and  $p = 5$  (T36 freeze:  $3784 = 3906 - 122$ ) with assembler  $R = \sigma_3 - 1 - c(p^2)/8$ ; Type-A/B densities  $N_A = \min(240(1 + p^3), \#iso - 1)$  ( $p = 7$  live 82560/743040; injectivity  $p = 11, 13$ ); signed  $a_p = -c(p)/8 \Rightarrow (-4, -2, +24)$  and  $b \in \{24, 124, 368\}$ ;  $O(1)$  assembler for odd  $p \leq 31$ . Classical scaffolding named classical; claim = in-suite two-sidedness + closed compiler densities. Fences: Shell( $p^2$ ) for  $p \geq 7$  out of scope; fibre refinement scoped to  $\{5, 7, 11\}$ ; weight-4  $\rightarrow$  GL(1) stays [O]; no RH; no marker moves. Machine-checked: [verification/v536\_eichler\_trace\_layer.py]. [E] (mechanics) / [O] (weight drop)

**Half-integral bridge (HECKE.GEOM.HALFINT.01).** Consolidated promotion of discovery T38/T41/T44/T45 (20 checks,  $\sim 90$  s; companion to HECKE.GEOM.01/HECKE.GEOM.EICHLER.01). Unique signed Shimura preimage  $\text{Sh}_{t=2, \psi=1}(g) = -8 f_8$  in the 70-element compiler weight-5/2 theta-monoid (witness  $g = \vartheta_2(q^2)^2 \vartheta_3(q^2) \vartheta_4 \vartheta_4(q^2)$ ; three further monoid elements lift to scales  $+8, -16, +16$ ); exact  $T(p^2)$ -eigenform with  $\lambda = a_p(f_8)$  for  $p = 3, 5, 7, 11, 13$ ;  $U_4(g) = 0$  and level  $32 = 4 \cdot 8$  outside Kohnen 1982 (plus-route structurally closed — negative result part of the claim); AFE-based Waldspurger/Baruch–Mao quotient  $R(d) = b(|d|)^2/(|d|^{3/2} L(f_8 \times \chi_d, 2))$  constant on ten fundamental  $d \equiv 1 \pmod{8}$  ( $R = 23.1873585645\dots$ , spread  $\leq 10^{-12}$ , including  $d = 89$ ), with  $d \equiv 5 \pmod{8}$  vanishing via root number  $-1$ . Classical Shimura/Kohnen/Waldspurger/Baruch–Mao named classical; claim = unique monoid preimage + in-suite chain + *measured*  $R$ . Fences: GL(2) centre  $s = 2$  (not the  $\xi$ -line); no RH; ZETA.HP.CARRIER untouched; no marker moves. Machine-checked: [verification/v537\_halfintegral\_bridge.py]. [E] (Shimura/Hecke/Kohnen fence) / [C] (measured  $R$ )

**Compiler relative-trace identity (HECKE.GEOM.RTF.01).** Promotion of discovery T53 (18 checks,  $\sim 14$  s): HECKE.GEOM.01/HECKE.GEOM.EICHLER.01/HECKE.GEOM.HALFINT.01 are three projections of one finite relative-trace identity  $[E_8/\text{glue orbital counting}] = [\text{Newform/Waldspurger spectral sum}]$ . (R1)  $\text{Tr}_V(\nu_p \circ \pi)$  bilateral independent at  $p \in \{3, 5, 7\}$  across five projectors (anchors e.g.  $p=7$ : 727680 full / 19840 on  $f_8$ ). (R2) Eichler  $\lambda_{\text{geom}} = \lambda_{\text{Eis}} + a_p^2$  is the RTF orbit split ( $\lambda_{\text{Eis}} = 336/3780/19264$ ,  $a_p^2 = 16/4/576$ ). (R3) Period side = quaternary lattice counting with  $R = 23.1873585645$  (rel  $\sim 10^{-12}$ ). Verdict ONE-FORMULA. Classical Jacquet-type RTF / Eichler / Witt / Waldspurger named classical; NEW = compiler-native bilateral independence. Fences: identity *finite*; infinite RTF named open; no RH; not ZETA.HP.CARRIER-relevant; no marker moves. Machine-checked: [verification/v538\_relative\_trace\_identity.py]. [E] (R1/R2) / [C] (R3)

**Weil structure of the compiler family (RTF.GNS.WEIL.01).** Promotion of discovery T55/T61/T62/T63 (25 checks,  $\sim 6$  s): the compiler family possesses a fully identified Weil structure up to two *explicitly isolated* obstructions (obstructions are load-bearing verified content, not residual footnotes). (A) GNS = direct integral over the Gelfand spectrum of the character subalgebra (fibre-orthogonal; twist-mix per fibre  $\leq 5.037 \times 10^{-16}$ ). (B) Trivial Sato–Tate isotype = GL(1) core  $G_0 = (1 + Y)/(1 - Y) = \zeta_p(w - 3)^2/\zeta_p(2w - 6) = \sum 2^{\omega(n)} n^{-u}$  globally. (C) Exact linear relation  $Q_{\text{fam}} = 2Q_{\zeta}(g) - 2Q_{\zeta}(g_b) + \text{Arch} + \text{Corr}$ . **Obstruction 1:** minus sign on the doubling term (family

positivity does *not* imply  $Q_\zeta \geq 0$ ). **Obstruction 2:** Corr is non-automorphic (no polar/arch absorption; no finite Euler factor;  $e^{-\sum p^{-u}}$ -type). Finite-class  $Q_{\text{fam}} \geq 0$  measured [C]; dense-class positivity open / RH-adjacent — *not* claimed. Classical Weil/Gelfand/SU(2)/Sato–Tate/ $2^\omega$ /GNS named classical. Fences: not “almost RH”; ZETA.HP.CARRIER untouched; no marker moves. Machine-checked: [verification/v539\_weil\_structure\_family.py]. [E] (A/B/C/obstructions) / [C] (finite-class positivity)

**Amplitude route and positive linear carrier (RTF.GNS.AMP.01).** Promotion of discovery T67–T72 (34 checks,  $\sim 3$  s): the route out of the square plane behind the two RTF.GNS.WEIL.01 obstructions. (A) Amplitude Dirac  $D = \begin{pmatrix} 0 & V \\ V^\top & 0 \end{pmatrix}$  with  $D^2 = \text{diag}(VV^\top, K)$ ,  $V^\top V = K$  (rel  $2.7 \times 10^{-16}$ ), exists exactly and is Hecke-equivariant ( $p = 3, 5, 7$  intertwining exact; Shimura  $b(dp^2) = b(d)(a_p - \chi_d(p)p)$  on 644 pairs). (B) Geometric polarisation  $b = N_+ - N_-$  exact;  $\Theta = N_+ + N_-$  is a *pure* Siegel–Weil Eisenstein  $\sigma_3$ -eigenform ( $p \leq 13$ , cusp component zero) with Cohen seed  $\Theta(d) = -48L(-1, \chi_d)$  exact-rational on 159 live  $d$  (Cohen 1975). (C) Every coefficient bilinear form inherits the even- $k$  deletion (Cauchy–Littlewood window lemma, five channels; numerator  $1 - p^6 X^2$  pair-independent); the deletion is exactly the square-class double counting (fam +  $2b + \delta_{k1} = [2, 2, 2, 2]$ ; Möbius square sieve exact). (D) Positive linear carrier  $\ell^2(d, \mu)$ ,  $\mu = 48|L(-1, \chi_d)||d|^{-a}$ : full weights  $\lambda_k = 1 + p^{3k} - (\chi p)^k$  (layer  $[1, 1, 1, 1]$ ) and plus balance  $Q = Q_\zeta(g_-) + Q_\zeta(g_+)$  (rel  $2 \times 10^{-15}$ ); the  $b$ /doubling kernel enters *with plus* (sign inverted vs the square plane). (E) Completed FE  $\Lambda_\Theta(s) = 8^{1-s}\Lambda_{\Theta^+}(5/2 - s)$  exact (rel  $\sim 10^{-40}$ ; Fricke closed); the guaranteed cone is FE-self-dual. **Open boundary as content:** Euler-region positivity only (edge  $L$ -values, not the central line); the residual distance to the Weil cone is the FE-covariant gap functional  $\lambda^*$  on  $n \equiv 6 \pmod{8}$  (one exact Farkas witness; no finite signed theta library removes it). Classical Cohen 1975 / Siegel–Weil / Jacobi–Fricke / Cauchy–Littlewood / Shintani / Weil 1952 / Farkas named classical. Fences: not “almost RH”; ZETA.HP.CARRIER untouched; no marker moves. Machine-checked: [verification/v540\_amplitude\_linear\_carrier.py]. [E] (A–E exact) / [C] (sampled cone facts;  $\lambda^*$  named open)

**Matching lemma and transport ledger package (RTF.GNS.LEDGER.01).** Promotion of the discovery proof package T78/T79/T80/T81/T82/T85 (33 checks,  $\sim 10$  s; every core check recomputed, not cited). (A) The matching lemma is *proved exact-integer* on  $[4, 10^6]$ :  $7S(j) < 40A(j)$  at every atom (0 violations over 939 870 clash atoms; proven no-overflow; big-integer rechecks); exact margin  $X = 22242575693/270725400000 = 0.082159$  at  $j^* = 554400$ ;  $\rho_{\text{crit}} = 1.1440 > 21/20$ ; maximal-greedy identity  $\Rightarrow$  uniform in  $(M, \lambda)$ ; structure laws (8-law, seed-tower,  $w$ -table, Cohen seeds  $\{-48, -80, -8, -8\}$ , bracket  $\Theta \leq 2|\psi| \leq 3\Theta$ ) at 0 tolerance. (B) The transport ledger closes exactly:  $Q_{\text{Weil}} = Q_{\text{cert}} + \Delta_{\text{arch}} + \Delta_2$  with  $\Delta_{\text{pole}} \equiv \Delta_{\text{conv}} \equiv 0$  *proven* (kernel collapse sympy-exact; residual  $4.4 \times 10^{-16}$  on 24 rows); the odd-prime side of the Weil functional equals the certified plus combination  $P_\zeta(g_-) + P_\zeta(g_+)$  exactly (rel  $3.9 \times 10^{-16}$ ). (C) The signed envelope is character-exact:  $-\psi = (\chi_{-4} + \frac{1}{4}\chi_8 + \frac{1}{4}\chi_{-8}) \cdot \Theta$  (all odd  $n \leq 50000$ , both builds); signed certificate 0 violations (margin  $8.86\times$ ); confinement set equality zero-credit  $\Leftrightarrow$  coherent on  $10^6$ . (D) The archimedean term is *internal*: Legendre duplication  $(2\pi)^{-s}\Gamma(s) = \frac{1}{2}\Gamma_{\mathbb{R}}(s)\Gamma_{\mathbb{R}}(s+1)$ ; kernel identity  $k_\zeta = K_{\text{fam}} - K_{\text{shift}}$  pointwise ( $7.9 \times 10^{-31}$ );  $\Delta_{\text{arch}} = A_{\text{fam}} - A_{\text{shift}}$  at  $9.9 \times 10^{-16}$ . (E) The coherent class ( $= \mathbb{Z}[i]$ -norms, set equality) is closed by the  $\lambda$ -equivariant channel: CM carrier  $g_\lambda$  exact in two independent routes;  $\mu_1 = c_1/(c_0 d^2) \in [-1, 1]$  exact;  $\lambda$ -window certificate 0.0653 vs 0.2365 unlifted ( $3.6\times$ ), exact Fraction rechecks. (F) Frontier facts [C]: coherent-chain crossing  $k^* = 14$  ( $\log_{10} N^* = 23.1$ , exact Fractions); Robin tail factor 6.16 — the constant route stays open. **Named limits as load-bearing content:** (i) *one* classically-shaped open lemma (correlated cancellation on credit-rich non-coherent tail atoms, uniform form; provably-shaped  $\neq$  formal proof); (ii) I5 in one-family form ( $Q_{\text{cert}} + \Delta_2 + A_{\text{fam}} - A_{\text{shift}} \geq 0$ ) as the single remaining TFPT-specific object — by the closed ledger *equivalent* to Weil positivity  $\Leftrightarrow$  RH (equivalence typing, no progress claim). Classical Gronwall / Robin 1983 unconditional / Cohen 1975 / Hecke 1918–1920 / Cornacchia / Dirichlet / Mertens-AP / Landau / Legendre duplication / Weil 1952 / Alaoglu–Erdős named classical. Fences: no RH-evidence language; not “almost RH”; ZETA.HP.CARRIER untouched; no marker moves. Machine-checked: [verification/v541\_matching\_lemma\_ledger.py]. [E] (A–E exact) /

[C] (frontier facts; two named limits open)

**The margin-chain identities of phase 2 (PRIME.MARGIN.IDENT.01).** Promotion of the *identity/theorem block* of the discovery parts T128–T131 (44 checks,  $\sim 0.2$  s; every check recomputed on small frame-A windows). Scope fence, declared with the row: identities and *per-instance* theorems only — no fit, no graded floor, no Lanczos value, no uniform-in-zone statement. Battery: 24 seam instances ( $n = 3 \dots 97$ ,  $\rho = 1.25 \dots 4.00$ ; every inverted or diagonalised matrix  $\leq 300$ ), 43 graded/fine grid pairs and 400 randomised profiles; every identity is a residual against a *preregistered* tolerance plus a mutation control that must fail by  $\geq 10^{-3}$  relative. (1)–(2)  $\sum_i (1 - g_i) = t_l + t_r$  exactly ( $5.5 \times 10^{-13}$ )  $\Rightarrow$  the  $\kappa$  weights are a probability vector and  $\kappa_{\text{flat}} = 1$ ; the (T) four-term identity  $\tau_{\text{dn}} = \|y\|^2 - m_{\text{prot}} + m_{\text{fill}} - V_{\text{bord}}$  at  $2.2 \times 10^{-16}$  (two independent routes for  $\tau$ ). (3)–(4)  $p_{N-1} = 2 - p_0 + 2E$  at  $2.3 \times 10^{-16} \Rightarrow$  the equivalence  $2E > p_0 \iff p_{N-1} > 2$  (0 exceptions, both branches realised); the Abel/Hölder chain  $\kappa \leq 2 - p_0 + \frac{N-2}{N} \sum_j |\Delta p_j - \overline{\Delta p}| + (p_{\text{max}} - p_{N-1})$  per instance (0 violations; the same chain *without* the curvature term is violated on 88/400 profiles). (5)–(6) Matrix-free two-scale assembly  $= J^T Q J$  ( $5.2 \times 10^{-16}$ ), graded overlap  $= J$ ; the C ea/Strang defect identity  $S_{\text{graded}} - S_{\text{uniform}} = R^T X^{-1} R$  as a *matrix* identity with PSD defect on all 43 grids ( $2.1 \times 10^{-14}$ )  $\Rightarrow$  the compression error is one-sided by construction. (7)–(9) The secular sandwich  $\varepsilon/(\|u\|^2 + \varepsilon/\mu_1) \leq \lambda_{\min}(Q) \leq \varepsilon/\|u\|^2$  on all 48 sections where  $A > 0$  and  $\varepsilon > 0$  are *checked* (width  $\leq 1.4448$ ) with Albert’s sign equivalence in *both* directions;  $S^{-1} > 0$  entrywise (measured 48/48)  $\Rightarrow$  sign-constant and simple ground vector (Perron–Frobenius), with a control where hypothesis and conclusion fail together; and  $n_{\text{run}} \leq n_{\text{cross}} + 1 \leq s_2 + 2$  with  $\sum_j |\Delta p_j - \overline{\Delta p}| = \text{TV}(P)$ , 0 exceptions on a non-vacuous battery ( $s_2$  up to 31). **Named limits as load-bearing content:** (i) per instance on *small* windows — nothing uniform in the zone index; (ii) the  $\kappa$  law is *not* claimed (T129 found it false in general), only the chain underneath it; (iii) positivity of the pole-free section at depth (the open M25d) is neither assumed nor extended and no floor value is carried; (iv) the Perron step is an *implication* with a measured hypothesis; (v) the defect identity says nothing about the defect’s *size* (the open T130 exponent). Classical Abel / H lder / Schur–Haynsworth 1968 / Albert 1969 / C ea 1964–Strang / Yserentant 1986 / Perron–Frobenius–Ostrowski 1937 named classical; Weil 1952 cited, never used as a criterion. Fences: zero-firewall AST-checked; no RH-evidence language; not “almost RH”; ZETA.HP.CARRIER untouched; no marker moves. Machine-checked: [verification/v542\_margin\_chain\_identities.py]. [E] (per-instance identities; nothing uniform)

**The lumped M-matrix pair of phase 2 (PRIME.MMATRIX.IDENT.01).** Promotion of the *identity/theorem block* of the discovery part T136 (35 checks,  $\sim 1$  s; every check recomputed on small frame-A windows). Scope fence, declared with the row: matrix identities, per-instance theorems with *checked* hypotheses, and one *negative* result — no fit, no  $D$ -exponent, no uniform-in-zone statement, no floor value. Battery: 20 seam windows ( $n = 4 \dots 59$ ,  $\rho = 1.25 \dots 4.00$ ,  $h = 25 \dots 299$ ; every factorised or diagonalised matrix  $\leq 300$ ) giving 39 matrices with positive off-diagonals, 312 shifted ladder rungs, 8 whitened two-grid rungs; every statement is a residual against a *preregistered* tolerance plus a mutation control that must fail by  $\geq 10^{-3}$  relative. (1)–(2) With  $L_\Delta = \text{diag}(\Delta \mathbf{1}) - \Delta$  the Laplacian of the positive off-diagonal part,  $A_B = A + L_\Delta$  is Stieltjes, row-sum preserving ( $1.6 \times 10^{-15}$ ) and  $A_B \succeq A$  — an *upper* bound on  $\lambda_{\min}(A)$ , a floor never; the congruence  $A = A_B^{1/2}(I - W)A_B^{1/2}$ ,  $W = A_B^{-1/2}L_\Delta A_B^{-1/2} \succeq 0$ , is a matrix identity ( $4.4 \times 10^{-15}$ ) with  $\lambda_{\max}(W) = \rho(L_\Delta A_B^{-1})$ , and the *declared* equivalence  $\lambda_{\max}(W) < 1 \iff A \succ 0$  is checked in *both* directions — the Loewner sandwich is circular by itself, which is the fence. (3)–(4) With  $x = A_B^{-1} \mathbf{1} \geq 0$  and  $A_B^{-1} \geq 0$  entrywise (Ostrowski 1937, *checked*),  $\lambda_{\min}(A_B) \geq 1/\max_r x_r$  on all 39 rows (min slack  $3.0 \times 10^{-3}$ , sharpest ratio 0.9070) — one solve, one sign check, no eigenvalue;  $(A - sI)_B = A_B - sI$  ( $2.5 \times 10^{-16}$ ) and  $(A_B - sI)^{-1} \geq A_B^{-1} \geq 0$  entrywise (Berman–Plemmons 1979)  $\Rightarrow$  the shifted ladder is *structurally empty* (0 exceptions of 312). (5)–(6) For the regular splitting  $A_B = D_0 - N_0$  (all three hypotheses checked), Varga’s  $\rho(J) = \tau/(1 + \tau)$  holds at  $4.8 \times 10^{-15}$  ( $\rho(J) = 0.7347\text{--}0.9860$ ), and the Collatz–Wielandt bracket at the anchor vector gives an *a-priori*  $\rho(J) \leq 0.9861 < 1$  with no eigenvalue anywhere, sharp on the measured gap by 1.0062–1.0531. (7) The  $k$ -scanned M18 majorant is a *valid* bound at every  $k$  (39 (rung,  $k$ ) pairs, 0 violations) and does *not* reach the *preregistered* bar  $1/2$  on any rung (best 3.080) while the measured bad mass is

$\leq 0.0120$ : a *negative* result, the localisation term alone exceeding the bar (0.542), and the bar is not moved. **Named limits as load-bearing content:** per instance on small windows, nothing uniform in the zone index, in  $D$  or in the gap; the T136  $D$ -exponents (anchor  $D^{-0.56}$ , gap  $D^{2.72}$ , slack  $D^{0.12}$ ) are fits and stay in the sandbox; (2) carries no floor and M25d stays open; (3) and (5) are implications with hypotheses checked at the instance; (4) kills a route and says nothing about the true  $\lambda_{\max}(W)$ . Classical Stieltjes / Ostrowski 1937 / Fan 1958 / Berman–Plemmons 1979 / Varga 1962 / Collatz 1942–Wielandt 1950 / Haynsworth 1968 named classical; Weil 1952 cited, never used as a criterion. Fences: zero-firewall AST-checked; not “almost RH”; ZETA.HP.CARRIER untouched; no marker moves. Machine-checked: [verification/v543\_lumped\_pair\_identities.py]. [E] (per-instance identities) + [X] (one closed route)

**The long-lag support structure of phase 2 (PRIME.LONGLAG.SUPP.01).** Promotion of the *structure block* of the discovery part T137 (24 checks,  $\sim 0.2$ s) on the same 20 windows (3–25 prime-power atoms, 20555 positive off-diagonal edges, 455 loaded stripes, 39 bracketed lumped comparisons). Scope fence: support and amplitude structure, one two-sided bracket, and one *death certificate* for a whole family of bounds — no fit, no decay exponent, and *no* structure bound that closes. (1)–(3) The kernel splits exactly as  $c = c_{\text{arch}} - c_{\text{comb}}$  with  $c_{\text{comb}} \geq 0$ , and the archimedean section alone has *all* off-diagonals  $\leq 0$  on 20/20 windows  $\Rightarrow$  every positive off-diagonal of  $A$  is comb-generated (control: the full section carries  $+2.0 \times 10^{-1}$ ); all 20555 positive edges have their Hankel index  $M - 1 - r - t$  inside the atom-hat index set (per-window fraction exactly 1.000000)  $\Rightarrow$   $\text{supp}(\Delta)$  is a union of anti-diagonal stripes at prime-power atom positions (converse reported, *not* claimed); each stripe is a perfect matching (0 endpoint collisions), so its edge vectors are exactly orthogonal and  $L_{\Delta}$  per stripe is a direct sum of  $2 \times 2$  blocks. (4)  $\Delta_{rt} \leq c_{\text{comb}}[M - 1 - r - t]$  on all 20555 edges (excess  $-1.9 \times 10^{-6}$ , sharpest ratio 0.99998532) — no eigenvalue, no norm; the Toeplitz lag in place of the Hankel index fails on 20415/20555. The tighter one-atom refinement is *honestly undecided* at  $h \leq 300$  (the atom hats never overlap) and is not promoted. (5) For the lumped comparison  $B = D_0 - N_0$  the anchor identity  $q_J = 1 - \min_r 1/(d_r x_r) < 1$  ( $8.2 \times 10^{-16}$ ) needs no measurement; the Neumann partial sum to  $K = 8$  is then an *entrywise* lower bound for  $B^{-1}$  *because the discarded terms are nonnegative*, and partial sum plus anchor-weighted tail an entrywise upper bound, on 39 rows. (6) Since Gershgorin, every row-sum norm and every Collatz–Wielandt quotient at every *positive* weight are upper bounds for  $\rho(|E|)$ ,  $E = B^{-1}L_{\Delta}$ , one rigorous *lower* bound decides the family at once:  $\rho(|E|) \geq 1.3141 \geq 1$  on 20/20 rows (cheapest four members  $\geq 1.5937$ ), so no member can *ever* certify  $\lambda_{\max}(W) < 1$  — while the *signed*  $\rho(E) = \lambda_{\max}(W) \leq 0.99995781$  on the same rows: the absolute value alone destroys the cancellation. **Named limits as load-bearing content:** per instance on small windows, no decay profile and no amplitude law (T137’s decay measurements are fits); (2) is one inclusion; (4) bounds the amplitude and *not*  $\lambda_{\max}(W)$ , and the structure bound assembled from (2)–(4) does *not* beat the margin (T137: 0 of 35 windows) and is not promoted; (5) brackets the *lumped* comparison, not the target, whose entrywise inverse positivity at depth stays open; (6) closes a route, not the question. Classical Gershgorin 1931 / the Neumann series / Haynsworth 1968 named classical; Demko–Moss–Smith 1984 cited as a model and *not* applied (the comparison is dense); Weil 1952 cited, never used as a criterion. Fences: zero-firewall AST-checked; not “almost RH”; ZETA.HP.CARRIER untouched; no marker moves. Machine-checked: [verification/v544\_long\_lag\_support.py]. [E] (per-instance structure) + [X] (one family closed)

**PRIME.HARDY.IDENT.01 — the Hardy-core identities of phase 2 (T140, T141).** (1) The form lifting  $\text{Gram} = CHC^T$  with  $C = \text{diag}(\sqrt{\Delta})M$  ( $3.0 \times 10^{-15}$ ), hence  $\text{rank}(\text{Gram}) \leq h - 1$ . (2) The exact finite-core reduction: the closed geometric covering kernel  $K_{rs} = W([r \wedge s, r \vee s])$  equals  $M^T \Delta M$  ( $1.3 \times 10^{-15}$ ) and  $\rho(W) = \lambda_{\max}(K^{1/2} H K^{1/2}) = 1 - \lambda_{\min}(A, A_B) = \lambda_{\max}(\text{Gram})$  on three independent routes. (3)  $H = \text{diag}(s) + L_N$  ( $4.7 \times 10^{-16}$ ). (4) The checkerboard split  $\text{Band}_b = D + A_{\text{even}} + A_{\text{odd}}$  entrywise exact with a three-step Weyl sequence. (5) The four Hardy identities — the  $K^+$  Rayleigh form,  $L_N = D^T B D$  with the signed crossing kernel,  $Q = B_+ \mathbf{1}$  and  $DKD^T = L_{\Delta}$  on the interior nodes. (6) The closed Muckenhoupt quantity  $A_{M0} = \max_k Q_k(\Delta \mathbf{1})_{k+1}$  by two routes. (7) The additive Hardy chain and the joint shape  $\Lambda(H, Y) \cdot \Omega(Y)$  as *valid* upper bounds on 11/11 windows — dominance only. (8) As a negative typing, the sign-versus-absolute

separation of the measured spreads ( $1.36\times$  against  $\geq 54.94\times$ ). **Named limits as load-bearing content:** per instance on small windows, nothing uniform in the zone index or in  $D$ ; every exponent is a fit; (7) promotes validity and not size — neither shape beats the target, the additive shape is dead as a *shape* by its own exact Weyl floor and the joint shape is decided at the normalisation  $\Omega$ ; (2) is an identity and not a bound, so the zone-uniform discrete Hardy inequality stays *open*. Fences: Muckenhoupt 1972 / Bradley 1978 / Opic–Kufner 1990 / Gantmacher–Krein 1950–1960 named classical; Weil 1952 cited, never used as a criterion; zero-firewall AST-checked; no marker moves. Machine-checked: [verification/v545\_hardy\_core\_identities.py]. [E] (per-instance identities) + [X] (one route typed)

**PRIME.CAPCHAIN.IDENT.01 — the capacity-chain identities of phase 2 (T142–T145).** (1) The capacity decomposition  $K^{-1} = D^T J^{-1} D + xx^T / \text{cap}$  with  $J = DKD^T$  (entrywise,  $3.8 \times 10^{-15}$ ),  $x = K^{-1}\mathbf{1}$ ,  $\text{cap} = \mathbf{1}^T K^{-1}\mathbf{1}$ , as a matrix identity ( $1.4 \times 10^{-13}$ ) whose Dirichlet half is an *orthogonal projection* under the congruence —  $\Omega = 1$  exactly ( $5.1 \times 10^{-12}$ ). (2) The exact capacity–Rayleigh identity for  $1 - \rho(W)$ : numerator/denominator assembly equals the direct quotient for every  $v$  ( $9.5 \times 10^{-14}$ ), eigenvalue reproduced under the declared  $1/\text{gap}$  conditioning, inf property verified, the constant vector carried entirely by the capacity term. (3) Cholesky nestedness:  $\text{cap}_E([a, a+j])$  as a prefix sum of squared forward-substitution entries of one triangular solve (3404 interval capacities at  $2.0 \times 10^{-14}$ , exhaustive on the median window; the reversed ordering fails at  $\geq 4.6 \times 10^{-1}$ ). (4) The whitening certificate  $\lambda_{\max}(W) \leq \kappa_{\text{up}} = 1.2245\text{--}1.4236 \leq \text{Gershgorin } 2.1213$ , direction-correct, refusing at the shrunken shift. (5) The T145 M4 split: the pointwise mass-half theorem ( $\sum_k f_k^2 \leq \psi^2$  exact),  $\sigma_{\text{tot}} = 0.2197\text{--}0.4425 < 1$  per window, licence 4 with  $|R|$  and not  $R^+$  (witness  $[[1, -\frac{1}{2}], [-\frac{1}{2}, 1]]$ ,  $3 > 2$ , built in), the ordered sign-free chain, and  $S1'$  *certified per window on both routes* (better route  $c_0 = 2.5154\text{--}4.2265$  against Maz’ya’s Dirichlet value 4), with Charikar’s cited density constant bracketed by exhaustive enumeration. **Named limits as load-bearing content:** per instance on small windows, nothing uniform in the zone index or in  $D$ ; the  $c_0$  of (5) is *measured* at the minimiser’s own layer cake and the a-priori level lemma (T145’s L1) stays open, unclaimed; Maz’ya’s strong-type half is not used as a theorem anywhere — it is the thing the cycle transcribes. Fences: Maz’ya 1985 / Muckenhoupt 1972 / Fukushima–Ōshima–Takeda 1994 / Miclo 1999 / Charikar 2000 / Goldberg 1984 / Gershgorin 1931 named classical; Weil 1952 cited, never used as a criterion; zero-firewall AST-checked; no marker moves. Machine-checked: [verification/v546\_capacity\_chain\_identities.py]. [E] (per-instance identities and certified inequalities)

**PRIME.LEVEL.LEMMA.01 — the level-lemma identities of phase 2 (T146).** (1) The union identity for nested cakes:  $\sum_{j,l} c_j c_l |S_j \cup S_l| = 2T \|\psi_t\|_1 - \|\psi_t\|^2$  with  $T = \sum_j c_j$  and  $\psi_t = \sum_j c_j \mathbf{1}_{S_j}$  — form-independent and  $O(m)$ , exact (0.0) against the  $K \times K$  double sum and against union sizes computed with no nestedness used (controls: one shuffled level set  $3.2 \times 10^{-2}$ , a 1% coefficient perturbation  $4.8 \times 10^{-3}$ ). (2) The layer-cake counting lemma at a general base:  $G(b) \leq 2b^2 \|v\|_\infty \|v\|_1 / \|v\|^2 + \varepsilon$  for *any* vector at all 6 preregistered bases; the classical dyadic 8 at  $b = 2$  exactly, the base free (gain  $4/b_*^2 = 3.7612$ ; the one-level-down cake breaks the domination on every draw). (3) The resolvent delocalisation bound:  $\psi = \lambda R \psi$  ( $9.0 \times 10^{-14}$ , an identity — no Davis–Kahan transfer anywhere)  $\Rightarrow \|\psi\|_\infty \leq \lambda_{\text{up}} \max_j \|Re_j\|$  with Rayleigh–Ritz  $\lambda_{\text{up}}$  (1.0592–1.2883) and residual-paid certified column norms, within 1.061–1.337 of the truth; the lever: the certified columns sit 2.85–10.17 below the trivial ceiling  $1/\lambda_{\min}$ , giving  $\Gamma = 1.8356\text{--}2.2797$  against the useless  $\sqrt{m}$ . (4) The composite level lemma:  $c_0^{\text{ap}} = 2b^2 \Gamma \min(1, \Gamma_1) + \varepsilon = 3.9042\text{--}4.8488$  (slack 1.986–2.418 over the measured constant, under the ceiling  $16 = 4 \times \text{Maz’ya’s Dirichlet } 4$ ), the certified window inequality  $\text{cert\_lam\_max}(R) \leq c_0^{\text{ap}} \Psi_{\text{abs}}$  on 12/12 windows, chain loss 0.0423–0.1272 of the exact gap; the built-in no-go discriminator ( $R = aa^T + \varepsilon I$ ,  $a_i = i^{-1/2}$ :  $\Gamma$  grows  $3.676 \rightarrow 6.798$ , ratio  $1.85 \geq 1.5$ ,  $c_0^{\text{ap}} = 11.10$  exceeds every window value) fails exactly where it must, the theorem half surviving on every no-go size and the Dirichlet path-Laplacian control flat (1.0063). **Named limits as load-bearing content:** per instance on small windows — a finite list of certified window inequalities with an explicit maximum; nothing uniform in the zone index or in  $D$ , and explicitly *no* statement for all  $D$ : the asymptotic delocalisation of the Green columns (T146’s remainder)

stays *open*, typed open; only the sign-free / Green constant is bounded a priori — the energy-route  $\sigma_{\text{tot}}$  stays measured; every T146 exponent is a fit, sandbox. Fences: Maz'ya 1985 (strong-type half not used as a theorem anywhere) / Rayleigh 1877–Ritz 1909 / Cauchy–Schwarz / Miclo 1999 / Davis–Kahan 1970 (cited, deliberately not used) / Perron–Frobenius (not applicable to  $E$ , not used) / Fukushima–Ōshima–Takeda 1994 / Charikar 2000 / Goldberg 1984 / Gershgorin 1931 named classical; Weil 1952 cited, never used as a criterion; zero-firewall AST-checked; no marker moves. Machine-checked: `[verification/v547_level_lemma_identities.py]`. [E] (per-instance identities and certified inequalities)

**PRIME.GREEN.SZEGO.IDENT.01 — the Green/Szegő identities of phase 2 (T147, T148).**

(1) The factorisation identity  $\Gamma = \sqrt{m} \lambda_{\text{up}} \max_j \|Re_j\| = \sqrt{Q^*} \cdot S_w$  with  $S_w = \lambda_{\text{up}} \|R\|_F$  purely spectral and  $Q^* = m \max_j (R^2)_{jj} / \text{tr}(R^2)$  purely geometric — exact at  $2.2 \times 10^{-16}$  on 12 frame-A windows ( $m = 26 \dots 285 \leq 300$ ) and 7 random positive definite forms; the certified factorisation (residual-paid column norms, Frobenius bracket) dominates the true  $\Gamma$  everywhere ( $S_w = 1.1186\text{--}1.5777$ ,  $Q^* = 2.0879\text{--}2.8634$ ; controls  $5.0 \times 10^{-3} / \geq 10^{-3}$ ). (2) The bottom-block trace identity  $\sum_j (\Pi_B)_{jj} = |B|$  exact ( $3.3 \times 10^{-16}$ ;  $\theta = 10$ ,  $|B| = 3\text{--}5$ ), hence  $Q_B \geq |B|$  a theorem; the *measured*  $Q_B/|B| = 1.5091\text{--}1.7980$  inside the preregistered band  $[1, 2.2]$  against the Szegő/Widom sharp reference  $2|B|$  (the rank- $|B|$  coordinate projector overshoots by  $\geq 3.9$ ). (3) The second-difference  $\ell^1$  ceiling:  $|a_k| \leq \min(1, \|s_k\|_\infty \|L_0^p \psi\|_1 / \mu_k^p)$  for  $p = 1, 2$  against the true coefficients everywhere, the  $m$ -free closed form  $\|a\|_1 \leq 2\sqrt{\nu} + 1$  with  $\nu = (m+1)^{3/2} \|L_0 \psi\|_1 / (2\sqrt{2})$ , and  $m \|\psi\|_\infty^2 \leq 2\|a\|_1^2$ ; Dirichlet calibration  $\|a\|_1 = 1$  exact,  $\nu \rightarrow \pi$  (controls:  $\times 10^{-3}$  fails, a one-index  $\mu$  shift breaks the identity at  $1.5 \times 10^{-1}$ ). (4) The layer-cake  $S_w$  certificate: completed  $\text{LDL}^\top$  inertia counts equal the direct spectral counts on all 96 (window, rung) pairs (Sylvester 1852; Bunch–Kaufman 1977), and the layer cake certifies  $S_w \leq 1.3288\text{--}1.8220$  per window (true  $1.1186\text{--}1.5777$ , slack  $\leq 1.216$ ; dropping the bottom layer breaks it everywhere). (5) The Liouville transform  $m \|\psi\|_\infty^2 \leq \kappa_\Lambda m \|\hat{\varphi}\|_\infty^2 \leq 2\kappa_\Lambda \|a(\hat{\varphi})\|_1^2 \leq 2\kappa_\Lambda \text{ceil}(\hat{\varphi})^2$  with  $\varphi = \Lambda^{-1/2} \psi$  and a-priori  $\kappa_\Lambda = \max \Lambda / \min \Lambda = 1.3868\text{--}2.5805$ , stepwise on every bottom-block mode; the weight-class discriminator on a model section with a genuine order-2 zero ( $f(0) = 0$ ,  $f''(0)/2 = 1$  exact), ladder  $m = 64 \dots 288$  at matched  $\kappa_\Lambda$  (mismatch 0.0013): the bounded-variation multiplier keeps  $\nu$  flat (max/min 1.003) while the rough one (TV larger by 162.6) grows by 19.9 — the roughness, not the conditioning; the first step  $\times 10^{-3}$  fails provably (slack  $\leq \kappa_\Lambda^2$ ). **Named limits as load-bearing content:** per instance on small windows — a finite list of certified window inequalities with an explicit maximum; nothing uniform in the zone index or in  $D$ , and explicitly *no* statement for all  $D$ : the one open lifting input ( $\nu_L \leq F(\text{form})$  with  $F$   $m$ -free; equivalently bounded  $\text{TV}(\log \Lambda)$  or any weaker smoothness of the multiplier) stays *open*, typed open; the symbol-level order-2 hypothesis is *refuted* (T148:  $f(0) < 0$  on every surface window) and enters nowhere as a theorem;  $\sigma_{\text{tot}}$  stays retired as a route; every T147/T148 exponent is a fit, sandbox. Fences: Kac–Murdock–Szegő 1953 / Widom 1958 / Basor–Ehrhardt 2009 / Böttcher–Silbermann (the mechanism's address, never an authority) / Sylvester 1852 / Bunch–Kaufman 1977 / Euler 1735 / Rayleigh 1877–Ritz 1909 / Cauchy–Schwarz / Davis–Kahan 1970 (cited, computed vacuous, not used) / Gershgorin 1931 / Wilkinson 1968–Higham 2002 named classical; Weil 1952 cited, never used as a criterion; zero-firewall AST-checked; no marker moves. Machine-checked: `[verification/v548_green_szego_identities.py]`. [E] (per-instance identities and certified inequalities)

**PRIME.GAUGE.PARITY.IDENT.01 — the gauge/parity identities of phase 2 (T149, T150).**

(1) The gauge freedom: for *any* positive diagonal  $\tilde{\Lambda}$  the whitening  $\tilde{E} = \tilde{\Lambda}^{-1/2} A \tilde{\Lambda}^{-1/2}$  is an exact congruence, so  $\lambda_{\min}(A) \geq \text{cert } \lambda_{\min}(\tilde{E}) \cdot \min_j \tilde{\Lambda}_j$  is a valid certified lower bound for every gauge (an identity plus one Rayleigh step; family maximum valid) — verified on all 48 gauge assemblies over 12 frame-A windows ( $m = 26 \dots 285 \leq 300$ );  $\Gamma = \sqrt{Q^*} S_w$  holds *gauge-invariantly* at  $2.4 \times 10^{-16}$ ; the const gauge has  $\text{TV}(\log \tilde{\Lambda}) = 0$  *exactly* with the certified entrywise sandwich  $\sigma = 1.3868\text{--}2.5805$  (controls: a 1%  $Q^*$  perturbation breaks at  $4.99 \times 10^{-3}$ ; the completed-Cholesky floor refuses on the shifted indefinite form, 12/12). (2) The two-regime discriminator: the flutter smoother (moving geometric mean, half-width  $m/16$ ) removes TV by  $\geq 2.88\times$  while moving the Liouville-transported  $\nu_{\tilde{L}}$  by  $\leq 0.0088$  — even the matched-amplitude rough re-gauge moves it  $\leq 0.0046$  (bar 0.01) —

while the *untransported* functional jumps 0.41–9.76 under the same rough gauge: the roughness is in the form, not the multiplier (T149’s withdrawal of the TV hypothesis, wired per instance). (3) The explicit zone diagonal  $\Lambda_r = c_0 - c_{M-1-2r} + \sum_{s \neq r} \max(c_{|r-s|} - c_{M-1-r-s}, 0)$  at residual  $5.9 \times 10^{-16}$ ; the exact split  $c = c^{\text{arch}} + c^{\text{atom}}$  with  $c^{\text{atom}} \leq 0$  entrywise and the closed atom budget  $\|c^{\text{atom}}\|_1 \leq \sum_j \mu_j \leq 4B\sqrt{N}$ ,  $B = 1.038821$  the verified maximum of  $\psi(n)/n$  over the whole atom table (Chebyshev 1852, computed, never assumed; Abel summation); the archimedean section alone is *indefinite* on 12/12 (up to 5 negative eigenvalues, LDL<sup>T</sup>) while the full section has zero — the atoms are co-responsible for the positivity, the additive perturbation reading structurally empty (control: dropping the positive part breaks at  $4.25 \times 10^{-1}$ ). (4) The parity compression  $U_-^T T_M(c) U_- = T - H$  and  $U_+^T T_M(c) U_+ = T + H$  exact ( $3.9 \times 10^{-16}$ , cross block  $1.4 \times 10^{-16}$ ) on the 6 windows with  $M \leq 300$ : the reflection-odd section *is* the compression of the full symmetric Toeplitz section onto its antisymmetric parity sector, and the negative inertia of the sign-changing symbol ( $f(0) = -29.2 \dots - 6.3 < 0$ , re-read from T148, never re-derived) sits entirely in the *even* sector, the odd sector counting zero (control: the sign-flipped basis misses by  $3.99 \times 10^{-1}$ ). (5) The parity-sine calibration: the parity sines are exact eigenpairs of the parity Laplacian (corner 3 forced by antisymmetry; eigen-residual  $2.3 \times 10^{-13}$ ),  $\nu_k^P = \pi k^2$  on the exact model at  $m = 2001$  to  $3.0 \times 10^{-5}$ , and the per-window ladder  $\nu_k^P \leq Ck^2$  holds with certified  $C = 11.70$ – $17.28$  and the derived per-mode ceiling  $m\|\psi_k\|_\infty^2 \leq 2(2\sqrt{Ck} + 1)^2$  (controls:  $\times 10^{-3}$  fails, a one-index shift misses by 0.75). **Named limits as load-bearing content:** per instance on small windows — a finite list of certified window inequalities with an explicit maximum; nothing uniform in the zone index or in  $D$ , and explicitly *no* statement for all  $D$ : the one open input after T150 (an  $m$ -free constant  $C$  in the odd-parity-sector ladder  $\nu_k \leq Ck^2$  for a sign-changing-symbol Toeplitz section) stays *open*, typed open; TV(log  $\Lambda$ ) stays withdrawn as a hypothesis (T149);  $\sigma_{\text{tot}}$  stays retired as a route (T148); every T149/T150 exponent is a fit, sandbox. Fences: Kac–Murdock–Szegő 1953 (parity sector) / Widom 1958 / Basor–Ehrhardt 2009 / Böttcher–Silbermann (the parity sectors as the algebra’s natural objects; the sign-changing symbol is exactly what Basor–Ehrhardt does not cover, and that gap is the open input) / Sylvester 1852 / Bunch–Kaufman 1977 / Chebyshev 1852 (verified on the table) / Abel / Weyl 1912 / Bauer–Fike 1960 / Davis–Kahan 1970 (cited, computed dead, not used) / Rayleigh 1877–Ritz 1909 / Cauchy–Schwarz / Gershgorin 1931 / Wilkinson 1968–Higham 2002 named classical; Weil 1952 cited, never used as a criterion; zero-firewall AST-checked; no marker moves. Machine-checked: [verification/v549\_gauge\_parity\_identities.py]. [E] (per-instance identities and certified inequalities)

**PRIME.ODD.SECTOR.IDENT.01 — the odd-sector identities of phase 2 (T151).** (1) The grid step-over: the parity sines live on the *shifted* grid  $\theta_k = 2\pi k/N$ ,  $N = 2m+1$ , which never contains  $\theta = 0$ , with  $\mu_k^P = 2 - 2\cos\theta_k$  *exactly* ( $2.2 \times 10^{-16}$ ; eigen-residual  $2.3 \times 10^{-13}$ ); the symbol has  $f(0) = -47.6 \dots - 6.3 < 0$  (T148’s honest negative, re-read, never re-derived) with a nonempty negative window  $[0, \theta_c)$ , and  $\theta_c/\theta_1 = 0.371$ – $0.485 \leq 0.9$  on *every* window of the battery (12 frame-A windows,  $m = 26 \dots 285 \leq 300$ ): the odd grid *steps over* the negative window — odd-sector positivity is a grid fact. The honest negative beside it, promoted as content (T150’s rest item 3 answered negatively): the symbol-only pencil floor  $\kappa_{\text{sym}}(1 - \rho)$  is vacuous ( $\rho \geq 1$ ) on 12/12 and the symbol moves by 0.45–3.27 relative across *one* mode spacing — the local model is a *matrix* statement, not a symbol statement (controls: the Dirichlet corner breaks the eigenpairs at 1.0, a one-index shift misses by  $6.5 \times 10^{-3}$ ). (2) The certified bottom ladder:  $\lambda_k(A) \leq S\mu_k^P$  for  $k \leq 8$  by completed LDL<sup>T</sup> inertia counts (Sylvester 1852; a certificate, never a sorted list),  $S = 1.1019$ – $6.4725$ , the count equal to the direct spectral count on all 96 (window,  $k$ ) pairs — order 2 in  $k$ ,  $f(0) < 0$  absorbed by the Rayleigh floor  $L_P \geq \mu_1^P I$  (control:  $S \times 10^{-3}$  certifies no eigenvalue below it). (3) The discrete Sobolev step:  $\|\nabla\psi\|_2^2 = \langle \psi, L_0\psi \rangle - \psi_{\text{last}}^2$  *exactly* ( $7.6 \times 10^{-15}$ ),  $\|\nabla\psi\|_2^2 \leq \langle \psi, L_P\psi \rangle$ , and  $\|\psi\|_\infty^2 \leq 2\|\psi\|_2\|\nabla\psi\|_2$  — licensed by the odd sector’s virtual node  $\psi_{-1} = 0$ ; calibration: the model pencil is (1, 1) (certified  $1.1 \times 10^{-12}$  at  $m = 241$ ) and the Sobolev per-mode price is  $\pi$  to  $5.1 \times 10^{-8}$  at  $m = 2001$  — the honest,  $m$ -free cost of leaving the  $\ell^2$  world (controls: the boundary term  $4.8 \times 10^{-2}$ , the Neumann undercount 0.50). (4) The closed archimedean minimum (T150’s rest item 2 closed *and attained*): given the two per-window-checked shape properties ( $c_i^{\text{arch}} < 0$ , non-

decreasing on  $i \geq 1$ ),  $\min_r \Lambda_r^{\text{arch}} = c_0^{\text{arch}} - c_{M-1}^{\text{arch}}$  exactly (residual 0.0; value 2.1719–3.4674), attained at  $r = 0$  (control: the full atom-including kernel breaks the shape and misses by  $\geq 8.7 \times 10^{-2}$ ). (5) The linear per-mode bound: with the certified pencil floor  $\kappa = 0.1939\text{--}0.3599$  (one completed Cholesky per window, floor carried as  $\text{fl}/\mu_1^P$ , direction re-read on the actual bottom eigenvalue) and  $K_{\text{bot}} = S$ ,  $b_k = 2m\sqrt{K_{\text{bot}}\mu_k^P/\kappa} \leq C_S k$  with  $C_S = 12.081\text{--}26.953 \leq 2\pi\sqrt{K_{\text{bot}}/\kappa}$  on the bottom 8 modes of every window — *linear* in  $k$  where the quadratic  $\nu$  ladder grew, every  $m$ -power cancelled, nothing additive anywhere; the bound dominates the truth and the parity Dirichlet energy on every read mode; the bottom pencil ratio  $R = K_{\text{bot}}/\kappa = 3.7362\text{--}18.6607$  is the one scalar the chain still owes, typed *open* in  $m$ ; the discriminator: on the T145 no-go form the ratio explodes ( $1.686 \times 10^3 \rightarrow 2.667 \times 10^4$ , factor 15.8 over  $m = 64 \dots 256$ ) while the exact parity model holds  $R = 1$  to  $1.1 \times 10^{-12}$  (controls: the ladder  $\times 10^{-3}$  fails, the inflated floor  $1.5\kappa$  makes the Cholesky refuse on 12/12). **Named limits as load-bearing content:** per instance on small windows — a finite list of certified window inequalities with an explicit maximum; nothing uniform in the zone index or in  $D$ , and explicitly *no* statement for all  $D$ : the one open input after T151 (the  $m$ -freeness of the bottom pencil ratio  $R = K_{\text{bot}}/\kappa$  for the odd parity sector of a sign-changing-symbol Toeplitz section) stays *open*, typed *open* — its flat trend ( $x^{0.037 \pm 0.015}$ ) is a fit, sandbox; the symbol-only route stays certified vacuous;  $\text{TV}(\log \Lambda)$  stays withdrawn as a hypothesis (T149);  $\sigma_{\text{tot}}$  stays retired as a route (T148); every T151 exponent is a fit, sandbox. Fences: Kac–Murdock–Szegő 1953 / Widom 1958 / Basor–Ehrhardt 2009 / Böttcher–Silbermann (the address of the symbol question — the computed answer: the pointwise symbol language does not extend to the bottom mode) / Weyl 1912 / Parlett / Hardy–Littlewood–Pólya 1934–Agmon 1965 / Sylvester 1852–Bunch–Kaufman 1977 / Rayleigh 1877–Ritz 1909 / Cauchy–Schwarz / Wilkinson 1968–Higham 2002 / Chebyshev 1852 named classical; Weil 1952 cited, never used as a criterion; zero-firewall AST-checked; no marker moves. Machine-checked: [verification/v550\_odd\_sector\_identities.py]. [E] (per-instance identities and certified inequalities)

**PRIME.RITZ.CEIL.01 — the fixed-size Ritz ceiling certificate of phase 2 (T154).** (1) The direction lemma: for orthonormal  $Q$  and  $W = Q^T A Q$ ,  $\lambda_k(A) \leq \theta_k$  for *every*  $k \leq \dim$  — Ritz values are *upper* bounds for the eigenvalues of the *same index* (Courant–Fischer 1920 / Cauchy 1829), verified against the exact spectrum on the certificate subspace of every window and on 24 random subspaces (one-sided excess 0.0) — so the fixed-size ceiling needs *no* residual argument; Temple 1928 / Kato 1949 is a floor device and is used nowhere in the module (control: the reversed reading is recognised as false on 12/12 windows,  $\theta_k$  overshooting  $\lambda_k$  by  $\geq 18.6$  somewhere on every window). (2) The sixteen-column certificate:  $\text{span}\{t_1..t_8\} + A^{-1}L_P \text{span}\{t_1..t_8\}$  has fixed dimension 16, independent of  $m$  (one Green step by Cholesky back-solves, no inverse formed), and its ceiling  $K^F = \max_k \theta_k/\mu_k^P$  (8 completed Choleskys of matrices  $\leq 8 \times 8$ ) equals 1.1019–1.5845 with  $K^F \geq \max_k \lambda_k/\mu_k^P$  on every window and  $|K^F/K_{\text{bot}} - 1| \leq 9.9 \times 10^{-7}$  against the declared cap  $10^{-4}$  on 12/12 ( $K_{\text{bot}}$  the inertia-certified true ladder constant, count = direct count on all 96 pairs): the fixed-size ceiling *attains* the truth, retiring the size- $m$   $\text{LDL}^T$  counts from the ceiling step (controls: the parity sines alone overshoot 41.6–739; corrupting the Green columns breaks the attainment on 12/12; Kaniel 1966 / Paige 1971 quoted as the mechanism, never used as a bound). (3) The one-direction anatomy (angles measured, Björck–Golub 1973): the median of the seven aligned angles between the bottom eight of  $A$  and the bottom eight of  $L_P$  is  $0.03\text{--}0.14^\circ$  while the largest is  $83.79\text{--}89.70^\circ$  on every window of the battery (the 12 deepest in-cap frame- $A$  windows,  $n = 29 \dots 139$ ,  $m = 96 \dots 291 \leq 300$  — declared, not tuned) — the misalignment is carried by *one* nearly orthogonal direction, and that direction *is* the collapse price  $\lambda_1(A)/(t\mu_1^P) = 4.408\text{--}7.042$ , recovered in full per window by one completed Cholesky of  $A - \gamma I$  (size  $m$ , certified, not fixed size, labelled as such; control:  $A = L_P$  returns  $K_{\text{bot}} = K^F = 1$ , zero misalignment, price exactly  $1/t$ ). (4) The no-go discriminator: on the T145 reconstruction  $c(l) = 1/(1+l)$  the odd section is positive definite and the Schur floor survives ( $t = 0.05$  on all 4 sizes) but  $K_{\text{bot}}$  explodes ( $\times 35.4$ ) and the sixteen-column ceiling explodes *with* it ( $\times 57.5$ ) over  $m = 48 \dots 288$  — no manufactured closure; the Courant–Fischer direction is violated nowhere (16/16); the Dirichlet control (the Hankel half removed, arithmetic kept bit for bit) stops being positive definite ( $\lambda_1 = -5.60 \dots -3.55$  certified

on  $3/3$ ) — the reflection half is load-bearing for positivity. **Named limits as load-bearing content:** per instance on small windows — a finite list of certified window inequalities with an explicit maximum; nothing uniform in the zone index or in  $D$ , and explicitly *no* statement for all  $D$ ; the attainment is exhibited per window and *not* promoted to a theorem (T154’s no-go stress shows the certificate overshoots there by 1.12–2.19 — a property of the real section, not of the instrument); the two open inputs after T154 ( $R1'$ , an  $m$ -free floor for  $\lambda_{\min}(B_{HH})$  on the bulk parity block;  $R2'$ , an  $m$ -free lower bound for the  $L_P$  floor on the complement of the eight bottom Ritz directions — arithmetic-free, worth 91–100% of the collapse price) stay *open*, typed open; every floor consumed is certified per window at size  $m$ ; PRIME.ODD.SECTOR.IDENT.01’s ladder stays as stated and nothing is re-opened. Fences: Courant–Fischer 1920 / Cauchy 1829 / Kac–Murdoch–Szegő 1953 / Schur 1917–Haynsworth 1968 / Sylvester 1852–Bunch–Kaufman 1977 / Weyl 1912 / Temple 1928–Kato 1949 (a floor device, not used here) / Kaniel 1966–Paige 1971 (quoted, never a bound) / Björck–Golub 1973 / Wilkinson 1968–Higham 2002 / Chebyshev 1852 named classical; Weil 1952 cited, never used as a criterion; zero-firewall AST-checked; no marker moves. Machine-checked: `[verification/v551_ritz_ceiling_certificate.py]`. [E] (per-instance theorems and certified inequalities)

**PRIME.WEIL.OPERATOR.01 — the localized-Weil-operator identification (research contract, [O]).** The external operator shell for the  $C = 1$  line (v618/v619), preregistered 2026-08-02 after the third external review. *Named reference frame (verified citations):* Suzuki, *Weil’s quadratic form via the screw function* (arXiv:2606.09096, June 2026) realizes the localized Weil quadratic form on every interval  $[-a, a]$  as a canonical self-adjoint operator  $A_a$  via the screw function, proves the continuity of the lowest eigenvalue  $\lambda_a$  in  $a$  without assuming RH, and presents Yoshida’s criterion:  $Q_W(v) > 0$  for all nonzero odd  $v \in C_c^\infty(\mathbb{R})$  already implies RH; a first numerical realization is arXiv:2607.24830. *The lemma chain:* **(W1)** exact operator identification — the frame-A window bilinear form  $\hat{A} = B - S$  is the Galerkin matrix of the odd restriction of  $A_a$  on the declared piecewise bases, exactly or up to an explicit rank-one transformation and a known normalization (no asymptotic “looks similar”); **(W2)** form density — the nested window spaces are form-dense in the odd form domain of  $A_a$  with Mosco/norm-resolvent convergence of the Galerkin forms; **(W3)** uniform discrete positivity —  $\hat{A}_{a,h} \succeq 0$  for every  $a$  and every refinement  $h$  (the RH-hard step; the  $C = 1$  contraction reading of v618 — the normalized arithmetic residue as an operator of norm  $\leq 1/h$  — is the named candidate mechanism); **(W4)** passage to the continuum — W2 + W3 give  $Q_W^a(v) \geq 0$  on all odd test functions for all  $a$ , hence RH via the localized criterion. *Kill tests:* (K1) the W1 identification fails beyond a window-dependent fit; (K2) a complete, numerically stable fresh window stably exceeds the frozen  $C = 1$  bound (the v619 complete-comb surface is the definitional domain); (K3) the odd-sector restriction loses the explicit-formula bookkeeping. *Fence:* a research contract, not a claim — no RH statement, no positivity claim beyond the measured surface; the zero-side representation stays diagnostic (a proof must run prime-side or through the screw function, avoiding the critical-line circularity). Status: [O] (preregistered; the audit and the citation verification are machine-checked in `[verification/v624_external_lattice_audit.py]`). *First executed slice (2026-08-02):* the atomic half of W1 is closed — the TFPT atom table is the prime measure of Suzuki’s screw function, literally (positions log prime powers, weights  $\Lambda(n)/\sqrt{n}$ , exact; tent projection verbatim), and the first numeric Galerkin comparison measures a *non-scalar* conversion profile (−0.064..−0.070) — the preregistered rank-one/normalization freedom, now with data (`[verification/v630_suzuki_contact.py]`). On the geometric side the Hodge-chamber question is *measured*: all 67 complete windows land in the positive cone of the cover polarization lattice, on one sheet (`[verification/v627_hodge_chamber.py]`);  $P$ -canonicity and the chamber theorem stay open. *Second slice (2026-08-02, same day):* the non-scalar residual is *resolved* — it IS the zeta pole term. The Lerch collapse  $[e^{-t/2}\Phi(e^{-2t}, 2, \frac{1}{4})]'' = 4e^{-t/2}/(1 - e^{-2t})$  (exact) gives the structure theorem  $g''(t) = -2 \cosh(t/2) - 4e^{-t/2}/(1 - e^{-2t})$  off the atoms (in the  $\tilde{g}$ -normalization of this chain; see the fourth-slice erratum below — Suzuki’s own smooth layer is  $+\rho$ ): the chain-kernel smooth layer is  $-4 \times$  the TFPT arch density *minus* the pole block  $2 \cosh(t/2)$  (the  $s = 0, 1$  pole weights — TFPT’s separate rank-one, v591). The closed per-lag ratio law  $-D[4 + (e^t + 1)(1 - e^{-2t})]$  is verified,

and after pole subtraction the smooth conversion is the *scalar*  $-4D$  (monotone to 1.0006); the atom-layer lattice constant is  $D^2$  exactly. **W1 status: closed at the measured level on both halves** —  $\hat{A}_{\text{arch}} = -(1/4D)$  Galerkin<sub>smooth</sub> + pole rank-one, with the typed remainder (second-order discretization term,  $d \leq 2$  boundary cells,  $L_0^2$  projection) named for the theorem-level statement ([verification/v631\_w1\_dictionary.py]). *The implication map stays explicit:* W1 (now measured-closed)  $\rightarrow$  W2 (form density, OPEN)  $\rightarrow$  W3 (uniform discrete positivity, OPEN — the RH-hard step)  $\rightarrow$  W4 (continuum passage, classical given W2+W3). Closing W1 does not move W3; no step is silently upgraded. *Third slice (2026-08-02, same day):* the typed W1 remainder *closes*, on three fronts. **(i) Boundary closure, exact:** the duplication identity  $\psi(\frac{1}{4}) = -\gamma - 3 \log 2 - \pi/2$  and the scheme conversion  $k = \pi/4 + \frac{3}{2} \log 2$  are sympy-exact, so the two  $\delta_0$  conventions are one object in one scheme, and the exact boundary equation  $A_{\text{arch}} = \frac{1}{4}g''_{\text{smooth}} - \frac{5}{4}(\log \pi - \psi(\frac{1}{4}))\delta_0$  verifies at relative residuals  $< 10^{-22}$ ; the  $d \leq 2$  cells become three computable constants and the second-order term is the derived, window-independent moment law  $1 + 1/(6d^2)$  ([verification/v640\_w1\_boundary.py]; the exact dictionary identities additionally machine-proved in Lean, `WeilDictionary.lean`). **(ii) Portability:** the frozen dictionary (closed formulas in  $D$ , no knob) executes unchanged on three fresh windows  $h = 285/540/997$  — atom constant  $D^2(\text{window})$  at machine precision — and the preregistered kill criterion is met nowhere ([verification/v641\_w1\_portability.py]). **(iii) Matrix level:** the full quadratic form closes at the operator level (16-mode block norm distance  $2.7 \times 10^{-3}$  fully derived,  $6.3 \times 10^{-6}$  after measure-level atom re-binning; the separated pole block numerically rank 1); the literal per-vector 1% bar fails only on two *typed* lattice residuals (near-cancellation denominators; tent-vs-B-spline atom profiles), certified as typed-residual checks ([verification/v642\_w1\_matrix.py]). *Fourth slice (2026-08-02, second round of the day): erratum + theorem + W2 start + W3 tool diagnosis.* **Convention erratum (corrected the same day):** the second and third slices read Suzuki’s eq. (1.3) with Lerch coefficient  $-1$ ; the paper’s own §2.2 data lock  $+1/4$  ([verification/v643\_w1\_theorem.py], C0.1). All slice-two/-three identities are correct identities of the chain kernel  $\tilde{g} = g - \frac{5}{4}$  Lerch, and every measured number transfers *verbatim* via the exact identity  $c_{\text{gal}}^{\text{sm}}(\tilde{g}) = -4 c_{\text{gal}}^{\text{sm}}(g)$  — only the labels change: Suzuki’s own smooth layer is  $+\rho$ , the two-layer dictionary collapses to the *single* scalar  $+1/D$  on both layers (sign-compatible with positivity), and the origin constant vanishes,  $\kappa = 0$  exactly (the  $-\frac{5}{4}L\delta_0$  term was an artifact of the misreading; the v563 near-cell scheme *is* Suzuki’s origin bookkeeping). **W1 is THEOREM-closed at the measure level** ([verification/v643\_w1\_theorem.py]): the projection lemma closes the last remainder — Suzuki’s  $L_0^2$  mean-zero condition is *automatic* on the  $u$ -side ( $D : H_0^1 \rightarrow L_0^2$  a bijection), the window parity sector is a subspace of the form domain,  $A_{\text{arch}} = -g''_{\text{smooth}}$  exactly at every lag ( $3.4 \times 10^{-52}$ ), and the full form equality holds at  $1.28 \times 10^{-10}$  on the common odd sector. **W2 is started** ([verification/v644\_w2\_form\_density.py]): the classical FEM-density slice is verified at rate (with  $Q_W$   $H^1$ -bounded and  $H_0^1$  a form core the density chain is classical, cited); Rayleigh–Ritz on the nested window spaces is monotone from above on all six sequences, and the  $\lambda_a$  scan is consistent with Suzuki Thm 1.3 ( $\lambda_a = 0^+$  within  $\sim 10^{-9}$ , no sign statement); the Mosco/norm-resolvent remainder is *named*, not claimed. **W3 stays open, now with a tool diagnosis** ([verification/v648\_sign\_uncertainty.py]): the sign-uncertainty / LP-ceiling toolbox has a real, 25-digit-anchored dictionary to the critical strip (the  $d = 1$  strip multiplier’s phase slope is exactly  $L = \log \pi - \psi(\frac{1}{4})$ , the W1  $\delta_0$  weight), but its quantitative mechanism is powered by the strip half-width and dies at  $d = 1$  ( $e^{-E^*} = 0.967$  where  $< \frac{1}{2}$  is needed;  $d \geq 21$  would be required; the window family’s negative structure sits  $5.2\times$  outside the maximal BCK ball): a BCK-type mass obstruction cannot certify  $\hat{A}_{a,h} \geq 0$  — any W3 route through this toolbox must replace the dimension lever by a different large parameter; empirically the W3 surface is positive on all 67 complete windows (min  $\lambda_{\min} = +8.26 \times 10^{-4}$ , float-floor certified). **W4 unchanged;** no RH statement. *Fifth slice (2026-08-02, the closing bundle): the consolidated W2/W3 state.* **W2:** the envelope inequality holds *measured*  $a$ -uniformly — the total window symbol grows sub-log with family-uniform constants  $s_{\text{tot}} \geq c_0 w^*(t) - C_0$ ,  $(c_0, C_0) \approx (0.021, 0.055)$ ; the earlier “flat” read was a  $t$ -range artifact ([verification/v663\_garding\_envelope.py]); the Mosco/norm-resolvent preparation stands — the resolvent family is Cauchy with uniformly bounded  $H_{\log}$  norms, and the eigen-scale col-

lapse is typed: W2 routes must run through resolvent observables ([verification/v655\_w2\_mosco.py]); the two named lattice remedies for the Gårding drift are *refuted* — the drift lives in the total symbol, not in the lattice edge ([verification/v661\_garding.py], [verification/v662\_garding\_edgeband.py]) — and the drift itself is *intrinsic* to the positivity problem, not Suzuki-frame-made: the external Baez–Duarte control frame shares it ([verification/v667\_baez\_duarte.py]). The remaining W2 steps are a Fejér spectral-density bound replacing the per-mode symbol read, plus the symbol-level proof. **W3**: reduced on the measured surface to *lock sign plus the uniform bound*  $q \tan^2 \theta \leq 1 - \delta$  via the exact 2D formula  $r_{\text{id}} = (1 - q \tan^2 \theta) / (1 - \tan^2 \theta)$  ([verification/v657\_rid\_alignment.py]; the margin-bridge correlation typed as a lock identification, [verification/v656\_margin\_link.py]); the resonance landscape is mapped *without positivity loss* —  $\lambda_{\min} > 0$  on all 635 surveyed points, every  $P > 1$  excursion a rotation artifact ([verification/v659\_w3\_landscape.py]; the risk survey [verification/v658\_w3\_uniform\_bound.py]); the rotation predicate is typed honestly (necessary, not sharp; successor: block deflation; [verification/v660\_theta\_predicate.py]); and the *ground-truth recalibration* lands ([verification/v668\_ground\_truth.py]): on a proven RH analogue true positivity has *no uniform margin* —  $\delta(K) \rightarrow 0$  exactly beyond the support-resolution depth — so shrinking margins on deeper windows are the *expected* behavior of a true positivity, and the alarm signal is only  $\lambda_{\min} < -\text{floor}$ , which has never been observed; conversely the machinery breaks by  $\sim 13$  orders of magnitude the moment the Euler product is removed (the Epstein fire-wall), so the family positivity read does carry arithmetic information; the needles are frame-side ([verification/v667\_baez\_duarte.py]). *Kill criteria K1–K3 unchanged; no step silently upgraded; no RH statement. Sixth slice (2026-08-02, the Fejér-density round): the density plane closes; the needle search narrows threefold; the Li corollary. W2 — the density plane is closed* ([verification/v669\_fejer\_density.py]): the total window symbol *is* the Fejér-smoothed zero-counting density — the exact identity  $s_{\text{tot}}(\cdot; a) = 2\pi(F_a \star dN)$  (Weil explicit formula on the tent test pair, verified  $2.9 \times 10^{-6}$  pointwise /  $7.2 \times 10^{-8}$   $t$ -averaged against the Turing-certified comb) — and the *unconditional* chain RvM + Trudgian yields  $\rho_{a,\delta}(t) \geq 0.306 \log(2+t) - C_0$   $a$ -uniformly at  $\delta \in \{4\pi, \pi a\}$  with 10/10 finite certificates on the zero-free grid (min margin +1.36). The honest floor: the literal  $1/a$  and the plane-wave  $\pi/a$  widths sit *below* the RvM threshold  $4\pi A_1 = 1.4074$  — per-mode symbol control is not certifiable this way at any height (the structural explanation of the Gårding drift), so the remaining W2 piece is the packet-to-point translation below the floor. **W3 — the needle search is narrowed threefold**: the needles are *not low-rank* (the  $(1+K)$ -block deflation fails honestly: leak median 1.000 at every  $K \leq 8$ , fidelity  $p_{90} \approx 49\%$ ; [verification/v670\_w3\_block\_deflation.py]), *not zero-comb-driven* (the Lehmer-pair link is a clean null:  $h^3$ -partials all  $p > 0.1$ , placement control unremarkable; [verification/v671\_lehmer\_resonance.py]), and *not in the BD frame* ([verification/v667\_baez\_duarte.py]) — the mechanism search moves frame-/assembly-side. **New — the Li corollary and the  $E_4$  lock**: the W1 window positivity *contains* 30/32 finite Li coefficients ( $g_n = \frac{1}{2}e^{-|u|/2}L_{n-1}^{(1)}(|u|)$  derived exactly; band  $2.4 \cup 6..32$  on  $h = 1433$ , pole-free hyperplane, odd sector certified by  $\lambda_{\min} = +1.516 \times 10^{-3} > 0$ ; [verification/v672\_li\_corollary.py]), and the  $E_8$  shell data and the  $\zeta$  comb agree in *one* Li sequence ( $\Lambda(E_4, s)$  completed with residues  $\pm 1/240$  — the  $E_8$  240 is the residue normalizer;  $\lambda_n^L = 2\lambda_n^\zeta$  exact; budget usage 0.124; [verification/v673\_li\_e4.py]). *Kill criteria K1–K3 unchanged; no step silently upgraded; no RH statement. Seventh slice (2026-08-02, the consolidating contraction): the W2 end-state; the W3 contraction to the carrying form; two theorems and the honest typing of the wall. W2 — the end-state*: the density plane is closed ([verification/v669\_fejer\_density.py]); the Gårding statement now stands in *frame form with theorem constants* ([verification/v674\_packet\_garding.py]): the frame inequality  $Q + C_0 G \geq c_0 Y$  holds at every  $M$  with the v669 constants rebuilt zero-free,  $c_0 = 0.3058$ , margins  $+1.52.. + 2.29$ ,  $a$ -uniform — while the packet *projection* norm honestly does not escape the  $1/\log$  drift: typed residual with inverted expectation, the minimizer single-mode tight at 0.90); the Mosco mechanism is numerically complete (8/8 sequences, same module); the named remainder is *within-packet equidistribution* ( $\theta \rightarrow 0.20$ ) = unconditional zero-gap information. **W3 — the contraction to the carrying form**: after the needle mechanism search *saturates* (all four assembly-side candidates miss: lag-quantization weakly significant  $p = 0.007$

but too dense, phase gradient wrong sign, lock rotation cleanly killed, pole weight significant  $p = 2 \times 10^{-4}$  but inverted — the needles are fourfold characterized as an emergent bulk/frame property; ([verification/v675\_needle\_mechanism.py]), the contract W3 form becomes “*MARGIN-regime generic bound + needle risk map*” rather than a uniform bound with an exception predicate. The candidate mechanism  $C = 1$  is demystified *and* grounded as a *quadrature theorem* ([verification/v676\_c1\_mechanism.py]:  $q_r$  = the exact zero-side reading of the lock profile, median residual  $1.1 \times 10^{-2}$  on 67 windows;  $h^{-1}$  = the DST normalization; the constant = the zero-free RvM density integral at ratio 1.005 — *no RH lever beyond the density, but computable*;  $\sup \leq 1$  is not implied by the mean, the  $\sin^2$  sampling budget is the precise remainder with factor-2 headroom; the v596 exponent corrected:  $-1.01$  was flip-contaminated, clean  $-1.28$ ). And the *structure theorem* ([verification/v677\_w3\_structure\_theorem.py]) is the exact identity of the rung: window positivity  $\Leftrightarrow$  alias positivity of the zero symbol (Cantoni–Butler compression + sandwich + master identity with  $T_x \geq 0$  on the line)  $\Leftrightarrow$  an off-line detector whose Epstein-lab census (exactly 12 off-line zeros,  $\zeta_K$  control 0) predicts the form break at factor-2 level (ratio 0.803), with the threshold map  $2\alpha s_{\min} \approx 1.4..2.6$  (the Ihara echo) — so *W3-on-the-family* is a theorem plus a detector certificate, while *uniform W3* over all windows is the conjecture itself (Weil 1952 / Yoshida 1992, RH via the localized criterion), typed as such and *not* an RH statement. *Kill criteria K1–K3 unchanged; no step silently upgraded; no RH statement. Eighth slice (2026-08-02, the zero-gap round): the W2 end-state, with the last gap carrying a number. The frame chain is now fully unconditionally grounded:* the best explicit unconditional zero-gap theorem ([verification/v678\_zero\_gap\_theorem.py]: the  $S$ -difference route with  $S_{\text{bound}} = \min\{\text{Platt } |S| \leq 2.5167 \text{ to } 3.06 \times 10^{10}; \text{Trudgian 2014; Bellotti 2025 Cor. 1.5, asymptotic floor } 4\pi 0.10076 = 1.26619\}$ , machine-verified on all 1999 Turing-certified comb gaps with 0 violations and min air  $2.05\times$ ) makes the v674 frame packet floor *unconditionally* positive: adaptive bands  $\delta_p = \kappa H_{\min}(b_p)$  hold  $\geq \kappa$  ordinates per band unconditionally, and the theorem-shaped floor  $0 < B_p \leq V_p$  holds on every packet (median worst-case discount 0.037) — zero-gap theorem + Fejér chain ([verification/v669\_fejer\_density.py]) ground the entire frame-Gårding statement in cited unconditional inputs. **The last gap (pointwise/projection form) carries a concrete number:** the packet projection ladder does not stabilize under adaptive re-partitioning either (typed residual with inverted expectation:  $6/6$  ( $\kappa, C$ ) cells miss, the minimizer stays single-mode at 0.898 — the obstruction is re-partition-invariant), and the pointwise main-lobe capture of a guaranteed zero *closes iff the explicit  $S(t)$  constant  $A_1$  improves from 0.10076 (Bellotti 2025) to  $< 0.09017 = 1/(4a_0)$  — an 11.7% improvement of a named constant at the frontier of explicit analytic number theory, independent of RH*; the Platt branch bottoms out at  $1.4183 > \pi/a_0 = 1.1331$ , so no current chain reaches it at any height. This is the honest end-state of the W2 packet route: an unconditional frame form, a quantified external research target for the pointwise form, and within-packet equidistribution as the alternative input (W3/W4 territory). *Kill criteria K1–K3 unchanged; no step silently upgraded; no RH statement. Ninth slice (2026-08-03, the pinch-break and hole-closure round): the W2 end-state — the pointwise map at the anchor is gapless. The pointwise density floor at the anchor window is now a gapless split-type map* (comb-certified / scan-certified / abstract-Weil / unconditional-asymptotic), closing the 11.7% pinch and the  $(2500, 7.27 \times 10^6)$  hole. **The pinch was a bookkeeping artifact** ([verification/v680\_pinch\_attack.py]): centered capture — the same zero-gap theorem spent two-sided,  $\text{dist}(t, Z) \leq H_{\min}(t - 26)/2$ , comb-verified on 6595 grid points with 0 violations — *doubles the threshold to  $A_1 < 1/(2a_0) = 0.18034$ , so the existing Bellotti constant 0.10076 passes with 79% headroom; and the Beurling–Selberg minorant eliminates  $A_1$  from the counting chain entirely* (the loss is exactly the  $\pi/a_0$  minorant mass plus the explicit prime term  $P = 2.534$ ; positive from  $t^* = 1.11 \times 10^7$ , every window including  $a \rightarrow \infty$ ; Weil identity on the comb at  $9.0 \times 10^{-5}$ ). **The hole is closed** ([verification/v681\_coverage\_hole.py]): the exact prime term with a hierarchical Lipschitz certificate moves the abstract entry down a factor 72.5 to  $t_x = 1.53 \times 10^5$ , and a budget-certified Riemann–Siegel scan (223 949 zeros, Gabcke remainder, correctness 603/603 against the certified strip, found/expected 0.99774) carries  $(2515, 1.56 \times 10^5]$  with floor 0.01664 — the statement:  $s_{\text{tot}}(t; a_0) \geq 0.02259 \log(2 + t) - 0.5185$  for *all*  $t \geq 10$ , every region with named support and type

(H1 honestly adjudicated: the Platt-constant chain reproduces the hole boundary exactly and cannot enter). **The W2 chain at the anchor:** density ([verification/v669\_fejer\_density.py]) + frame-Gårding ([verification/v674\_packet\_garding.py] / [verification/v678\_zero\_gap\_theorem.py]) + pointwise ([verification/v681\_coverage\_hole.py]) — complete, with the mixed typing ledger explicit. *Remaining:* deep family windows  $a > 4.43$  (Selberg-only entry), the formal Mosco writeup, the projection-norm form,  $a \rightarrow \infty$  (= W3/W4 territory). *Kill criteria K1–K3 unchanged; no step silently upgraded; no RH statement. Tenth slice (2026-08-03, the great bundling): the master contract PRIME.UNIFPOS.01 and the five proof offensives. The master contract — the Uniform Positivity Theorem (PRIME.UNIFPOS.01, registered as a new contract).* Target statement, verbatim: *for every  $a > 0$ , Suzuki’s localized Weil operator  $B_a$  on  $H_1^0(-a, a)$  is positive semidefinite, with the positivity derived from an operatoric/geometric representation whose definition uses neither Riemann zeros nor an RH-equivalent presupposition.* The machine-checkable intermediate form:  $B_a = L_a + K_a$  with  $L_a \succeq 0$  from a structural source and the relative bound  $|\langle K_a v, v \rangle| \leq \theta \langle L_a v, v \rangle$  with  $\theta < 1$  uniform — then  $B_a \geq (1 - \theta)L_a \succeq 0$ . **Annotated with the v682 verdict** ([verification/v682\_lk\_split\_theta.py], LK-SPLIT-DIES): the two-sided form for smooth zero-free  $L$  is structurally dead — measured  $\sup \theta$  per split 11.3–58.0 with every smooth  $L$  indefinite on the odd sector, and the death mechanism is a *positive* find: the extremal pencil direction sits at  $t = \gamma_1 = 14.13$  (found from primes + digamma alone, no zero loaded),  $\mu_{\max}$  follows the spike law  $1 + 2\alpha/\Omega(t_{\text{peak}})$  (median dev 0.127),  $\theta$  grows  $\sim 2.47\alpha$  — *the deployed windows resolve individual zeros*, so the surviving directions are *one-sided* forms (the lower side alone — which must beat the pure-renaming floor  $\theta = 0.998$  of the trivial spectral split) or *multi-resolution* splits (an  $L$  carrying the comb spikes at matched resolution; untested, the open lever). *Kill criteria:* (K1) a  $C_a/L_a$  defined through Riemann zeros = renaming; (K2) a contractivity statement that is itself RH-equivalent = renaming; (K3) any candidate binding at the trivial-split floor adds nothing beyond measured window positivity. **The rank-3 chain reaches theorem shape at the sign level:** [verification/v683\_rank3\_functionals.py] lifts the prime influence on the load-bearing  $2 \times 2$  block into three exact functionals ( $4.9 \times 10^{-16}$ ) and adjudicates every entrywise route as class (c) (the T-B absorption margin demands  $\sim 10^6 \times$  more than any per-entry input; no Cauchy–Schwarz win: 99.7% of atoms violate pointwise; even a nonnegative pretentious measure breaks  $\det S$ ); [verification/v684\_rank3\_zeroside.py] treats the one functional by the explicit formula (residuum  $\leq 6.9 \times 10^{-6}$ , 22 491 budget-certified zeros) and delivers the *unconditional*  $\kappa_{\text{unc}} = 0.039\text{--}0.190 < 1$  (class (a); the pretentious escape blocked  $\times 634$  by the classical zero-free strip); [verification/v685\_rank3\_uniformity.py] closes the uniformity quantifier: the symbolic envelope  $K_{\text{env}} \leq 0.9798 < 1$  for all 62 complete windows with  $a \geq 3.434$  plus five finite certificates  $\kappa' \leq 0.11$  —  $\det S > 0$  *unconditionally on the entire declared surface* (SURFACE-CLOSED; a surface theorem on the declared family, *not* uniform in all  $a$ ; T-B stays prob:R1). [2026-08-03 prob:R1 update: T-B is now typed and largely closed. [verification/v692\_rank3\_lockgram.py] (7/7) types T-B(block) exactly:  $\hat{A}_2 =$  positive zero-Gram + psd rank-1 pole layer, the margin is the transverse zero mass (a sum of squares), on-line zeros only help (psd superadditivity) — “E kippt nicht” is a quantified partial-RH tail statement, neither a 2c-envelope question nor a W3 renaming. [verification/v693\_rank3\_density\_close.py] (5/5) builds the cited explicit bounds (Platt–Trudgian  $3 \times 10^{12}$ ; Hasanalizade–Shen–Wong count; functional-equation halving; the explicit Ingham-form zero density, arXiv:2507.15184) into a sharpened safe-side envelope (sinh difference pairing,  $V(0) = 0$ , on-line subtraction; PEN down  $\times 6.5\text{--}14.4$ ): 60/70 **complete windows close unconditionally-modulo-citations**; the exact open remainder is  $h \in \{1359, 1721, 1868, 2018, 2093, 2472, 2475, 2630, 2656\}$  with  $T^* \in [1.0 \times 10^{13}, 3.0 \times 10^{13}]$  plus  $h = 5690$  at  $T^* = 8.5 \times 10^{14}$ ; extension beyond the declared family grows  $\sim e^{3.74a}$ .] [2026-08-03 prob:R1 update II: [verification/v712\_rank3\_transverse\_deck.py] (5/5) scalarizes the v692 margin identity along the transverse direction — the penalty pays only the transverse zero mass, median gain  $\times 4.23$  on the open ten — and the census closes nine of them: 69/70 **complete family windows now close at the cited  $3 \times 10^{12}$**  (unconditional-modulo-citations); the single remainder is  $h = 5690$  with  $T^* = 1.6 \times 10^{14}$  (down from  $8.5 \times 10^{14}$ ); the deck-triplication lever is honestly dead

(triangle inequality: channel sups can only lose). *T-B itself stays the conjecture.* **The kernel class collapses** ([verification/v687\_extremal\_kernel.py], KERNEL-CLASS-TOO-NARROW): the pinning certificate  $2\pi F_{a_0}(m\pi/a_0) = 0$  (sympy-exact) plus Shannon sampling collapse the entire admissible band-limited positive class onto  $\{\lambda \text{ Fejér}\}$  with class supremum  $\theta_{\text{req}} = 0.0901684 < 0.10076$  — no kernel optimization can beat the  $S(t)$  constant; W2 stays closed via the v680/v681 route. **The W3 falsifier is constructive** ([verification/v688\_interpolation\_detector.py], INTERP-FALSIFIER-CONSTRUCTIVE): every isolated off-line quadruple is detected from  $2\alpha\delta \geq 1.974$  ( $C_{\text{up}} = 0.987$ ), a  $\times 2$  improvement over the v677 mode map, with masking fully adjudicated (46 broken + 2 provably sub-resolution of 48). **The geometric source opens a contract** ([verification/v686\_geometric\_sos.py], PRIME.GEOMSOS.01):  $\Lambda(n)$  reconstructed circle-free from E8 shell counting ( $7.1 \times 10^{-31}$ , Satake data from counts, K1 passes), canonical dim-3 cover-SOS, K4 firing honestly; missing: the bundle  $H$  and the transport  $\Phi$  (top-3 eigenvalue mass only 91.4%), milestones M1–M4 with kill criteria K1–K5. **The Hecke-SOS mechanism is extracted** ([verification/v691\_hecke\_sos.py], HECKE-SOS-IHARA-MECHANISM-EXTRACTED): the target factorisation  $A = B^*B + P$  exists *exactly* in the Ihara lab ( $B$  by Chebyshev recursion, no Cholesky;  $P \succeq 0 \iff$  Ramanujan), the deployed  $\zeta$  window form is exactly the sine/defect half of the canonical split (the part whose lab positivity is the RH analogue), the Euler mechanism is measurable (resonance lattice  $2\pi j/\log 2$  exact, depth  $\times 1.98$ ), and the missing ingredient is exactly named: Z1 — a self-adjoint geometric operator whose polynomial traces are the window moments (Hilbert–Pólya); Z2 would be RH itself. *Kill criteria K1–K3 unchanged; no step silently upgraded; no RH statement.*

**The review contracts of the great bundling (registered 2026-08-03, brief).** **E8.GAUSSIAN.CODE.01** —  $E_8(\mathbb{Z}[i])/(1+i)$  is the information space of  $\text{RM}(1,3)$ : first slice executed by [verification/v689\_gaussian\_code\_bridge.py] (26/26, GAUSSIAN-CODE-BRIDGE-EXACT: SNF-exact quotient  $\mathbb{F}_2^4$ ,  $240 = 15 \times 16$  census with the zero class provably empty,  $\sigma$  = the  $\text{RM}(1,3)$  family permutation with residual gauge of order 18, coordinate block =  $F_1 + F_2 + F_3$ ; all must-fail controls fire); remaining: the theorem writeup. [2026-08-03 addendum: the algebraic core is now formally verified — kernel-checked in Lean 4 (TfptCarrier/GaussianCodeBridge.lean, lake build green, no sorry/native\_decide): the  $\mathbb{Z}[i]$ -module with the SNF certificate (unimodular  $P, Q$  both directions), the 240-root census  $15 \times 16$  with the zero class provably empty,  $\mu_4$  stability, and the  $\sigma$  family action with coordinate block =  $F_1 + F_2 + F_3$ .] **ARF.BOUNDARY.CODE.01** (registered 2026-08-05, [O]) — boundary code uniqueness: the non-circular selection chain boundary  $\rightarrow$  code  $\rightarrow E_8 \rightarrow \mathbb{F}_2^4 \rightarrow$  five slots *exists* with  $g_{\text{car}} = 5$  as output ([verification/v776\_boundary\_hamming\_uniqueness.py], 31/31, BOUNDARY-CODE-UNIQUE: census  $135 \rightarrow 30 \rightarrow 14 \rightarrow 2 \rightarrow 1$  without Hamming/ $E_8$  input; the flag datum is *one* deck-invariant  $\mathbb{Z}_2$ , the product of the four side orientations; Construction A of the selected code *is*  $E_8$ ; the Arf selector is unique; packets  $1 + 4$ ). *The open contract form as registered:* derive evenness/self-duality of the mark bookkeeping from the seam axioms — the R1 *universe premise* (“the boundary realizes a doubly-even self-dual binary code on its 8 sides”) was the relocated and shrunk circularity, joined by the typed residuals R2 (deck placement is semantic, not selective) and R3 (the anchor naming  $\bar{b}(F_\Sigma, A) = 1$  is part of the frozen selector, not a theorem). [2026-08-06 note, markers unchanged: all three residuals moved — R1 is replaced by R1' (the three physical axioms of SEAM.CODE.TYPEII.01 below, [verification/v799\_seam\_code\_typeii.py]); R2 is placed and R3 is named by the metaplectic  $S$ -lift (SEAM.CLIFFORD.MODULAR.S.01 below, [verification/v798\_seam\_clifford\_modular\_s.py]).] [2026-08-06 round-22 note, markers unchanged: the attack surface sharpens — affine functions on the NS/R torsor reconstruct  $\text{RM}(1,3) = H_8$  with  $\text{Stab}(\chi_{\text{NSR}}) = \text{AGL}(3,2) = \text{Aut}(H_8)$  in its natural 8-point action ([verification/v803\_e8\_jetcode\_affine\_nsr.py]); a typed internal reconstruction, not a P2 removal — R1' unchanged.] *No P2-removal claim; kills:* a second boundary-legal code surviving the flagged selection, or a seam-axiom derivation forcing a non-doubly-even universe. No marker move. **SEAM.CODE.TYPEII.01** (registered 2026-08-06, [C]) — Type II from the seam axioms: at length 8 (four  $J$ -paired marks  $\times$  two sides) the finite chain is exhaustively machine-verified over all 308,993 subspaces of  $\mathbb{F}_2^8$  of  $\dim \leq 4$  ([verification/v799\_seam\_code\_typeii.py], 19/19, TYPEII-FORCED): mutual locality of the Construction-A boundary vertex words  $\iff C$  self-orthogonal; integer

conformal spin (trivial  $T$  sector phases)  $\iff C$  doubly even; holomorphy/index one ( $|\text{disc}| = 1$ ,  $S$ -closure of the single character)  $\iff C$  self-dual — the survivor set is exactly the 30  $\hat{e}_8$  copies and no boundary-legal non-Type-II code exists. Chained to `ARF.BOUNDARY.CODE.01`, the full route seam axioms A1–A3 + deck + orientation bit  $\rightarrow$  unique code  $\rightarrow E_8 \rightarrow \mathbb{F}_2^4 \rightarrow g_{\text{car}} = 5$  stands with the single **[C]** residual R1': *the physical seam satisfies* A1 mutual locality (bose statistics / trivial monodromy), A2 integer conformal spin (modular  $T$ -alignment), A3 holomorphy/index one (a single modular-closed sector) — each a standard physical demand, not a coding-theory premise. Imported lattice-VOA/theta ingredients are classical, cited and machine-verified in every instance used; the OPE-linearity scope note is typed. Kill: a seam construction violating A1–A3 while staying boundary-legal. No P2 marker move. `SEAM.CLIFFORD.MODULAR.S.01` (registered 2026-08-06, partial state) — the missing Clifford bit as the metaplectic lift of modular  $S$  at  $\tau = i$  (`[verification/v798_seam_clifford_modular_s.py]`, 27/27, MODULAR-S-PARTIAL): *coset decision positive* — the per-binary-factor Hadamard lands exactly in the missing coset  $\zeta_8 G_{31}$  with the exact fixed-point relations ( $S^2 = I$ ,  $T^2 = \text{Pauli parity}$ ,  $(TK)^3 = \zeta_8 I$ ); the deck  $J = iI$  is the metaplectic anomaly of the total symplectic Fourier (R2 placed) and  $\zeta_8 = (1+J)/\sqrt{2}$  the real half-deck swapping the seam pair; the strict uniqueness census over all 46080 coset elements is *empty* — no  $\Theta$ -real metaplectic involution fixes  $q^*$ , the Arf defect  $q^* \mapsto q^* + \bar{h}(\cdot, c)$  is unavoidable — with exactly one  $G_{31}$ -conjugacy class of retyped survivors, and the  $\sigma$ -orbit Arf defects are exactly  $\{A + F_\Sigma + F_i\}$  with XOR = the anchor  $A$  (R3 *named*, not fixed). *What remains:* derive the physical selection of the defect class (the anchor naming from seam data) or type the two-variant/defect residue as intrinsic gauge. Kill: a  $\Theta$ -real metaplectic coset element fixing  $q^*$  (excluded here exhaustively) would falsify the defect law. No marker move. `CURVE.CODE.OUTERSPIN.01` (registered 2026-08-06, *decided* CANONICAL) — the outer bridge between the code duad side and the curve syntheme side (`[verification/v796_duad_syntheme_bridge.py]`, 24/24, BRIDGE-CANONICAL): the Sylvester duality  $\beta$  with  $\beta\sigma = \tau\beta$  and  $\beta(q^*) = S^*$  (the unique  $\langle \tau, \rho \rangle$ -fixed spread) exists and is unique up to the order-6 code joint stabilizer (one orbit including the  $\tau/\tau^2$  swap); it resolves the v784 Arf mismatch canonically ( $q_\Theta$  pulls back constantly to  $q_{\text{even}}^* = q^* + \bar{h}(\cdot, A)$ , the anchor shift), transports  $\bar{5}/10$  and  $3+2$  exactly, and sends  $t_A$  to  $\rho = J$ ;  $\chi_{\text{NSR}}$  transports as an honest exact two-variant torsor pinned by the v752 NS/R parity — that two-fold ambiguity is intrinsic and typed, not open work. The `CURVE.CODE.2TORSION.01` inner-identification kill stands unchanged; no matter semantics (ROOTCLASS-MIXED). *[2026-08-06 round-22 note: the bridge sharpens into an incidence-operator statement — every census bridge factors exactly  $N_{\text{code}} = P_\Delta^T N_{\text{curve}}^T P_\Gamma$ , a transpose duality, not an isomorphism (`[verification/v809_curve_code_doily.py]`).]* No marker move. `E8.RAMIFIED.JETCODE.01` (registered 2026-08-06, **[E]**; carrying `E8.AFFINE.NSR.01`) — the jet-code normal form of the Gaussian compiler (`[verification/v803_e8_jetcode_affine_nsr.py]`, 18/18 + 15/15, JETCODE-EXACT+AFFINE-NSR-EXACT): because 2 ramifies,  $\mathbb{Z}[i]/(2) \cong \mathbb{F}_2[\epsilon]/(\epsilon^2)$  (unique iso; the split ring fails all 24 bijections) and  $W = L/2L$  is *free* of rank 4 over the dual numbers; the deck is the jet unit  $J = 1+\epsilon$ ; among all 65536 sections of  $0 \rightarrow \epsilon W \rightarrow W \rightarrow V \rightarrow 0$  there are *zero*  $\mu_4$ -equivariant splittings — non-splitness is the v782 no-origin obstruction;  $240 = 15 \times 8 \times 2$  (root = address  $\times$  error jet  $\times$  orientation). Kill: a  $\mu_4$ -equivariant splitting, or a second iso class of the fiber ring. No marker move. `E8.SYNDROME.ALGEBRA.01` (registered 2026-08-06, *dead* as frozen) — the sharpest error-correction-hull hypothesis killed as preregistered (`[verification/v805_e8_syndrome_algebra.py]`, 13/13 incl. the must-fire controls, SYNDROME-DEAD): max rank 222/256 over all four frozen phase conventions; the  $\pm r$  proportionality obstruction structural for the whole diagonal-phase family; both post-hoc escapes score worse. Survives exactly:  $16 + 240 = 256$ , the 16 neutral kernels = the diagonal MASA, same-class products  $A_0$ -valued over  $\mu_4$ . Successor criterion named: root-individual operator data beyond diagonal phases. The Construction-A Hamming theorem untouched. `CURVE.CODE.DOILY.01` (registered 2026-08-06, **[E]** + **[O]**) — the doily incidence identity (`[verification/v809_curve_code_doily.py]`, 21/21, DOILY-EXACT):  $NN^T = B + 2I$  entrywise on the  $K_6$  doily (15 duads  $\times$  15 syntheses), spec =  $\{9, 4^9, 0^5\}$ ;  $N/3$  has singular values  $\{1, 2/3, 0\}$  — the recovery base rate  $2/3$  as the unique nontrivial singular value of the canonical duad-to-syntheme incidence; Petersen edges = syntheses 15/15; the Arf switch set  $T_6 =$

the six within-block duads. *Open target*: the six-step composition identifying  $(2/3)^6$  itself. [2026-08-06 round-23: decided — SIXSTEP-CLOCK-ONLY ([verification/v814\_k5\_sixstep\_transport.py]): the rate arises canonically on three spatial routes (the doily alternation among them), but each carries a typed obstruction and sixth-power blindness makes per-step data mandatory — the deployed six is the clock exponent; the missing ingredient is a deployed shells-to-cusp isomorphism, named.] Kill: any census bridge failing the transpose factorization. No marker move. **CARRIER.CUTCODE.01** (registered 2026-08-06, [E]; carrying CARRIER.MOMENT.INCIDENTCE.01) — the Arf split as the cut classification of  $K_5$  ([verification/v810\_cutcode\_moment.py], 17/17 + 16/16, CUTCODE-EXACT+MOMENT-INCIDENTCE-PARTIAL): the cut map injective-isometric on all 256 cells;  $16 = 1 + \bar{5} + 10$  is  $(0|5) + (1|4) + (2|3)$ ; Clebsch  $40 = 5 + 20 + 15$  over the Arf shells with within-10 = Petersen; roots/lines grade  $80 + 160 / 20 + 40$  — Arf/cut grading, not matter; the master moment  $\text{Tr}_S + X^2 = 4|X|^2 = 4h(E_8) = 120$  one symbolic identity. The frozen 210 decision:  $\binom{10}{4}$  decorative (equivariant bijection killed by the exhaustive fixed-point obstruction; the partial  $80 \rightarrow 5 \times 16$  miss-slot map and the Newton tie survive as audit). No marker move. **CARRIER.PETERSEN.RADIAL.01** (registered 2026-08-06, [E] + [H neu]; carrying TRANSPORT.SIXTHROOT.01) — the  $K_5$  edge machine ([verification/v808\_petersen\_sixthroot.py], 19/19 + 16/16, PETERSEN-RADIAL-EXACT+SIXTHROOT-SPECTRAL-ONLY): charges  $(-4, +1, +6) =$  the certified  $SU(5)$  decuple; the edge graph is Petersen  $\text{SRG}(10, 3, 0, 1)$ ; distance shells  $(1, 3, 6)$  charge-pure; the flavor hexagon is the distance-2 shell edge-by-edge;  $\text{spec}(P^6) = \{1, 64/729, 1/729\}$  is the deployed transport spectrum exactly; the 10-dim multiplicity reading ( $\times 5 = g_{\text{car}}$  fast,  $\times 4 = |\mu_4|$  slow) typed [H neu]. The sixth root:  $B = \frac{1}{18}[[13, 1, 4], [1, 13, 4], [4, 4, 10]]$  unique symmetric doubly stochastic PSD in the corpus basis,  $B^6$   $GL(3, \mathbb{Q})$ -conjugate to but bit-equal to *neither* deployed form — the corpus freezes the spectrum, not the basis (the basis-freeze question is the named open). [2026-08-06 round-23 dated correction + addendum ([verification/v814\_k5\_sixstep\_transport.py]): the sixth-root census had missed the deployed **v221** evaluation —  $T_{v221}$ , the exact spectral sum of **v221**'s transfer  $(1 - w)^6$ , equals  $B^6$  bit-exactly, so the verdict upgrades from SPECTRAL-ONLY to BITEACT-vs-v221 and the basis-freeze question is answered on the deployed route (the frozen SPECTRAL-ONLY stands as the probe's adjudication; ABEL-DEAD correction precedent); and the [H neu] 10-dim reading now has a promotion-ready canonical PROPOSAL:  $T_{10} = P_{10}^6$  in the  $\Lambda^2$  basis passes all gates with the equitable **v774** quotient  $= (Q_{\text{Pet}}/3)^6$  bit-exact (TRANSPORT10-PROPOSED, no deploy claim).] Kill: a second PSD sixth root in basis, or a shell/charge mismatch. No marker move. **E8.QUARTIC.HALF.01** — G31 as the holomorphic  $\mu_4$  shadow of the E8 compiler: first slice executed by [verification/v690\_quartic\_half.py] (22/22, QUARTIC-HALF-ALIVE) with the honest correction: the vanishing half  $\{2, 14, 18, 30\}$  is true but trivially  $\mu_4$ ; the substance is  $F_8|, F_{12}|, F_{20}|, F_{24}|$  as a system of basic G31 invariants (Chevalley cited; the holomorphic restriction generates the full G31 ring); remaining: the symbolic degree-60 Jacobian factorization and the writeup. [2026-08-03 addendum: the  $\mu_4$ -orbit vanishing mechanism and the quartic power sum  $\sum_{\alpha} \langle \alpha, x \rangle^4 = 576 (\sum_i x_i^2)^2$  over the explicit 240 roots are now kernel-checked in Lean 4 (TfptCarrier/QuarticHalf.lean, lake build green, no sorry/native\_decide); the Chevalley basic-invariant step remains cited, not formalized.] **PRIME.W3.INTERPOLATION.01** — the alias-interpolation falsifier: first slice executed by [verification/v688\_interpolation\_detector.py] (constants above); remaining: the lemma as a paper lemma. [2026-08-03 addendum: the two named building blocks are closed near-proof by [verification/v694\_interpolation\_lemma\_closure.py] (20/20, CLOSURE-BOTH-NEAR-PROOF): retention via an exact projection identity plus a closed-form  $O(k^2)$  certificate ( $r_f \geq 0.548$  on the family surface; straddle boundary case discovered,  $k = 0$  branch secured), and the separation law form-corrected to the additive  $\alpha^* = C_{\text{cell}}/\delta + C''/\Delta\gamma$  ( $C'' \leq 0.59$ , zero exceptions in 97 tests); the two open parent configurations detect at  $\alpha^* = 7.90$ ; declared residues: the explicit Ingham constant chain, the Dirichlet-kernel overlap bound, the  $C''(n)$  induction.] **PRIME.Z1.OPERATOR.01** — OPEN (Offensive 5 running): construct the self-adjoint geometric operator whose polynomial traces are the deployed window moments (the Hilbert–Pólya step of the extracted Ihara mechanism), without zeros as input; an operator defined through zeros or an RH-equivalent norm assumption = renaming; Z2 (the norm bound)

would be RH itself — the factorisation *localizes* RH, it does not bypass it. [2026-08-03 update, the 5–5d series ([verification/v695\_z1\_trace\_operator.py] / [verification/v696\_z1\_jacobi.py] / [verification/v697\_z1\_uvarov.py] / [verification/v698\_z1\_flow\_recursion.py]): the contract stays OPEN, with the ground taken recorded honestly — (5) the measure comes from counting [E]: seam and orbifold are dead as Z1 spectrum (null-calibrated), while the E8 counting route delivers the Weil measure exactly (atoms 0.0, arch/pole via the closed  $\Gamma$  duplication split, 7/8 target spikes at  $q = 0.000$ ); the comb sequence  $c + \text{pole}$  is positive-feasible — the signedness of the deployed sequence is the pole subtraction. (5b) the canonical operator exists (CMV/Jacobi, all moments reproduced at  $2 \times 10^{-13}$ ); no closed coefficient law yet, but the renaming clause is explicitly refuted: features on prime-power slots ( $q = 0.000$ ), amplitude-mass link  $r_{\log} = +0.76$ , positivity load-bearing (170/171 vs 2865). (5c) masses: atoms are lag insertions (duality proven), exact transfer identity  $\Delta\alpha = w_1/E$  ( $5.6 \times 10^{-17}$ ), and the inverted stabilization law — the  $\Gamma$  flow predicts the counting masses to  $\sim 10\%$  (median ratio 1.026, corr 0.86). (5d) positions forced (windows 0.5–2 cells, jitter null 30/30, adversary –516 lags); shooting recovers the masses at the true positions to 0.11% median. Remaining: lookahead autonomy (the greedy-saturation failure of autonomous reconstruction) and the continuum reading of the operator.] [2026-08-03 update II, the montage and the moonshot: the contract stays formally OPEN — the continuum theorems are missing — but both remainders moved. (i) The lookahead question is decided ([verification/v702\_z1\_lookahead.py], 13/13): the  $\Gamma$  flow is a full-depth comb verifier ( $913 \geq 898$ , every fake dies  $\geq 184$  lags earlier) but not a generator — both prescribed lookahead routes reach 2–4 slots, and the per-slot information demand  $g \sim 5\text{--}12$  bits is information-theoretic and parameter-robust (FLOW-VERIFIER-NOT-GENERATOR). (ii) The chain arc  $S\text{--}A\text{...}S\text{--}G$  ([verification/v703\_chain\_tolerance\_scaling.py] through [verification/v709\_chain\_zero\_layer.py] plus [verification/v710\_chain\_two\_stage\_hecke.py] / [verification/v711\_chain\_firstbirth\_scaling.py]) types the mass law honestly: tolerance induction dead (RH-grade), the arch is a deck-sector trace [E] ([verification/v705\_chain\_deck\_sector.py],  $3.4 \times 10^{-16}$ ), the flow  $\alpha$  are the Schur parameters ( $10^{-40}$ ), the section edge is an exact resolvent identity, and  $G1 = [\text{E-edge}] + [\text{E-Hecke}] + [\text{M-rate}]$  with the corridor position  $\sim 0.53$  as the named open selection question. (iii) The continuum candidate is assembled ([verification/v713\_l1\_montage.py], 16/16, L1-ASSEMBLED-MEASURED): cover limit  $6.2 \times 10^{-8}$ , point masses at the  $D^2$  rate, one UV slot, divergence exactly at  $s \leq 1/2$  (the critical line as the boundary of the construction — a consistency signature, not a positivity statement); the missing theorems (self-adjointness, trace identification, positivity) are named. (iv) The atom layer is groupoid-internal ([verification/v714\_moonshot\_hecke\_groupoid.py], 19/19, MOONSHOT-GO; see PRIME.Z1.MOONSHOT.01 below). What is missing is exactly the continuum SÄTZE: until they exist, no marker moves.] [2026-08-03 update III, the evening round: the moonshot stages 2–4 plus the Satz-1 slice and the  $K$  collapse are done and machine-checked (see PRIME.Z1.MOONSHOT.01 below for the six modules [verification/v716\_moonshot\_arch\_glue.py] through [verification/v721\_k1\_node\_capture.py]). The Hilbert–Pólya candidate is now structurally complete at measurement level — an operator family, a state, an exact finite trace formula with a closed term dictionary to Weil 1952, and a measured spectral limit onto the zeros (100% of 377 at tol 0.25, rate  $-1.61$ , SHA256 prediction firewall). The contract stays formally OPEN: the remaining hard substance is (L1) the identification of the vague limit at the line  $s = 1/2$  (no mass loss — the single place the zero fluctuation enters) and (L2) the node convergence as a theorem (the capture half is proof-near, [verification/v721\_k1\_node\_capture.py]; the  $\sim 900\times$  super-resolution is the open K1b question); no marker moves, no RH statement.] PRIME.Z1.MOONSHOT.01 — OPEN (registered 2026-08-03, the groupoid program): derive the full Weil measure (atoms + arch + pole) as the trace layer of one geometric groupoid built on the  $\mathbb{Z}[i]$ -E8 Hecke tower, with the arch layer *glued* (not direct-summed) and the E8 specificity established. Stage 1 (done): [verification/v714\_moonshot\_hecke\_groupoid.py] (19/19, MOONSHOT-GO) shows the atom layer is groupoid-internal — the primitive degrees of the tower are *exactly* the Gaussian prime norms (maximality census; no prime projector anywhere: the primes come *out* of the tower), the log generator + cell normalizer de-divisorize the tower traces at  $4 \times 10^{-13}$ , and the  $\sigma$  descent reproduces

100 atom positions, 100 masses and 368 moments at deviation 0.00; declared ingredients (named as inputs, not derived): the  $\sqrt{n}$  half-density and the  $\sigma$  descent. *The honest fence (typed)*: the finite counting facts of stage 1 are *E8-unspecific* — the specificity question belonged to stage 2, which has since been executed. *Stage 2 (done, 2026-08-03)*: `[verification/v716_moonshot_arch_glue.py]` (18/18, MOONSHOT-STAGE2-GLUED) — the arch *glues*: one object, *one* normalization (the free 3-scalar LSQ returns  $\kappa = (1, 1, 1)$  to  $< 10^{-12}$  on all five windows); the  $1/4$  in  $\operatorname{Re}\psi(1/4 + i\pi/2)$  is *derived* as the  $\mu_4$  fixed-sector offset;  $\log \pi$  enters only through the declared self-dual (Tate) UV cell — the exact arch mirror of the declared  $\sqrt{n}$  half-density; and E8 becomes forced *at the gluing* (the  $\mu_2$  tower, the wrong  $\mu_4$  class and the wrong deck twists all rip). *Stage 3 (done)*: `[verification/v717_moonshot_state.py]` (21/21, STATE-ON-TRUNCATIONS) — the Weil functional is *constructively* a GNS vector state on every truncation (Levinson PD); the naive tower reading misses by  $|\Delta|/W3\text{-margin} = 104\text{--}1563$  — the trace-formula boundary *is* the localized RH substance; KMS  $\beta = 1$  (Bost–Connes critical) measured. *Stage 4 (done)*: `[verification/v718_moonshot_spectral.py]` (10/10, SPECTRUM-CONVERGES-MEASURED) — the truncation spectra converge *measurably* onto the zeros: 100% of 377 frozen targets at tol 0.25 ( $h = 1433$ ), rate  $-1.61$ , monotone GNS  $K$ -ladder, SHA256 prediction firewall; the scramble control dies as a state. *The Satz-1 slice (done)*: `[verification/v719_moonshot_traceformula.py]` (24/24, SATZ1-LEDGER-EXACT) — the finite trace formula is exact Gauss quadrature and the term dictionary to the classical Weil formula *closes* (pole/atoms/arch/UV); the E8 census over  $\{\mathbb{Z}^8, E8\}$  (Mordell 1938): only the Gaussian E8 glues. *The K collapse (done)*: `[verification/v720_moonshot_k2k3.py]` (14/14, K2K3-COLLAPSED-CLASSICAL) — the convergence items collapse onto classical analysis (tightness measured, unconditional tail majorant on  $s > 1/2$ , the boundary exactly  $s = 1/2$ ); and `[verification/v721_k1_node_capture.py]` (14/14, K1-LEMMA-CAPTURE-ONLY) makes the capture half of the node question proof-near. *Kill audit (criteria from the strategy council, adjudicated)*: (K1) does *not* fire — the gluing needs no zeros (SHA256-typed diagnostics); (K2) half-standing — the  $\sqrt{n}$  half-density stays *declared*, its arch mirror  $\log \pi$  is now typed as the self-dual normalization; (K3) does *not* fire at measurement level — the counterfactual census rips (the uniqueness *theorem* stays open); (K4) does not fire — positivity holds on every deployed window. *The remaining substance, ordered*: (L1) identification of the vague limit *at* the line  $s = 1/2$  (no mass loss); (L2) node convergence as a theorem (the  $\sim 900\times$  super-resolution — the K1b question, exploration running); (L3) routine lemmas (tent/Jackson discretization, PNT tail majorant — both exhibited with constants). A research contract, *not* a claim; the contract stays OPEN; no marker move; no RH statement. *[2026-08-03 update IV, the keystone round — the contract at its final state of the day. Eight further modules ([verification/v727\_l1\_identification.py] through [verification/v734\_s1\_canonical.py]) localize the remaining substance maximally. (L2) is resolved classically: the  $\sim 900\times$  super-resolution is atom pinning of the raw measure — the memo’s locality/symmetry mechanism dies on two preregistered bars ([verification/v728\_k1b\_superresolution.py], 11/13, K1B-ATOM-PINNING; lab rate  $K^{-1.84}$ , real correlate 0.64), and the explicit pinning lemma is certified with one-line proofs and 0 violations at 598 peaks/535 zeros including all 377 of the largest window ([verification/v730\_strat2\_pinning\_lemma.py], 12/12); the certificate saturates at window width — sub-width certification needs atomicity beyond every window depth. The K1 diagram closes: capture theorem (v721) + pinning lemma (v730) + conditional gap rigidity ([verification/v731\_strat2\_gap\_universality.py], 5/5, GAP-UNIV-KILLED for the strong measured form; a fixed-window conditional candidate remains via Lubinsky/Levin–Lubinsky/Avila–Last–Simon/Totik) + atom structure (= L1). (L1) is typed in four equivalent languages: determinacy of the limit moment problem is classical and the identification chain contains RH nowhere ([verification/v727\_l1\_identification.py], 11/11, L1-DETERMINED-WALL-TYPED); THE WALL is PD persistence of the window ladder — infinitely many PD windows suffice for RH, each window finitely decidable, measured 9/9; equivalently: moments (Hankel-PD) / state (Levinson-PD) / symbol (Fejér  $\geq 0$ ) / Hamiltonian (Krein  $l_k > 0$  — [verification/v734\_s1\_canonical.py], 10/10, S1-CANONICAL-SUZUKI-BRIDGED: exact Krein correspondence at  $1.6 \times 10^{-15}$ ,  $H_h \geq 0$  measured on all nine windows, boundary phase =  $\gamma$ -side at scale exactly 1; Suzuki JFA 2021 as the fourth anchor, unconditional only for  $\omega > 1$ ). The  $h$ -universal quantifier is the RH*

substance. Honest kills, fenced: naive reflection positivity is dead on the truncations — the FE symmetry buys exactly the atom-free Connes–Consani 2021 regime  $|u| < \log 2$  and no more ([verification/v729\_strat2\_rp\_universality.py], 5/5); the direct Sonin-compression transfer is dead ([verification/v732\_strat3\_sonin\_prolate.py], 21/21 construction, SONIN-RANK-GROWS; the CC mechanism itself untouched, retuning forbidden). A first exact bridge into Connes’ framework: the  $\sigma$ -descended groupoid maps by an exact weak functor candidate into the  $S$ -local periodic-orbit subalgebra ( $S = \{\infty, 2, 3, 5\}$ , 6859 convolutions at deviation 0,  $\kappa \equiv 1$ ; [verification/v733\_strat3\_gate0\_census.py], 13/13, GATE0-WEAK-FUNCTOR; full-algebra extension open). The literature anchor as the dual council reads it: the wall has three names — Weil 1952, Connes 1999/Connes–Consani 2021, Suzuki 2021 — with Meyer 2005 the proven boundary line (unconditional spectral realization without positivity); TFPT adds the computable zeta-free truncation framework with emergent prime places and per-window decidability. The contract stays OPEN; no marker move; no RH statement.] [2026-08-03 update V, the closing slice — the L1 route portfolio is complete. The state-transport route also stagnates ([verification/v735\_strat3\_ucp\_inductive.py], 10/10 construction checks, UCP-STAGNATES): all eight adjacent GNS-ladder maps exist exactly as UCP quantile couplings — the Choi SDP reduces to a finite transport LP, unitality/state preservation to  $3.1 \times 10^{-12}/1.2 \times 10^{-15}$ , every Choi rank inside the sharp bound  $h + h' - 1$  with log-log slope 0.991 (linearly controlled) — but the preregistered shift/grade covariance gate fails honestly (error tail not monotonically decreasing): mere state transport is not the missing L1 mechanism, and the maps are quantile Markov couplings, not  $*$ -homomorphisms (multiplicativity defects  $\sim 5\text{--}8 \times 10^{-3}$ ). The portfolio now reads: Sonin dead (v732), UCP stagnates (v735), Gate 0 weak functor (v733), Suzuki bridge a measurement (v734) — what remains for L1 is atomicity/i-identification substance, not transport machinery. The contract stays OPEN; no RH statement.] [2026-08-05, SHARPENED — the diagonal-Gram closure hands over (PRIME.GRAM.DIAGONAL.01 is closed per its own frozen stakes; the qf offensive [verification/v769\_qf\_contract\_necessity.py] / [verification/v770\_qf\_spectral\_bundle.py] / [verification/v771\_qf\_representation\_census.py] / [verification/v772\_qf\_feshbach\_effective.py] / [verification/v773\_qf\_cell\_cocycle.py]). The sharpened target form: construct a self-adjoint geometric bulk operator with a canonical finite-dimensional boundary/threshold structure whose Weyl  $M$ -function delivers the qf block and whose relative trace formula reproduces the full Weil formula. The measured handover content — mandatory constraints any Z1 candidate must carry: (i) the Hecke/E8 bulk with the archimedean lift ([verification/v714\_moonshot\_hecke\_groupoid.py] / [verification/v716\_moonshot\_arch\_glue.py]); (ii) the moving-edge threshold structure with its measured cadence — mode entries at  $M \simeq 888/992/1108/1276$ ,  $\Delta X$  widening  $1.625 \rightarrow 2.625$  (spread 0.381, not regular at the 0.25 bar), avoided crossing 0.0039 at  $M = 1240$  ([verification/v772\_qf\_feshbach\_effective.py]); (iii) the ram-odd contact direction ( $\cos \theta_1 = 0.997$  rung-stable, [verification/v771\_qf\_representation\_census.py]); (iv) the settled per-function coupling levels (R2 family 0.225–0.358, R1 family 0.008–0.079, [verification/v770\_qf\_spectral\_bundle.py]); (v) the cell-ordered prime transport in domain-preserving (Nevanlinna/Loewner) form — the exact Möbius/Redheffer cell cocycle of [verification/v773\_qf\_cell\_cocycle.py]. Kills, active: an operator defined through zeros or an RH-equivalent norm assumption = renaming (unchanged); additionally, a candidate whose threshold structure cannot carry the measured widening cadence, or whose  $M$ -function block contradicts the settled qf coupling levels or the ram-odd contact direction, is dead on arrival. The contract stays OPEN; no marker move; no RH statement.] [2026-08-05, EXECUTED — the compactness split measured ([verification/v780\_z1\_compactness\_trilogy.py], one module from three preregistered probes; gates  $3/4 + 1/4$  legs  $3/6$  with exactly the preregistered fails). The N2 GNS/Jacobi family carries the ram-odd contact ( $\cos$  up to 0.9971), edge collisions at the source depth (0.0032 vs 0.0039) and settled couplings (14/14  $P$ ) natively — but not the widening cadence (its edge census is regular quadrature filling, the free-Jacobi type): the cell structure must be imported as boundary data. The finite boundary triple ( $A = 2 - J_K$ , imported Gram band) is exactly Herglotz through all typed transitions but dies as a whitened  $K = M/2$  import — the Gram near-kernel packets are delocalized (cell-mass exactly 0.500 beyond the half-window), so whitening

turns truncation loss into quasi-generic boundary noise (delivery 5.88; still arithmetic-specific: 5.9 vs GOE 13.0 vs random 132.5; the full-depth frame is provably out of reach by deepening,  $e^{27.75}$  atoms). With the raw  $\mu$ -weighted import the variable-edge Herglotz family is bounded (0.936) and equicontinuous (0.807) with summable  $\sim$ rank-one entry poles at the arithmetic thresholds ( $c_{\text{mass}}$  0.032  $\ll$  1; rank ratios 0.80–0.87), while the delivery metric loses discrimination (typed UNRESOLVED-BY-THIS-METRIC) and the moment functionals do not settle (osc 22.1). The measured unification: this contract and PRIME.KMS.INDUCTIVE\_STATE.02 share their compactness theorem (carried: Helly/Montel at finite level) and share their single remaining obstruction — state selection equals import-faithful boundary coupling (a coupling of the Gram band to the bulk that does not pass through the lossy half-window frame). The merged target is named Z1-COMPACTNESS (PRIME.Z1.COMPACTNESS.01); both rows stay, cross-referenced; any future module must attack the shared delta, not the two contracts separately. The contract stays OPEN; no marker move; no RH statement.]

**The review master contract (PRIME.KMS.INDUCTIVE\_STATE.01).** OPEN (registered 2026-08-04, from the external review of the strategy portfolio): extend the finite Hecke functor of Gate 0 ([verification/v733\_strat3\_gate0\_census.py], 13/13, GATE0-WEAK-FUNCTOR — exact on the periodic-orbit subalgebra, 6859 convolutions at deviation 0,  $\kappa \equiv 1$ ) to a positive global KMS state  $\omega$  at  $\beta = 1$  on the closed groupoid algebra, with embeddings  $\iota_h: A_h \rightarrow A_{\text{ES}}$  and  $T_h = \omega \circ \iota_h$  — then Weil positivity follows for the dense test space, and with it RH. The five separation steps, each individually named and decidable: (1) algebraic extension from the periodic-orbit subalgebra to the full convolution algebra; (2) topological closability; (3) complete positivity; (4) the KMS condition at  $\beta = 1$ ; (5) the Weil-limit identification. Kill criterion: complete positivity fails already on a finite extension — one finite counterexample retires the route. The contract names the gap; it does not close it. A research contract, not a claim; no marker move; no RH statement. [2026-08-04 update: the strong form is DEAD and the contract is refined. [verification/v740\_kms\_extension\_switch.py] (8/8, KMS-EXTENSION-DEAD):  $\beta = 1$  KMS is algebraically impossible on Gate-0's nonzero-grade Laurent unitaries —  $T(u_p u_p^*) = 1$  versus  $T(u_p^* \sigma_i(u_p)) = 1/p$ , exact defects 1/2, 2/3, 4/5 at  $p = 2, 3, 5$  — and exact covariant  $*$ -restrictions do not exist (0 of 141..1255 required boundaries); local extended Gram positivity survives (55,296 monomials/window). [verification/v741\_kms\_toeplitz\_semigroup.py] (9/9, KMS-TOEPLITZ-DEAD): the Bost–Connes semigroup Toeplitz correction repairs the algebraic KMS step completely (boundary-corrected identities exact; normalized bare error  $10^{-2} \rightarrow 10^{-9}$  along the directed system), but the covariant inductive step stays dead (the preregistered UCP window covariance is nonmonotone). Separation steps (1) and (4) are thereby settled at the finite level; the surviving demand is restated below as PRIME.KMS.INDUCTIVE\_STATE.02.]

**The refined master contract: operator-system compactness (PRIME.KMS.INDUCTIVE\_STATE.02).** OPEN (registered 2026-08-04, refined from INDUCTIVE\_STATE.01 after [verification/v740\_kms\_extension\_switch.py] and [verification/v741\_kms\_toeplitz\_semigroup.py]): keep the proven local blocks — the BC semigroup Toeplitz algebra per window (nonunitary isometries with boundary projections), the boundary-corrected  $\beta = 1$  KMS identities (exact), extended Gram positivity on the factorized slices, and UCP state transport ([verification/v735\_strat3\_ucp\_inductive.py]) — and state the single remaining global step purely as an operator-system compactness/subsequence problem: does the net of window states have a weak- $*$  cluster point that is a positive KMS functional on the closed algebra, without demanding a covariant inductive system? Kill criterion: a finite obstruction to any cluster point (a uniform positivity violation along a subnet). Do not claim a covariant inductive system until a new canonical window map is found. Literature anchor:  $\beta = 1$  KMS for the genuine Bost–Connes algebra is classical (Bost–Connes 1995); the TFPT burden is the compatibility of the glued window states with one global limit, not local KMS existence. A research contract, not a claim; no marker move; no RH statement. [2026-08-04 update: step (3) — complete positivity — is SOLVED on the factorized ramified 15-label Toeplitz layer. [verification/v756\_kms\_incidence\_stinespring.py] (7/7, KMS-INCIDENCE-CP-COVARIANT): 105 explicit Kraus terms realize the incidence channel  $K = B/7$  as a completely positive unital map, with all four compatibilities exact —  $\sigma$ -covariance

at deviation 0, half-weight covariance  $DU = 2^{-1/2}UD$  to  $1.1 \times 10^{-16}$ ,  $\beta = 1$  detailed balance to  $6.9 \times 10^{-18}$ , and Gate-0 compatibility to  $3.3 \times 10^{-16}$ . The local candidate state is  $\tau_{\text{inc}} \otimes \tau_{\text{Toeplitz}}$  (channel-preserved). The remainder to a global state is exactly the compactness step of this contract: compatible incidence channels at every finite place and level, the interaction with non-label ramified multipliers and the arch lift, a canonical UCP map between the nine window systems, compactness/normality of the limit, and the Weil-functional identification.] [2026-08-04 update (the diagonal gram round): the compactness step now carries MEASURED handoff evidence and an abstract Lean closure. Measured: the  $\varepsilon$ -regulated resolvent is the admissible operator-system evaluation with  $\varepsilon$ -robust falling local defects ([`verification/v766_handoff_bulk.py`], 17/17, HANDOFF-BULK-CONVERGES); a target-independent growing-frequency PSD source Gram passes the finite ladder at slope  $-0.369$  ([`verification/v767_handoff_frequency_gram.py`], 6/6, HANDOFF-GRAM-CONVERGES); the quadrature tail terminates algebraically ( $2.4 \times 10^{-13}$ , uniform in the window), cross-window compatibility is closed at  $\varepsilon = 10^{-3}$  on the deepest reachable ladder, and the source limit is identified with the deployed Weil functional on the frozen battery within 0.0485 ([`verification/v768_handoff_tail_weil.py`], TAIL-WEIL-PARTIAL+COMPAT-EPS3-CONVERGES). Abstract: the finite-stage Gram compactness/diagonal-extraction skeleton is machine-checked in Lean (`TfptCarrier/GramCompactness.lean`, 8 theorems). THE SURVIVING OPEN OBJECT is the  $q_f$  wall of [`verification/v763_yosida_handoff.py`] (QF-DEAD): the battery near-kernel pairing is non-Cauchy at honest maximal depth and survives orthogonalization by structural subordination — the weak-\* cluster question of this contract now lives in exactly that bounded quantity; the round's theorem form is registered as **PRIME.GRAM.DIAGONAL.01**. No marker move; no RH statement.] [2026-08-05 follow-up: the  $q_f$  object is adjudicated —  $q_f$  is representation gauge and the corner-native Gate 3' is dead ([`verification/v769_qf_contract_necessity.py`]), the levels settle positive ([`verification/v770_qf_spectral_bundle.py`]), fixed-d Feshbach closes at a moving spectral edge ([`verification/v772_qf_feshbach_effective.py`]), and the cell cocycle is domain-only ([`verification/v773_qf_cell_cocycle.py`]): **PRIME.GRAM.DIAGONAL.01** is CLOSED per its own frozen stakes, and the compactness remainder of this contract hands over to the sharpened **PRIME.Z1.OPERATOR.01** (the Weyl M-function / boundary-threshold form with the measured  $q_f$  constraints). No marker move; no RH statement.] [2026-08-05, the unification measured ([`verification/v780_z1_compactness_trilogy.py`]): the weak-\* cluster demand of this contract and the operator demand of **PRIME.Z1.OPERATOR.01** are measured to be the same remaining object on the reachable surface. The cluster-existence half is carried: the variable-edge boundary family is bounded (0.936) and equicontinuous (0.807) on the frozen compacts with summable  $\sim$ rank-one entry poles ( $c_{\text{mass}}$  0.032) — every subsequence has weak-\*/locally-uniform cluster points (Helly/Montel at finite level). The selection half is not carried (moment-functional tail oscillation 22.1), and its obstruction is the same measured object that fails the delivery legs: the lossy  $K = M/2$  half-window import. State selection (this contract) = import-faithful boundary map (**OPERATOR.01**): one object viewed twice; the merged target is named Z1-COMPACTNESS (**PRIME.Z1.COMPACTNESS.01**); both rows stay, cross-referenced. No marker move; no RH statement.] [2026-08-05 add-on: the Z1-COMPACTNESS frame is now also the named limit frame of demand (2) of the new **PRIME.POSITIVE\_DESCENT.01** contract below — the GL1-sector PSD persistence of the manifestly positive packet object is to be decided in exactly this weak-\* cluster frame.] [2026-08-06 add-on: the CP extension gate ([`verification/v794_cp_extension_gate.py`], CPGATE-UNDECIDED-FALLING) measures this demand sector-decorated: the  $\chi = \text{GL1}$  instance of the four-sector window net is this contract's object (anchor  $3.882 \times 10^{-6}$  at  $M = 1176$  reproduced on the same  $10^8$  comb), “state selection must settle” = “the sector margins never cross”; measured: GL1 falls power-law ( $\beta = -1.69$ ),  $\chi_4$  saturates at  $\sim 1.5 \times 10^{-6}$  — the first measured sector floor.] [2026-08-06 round-23 note: the SELECTION half is now finite-level solved — Mosco form convergence + Friedrichs minimality on the v762 dense core replace the failed moment selector on identical rungs (moment oscillation 22.08 vs Friedrichs resolvents 0.0004 / ledger-adjusted Weyl 0.0843 / raw  $\mu$ -face 0.0000; cofinal-unique  $2.5 \times 10^{-5}$ ), and the import diagnosis is completed at form level: nothing is whitened at the core ([`verification/v816_prime_mosco_selection.py`], MOSCO-SELECTS;

the new contract *PRIME.MOSCO.SELECTION.01* below). The remaining analytic object of this contract is the sector floor — restated as *PRIME.FLOOR.RATIO.01* below.] [2026-08-06 round-24 note: the floor contract is narrowed — the lock-block floor  $\lambda\tau$  now carries a complete certified three-piece lower-bound skeleton on the deployed battery, closed at citation grade for all  $h$  up to  $\alpha^* \approx 11$  ([*verification/v823\_prime\_lagrange\_floor.py*], [*verification/v824\_prime\_floor\_skeleton.py*]); battery-relative and  $\alpha$ -bounded, still open.]

**The positive descent (PRIME.POSITIVE\_DESCENT.01).** *OPEN* (registered 2026-08-05, from the positive-protocol round; finite-level facts machine-checked in [*verification/v791\_positive\_descent.py*], 30/30 incl. the five must-fire controls, verdict DESCENT-PARTIAL; the declared run-1  $\rightarrow$  run-2 calibration — the falsified linear state F1, the falsified all-PSD operator expectation F2, the corrected C3 constant — is carried verbatim in the module provenance). *The finite-level facts:* (i) the packet GNS state on the extended positive register  $\mathbb{N}[C_2] \otimes \mathbb{N}[\mathbb{F}_2^4] \otimes \mathbb{N}[\mu_4]$  is positive at every depth *by construction* (a sum of event squares; verified in exact rationals; the falling minimal sector is  $f_8^2$ -weight dilution, structurally unable to cross zero), while the naive *linear* pushforward of the same nonnegative data is *not* a state (F1, quantified:  $-6.6 \times 10^{-2}$ , depth-flat) — the squared/GNS form is the only positive aggregation found, and signedness is *observer projection*; (ii) the GL1 character sector of the packet correspondence operator is *bit-identical* to the deployed Weil window (lag identity at  $6 \times 10^{-16}$ ) and is its *unique* PSD sector of all 24 (margins  $+5.3 \times 10^{-5} \rightarrow +1.2 \times 10^{-5}$  over  $X = 4 \rightarrow 10$ , reproduced read-only); (iii) every damped/twisted sector breaks because the register-trivial continuum cannot follow the atom leg (the pole-atom cancellation is exposed; localization: arch-only flat, arch+pole falling) — the breaking tensor factor is the *continuum* ( $e$ -slot) leg, not  $C_2$ ,  $\mathbb{F}_2^4$  or  $\mu_4$ . *The contract therefore demands three objects:* (1) **sector-adapted continua** (the missing CP-functor data): for each character sector  $\chi$  the arch/pole layer of the *twisted* channel (the explicit-formula continuum of the  $\chi$ -twisted  $L$ -data; for the  $f_8$  sector: a weight-4 archimedean factor, *no* pole), such that the sector window  $T_{\text{cont}}^\chi + \sum_n \text{Re } \chi(x_n) Q_n$  is PSD at finite level — *finite-level falsifier*: exhibiting the  $f_8$ -sector continuum and finding its window non-PSD kills the operator face for good; constructing it PSD promotes the functor from state level to operator level. (2) **sector-PSD persistence** (the GL1 limit): the weak-\* limit — in the Z1-COMPACTNESS frame of PRIME.Z1.COMPACTNESS.01 / [*verification/v780\_z1\_compactness\_trilogy.py*] — of the GL1 sector functional stays in the cone as  $X \rightarrow \infty$ . This *is* Weil positivity, now typed as: “the sign sector of a manifestly positive packet object stays PSD” — the sign is never controlled directly, only the sector compression. (3) **the carrier intertwiner**: a positive identification, in the GNS limit, of the population register ( $\chi_-$  readout = the damped Hecke ratio  $a_n/\sigma_3(n)$ ) with the carrier register ( $\chi_-$  readout = the unit-modulus sign); at finite level this step is *not* CP-trivial (the signed count fails positivity) — the renormalization/amplification limit is the single non-automatic step. *Stop conditions:* no fixed- $d$  variants, no re-gating of the closed diagonal-Gram objects, no RH claim; demand (2) failing at any finite depth = DESCENT-DEAD for the route; demands (1) and (3) without (2) are decorative. *Inherited finite-level falsifier:* any rung where a non-GL1 carrier sector becomes PSD-binding while GL1 stays positive breaks the “GL1 is the unique carrier edge” reading measured here. A research contract, *not* a claim; no marker move; no RH statement. [2026-08-06 update, markers unchanged — the contract stays open for the floor: all three demanded objects are now delivered at finite level. (1-F) The demand sharpens to the continuum functor:  $\chi \mapsto \Lambda_\chi = q(\chi)^{s/2} \prod_j \Gamma_{\mathbb{R}}(s + \mu_j(\chi)) L_{\text{an}}$  with  $(\mu, q, \text{pole})$  given by one closed rule; verified instances GL1 ( $q=1$ , pole),  $\chi_4$  ( $q=4$ ),  $f_8$  ( $q=8$ ),  $f_8 \otimes \chi_{-4}$  ( $q=16$ , Atkin-Li + Fricke), the conductor datum load-bearing at exactly  $\ln 2$  and pinned from below by the PSD margins; the mirror-type sectors are structurally non-automorphic and carry no positivity demand ([*verification/v793\_f8sector\_conductor.py*], 15/15 + 16/16). (2-L/2-Z) The limit object acquires its exact form: per automorphic character the positivity of the sector Weil functional  $W_\chi$  on the union test space  $V_\infty$  —  $\inf_X \lambda_{\min}(T_{\chi,X}) \geq 0$  — with the net data a CP-extension theorem consumes measured (restriction exactness, interlacing, uniform bounds), and demand (2) at  $\chi = \text{GL1}$  identified with the Z1-COMPACTNESS obstruction, sector-decorated; measured trend: no crossing, GL1/ $f_8$ /twist power-law falling,  $\chi_4$  saturating at the first measured

sector floor ([verification/v794\_cp\_extension\_gate.py], 17/17). (3) The carrier intertwiner exists at finite level in Stinespring form with exact Choi and all covariances, its four automorphic channels landing on the four sector windows from one CP map, and the F1 failure retyped as a negative Choi eigenvalue ([verification/v801\_prime\_cp\_intertwiner.py], 29/29; the infinite-level demands live in PRIME.CP.INTERWINER.01 below). The remaining analytic core, restated: the per-sector positivity floor (GRH-type, unproven) plus the intertwiner demands D1–D3. No RH statement.] [2026-08-06 round-23 note: selection and uniqueness are now both finite-level delivered — selection by Mosco+Friedrichs ([verification/v816\_prime\_mosco\_selection.py], PRIME.MOSCO.SELECTION.01) and the unique positive limit by the master theorem ([verification/v817\_positive\_descent\_master.py], TFPT.POSITIVE\_DESCENT.MASTER.01: prime packet state defects summable at 0.701/doubling,  $\hat{m}_2(p) = -1/15$  exact for every odd prime; positivity a separate hypothesis by control  $K_4$ ). The floor itself reduces to the single ratio inequality  $\rho = \tau/\tau_{\text{pnt}} > 0$  with a measured  $h^{-3/2}$  envelope — the per-sector positivity floor is the one remaining analytic object, restated as PRIME.FLOOR.RATIO.01 below.] [2026-08-06 round-24 note: that remaining object is narrowed — exact Lagrange sum-of-squares identity, the certified fixed pair, the analytic fixed-pair bound, the family exhaustion, and the deep tail closed at citation grade for all  $h$  on the deployed battery ([verification/v823\_prime\_lagrange\_floor.py], [verification/v824\_prime\_floor\_skeleton.py]; the narrowed contract text below); battery-relative,  $\alpha$ -bounded, still open.]

**The carrier intertwiner (PRIME.CP.INTERWINER.01).** OPEN (registered 2026-08-06; demand (3) of PRIME.POSITIVE\_DESCENT.01 solved at finite level in [verification/v801\_prime\_cp\_intertwiner.py], 29/29, CP-INTERWINER-EXISTS). *Object:* the finite-level carrier intertwiner  $\Phi(a) = V^*\pi(a)V$  on the arrow  $*$ -algebra  $M_{15} \otimes \mathbb{C}[\text{tower}] \otimes (\text{packet slots at the odd places})$ ,  $V$  the 105-leg Kraus stack (unitality  $\sum_e E_{x_e x_e} = 7I$  exact), involution = arrow reversal, KMS half-weights  $7^{-1/2}$  per leg /  $n^{-1/2}$  universal. *Certified at finite level:* CP/unitality (exact Choi, rank 105, eigenvalue  $1/7$ );  $\sigma/\text{deck}/\text{half-weight}/\beta = 1$  covariances exact; the Choi census on the generating family exact (integer Gram factorization); the four automorphic channels of the one map land on the deployed GL1 Weil window (machine precision) and the  $\chi_4/f_8/\text{twist } \Gamma_{\mathbb{R}}$ -rule windows (PSD on  $X = 4..10$ ); the ramified conductor emptiness structural. *Demands (open):* D1 the  $X$ -compatible dilation family and its inductive limit (with the v794 margin trends as the honest obstruction data); D2 normal extension on the Z1-COMPACTNESS limit object (v780); D3 the limit identification of the GL1 channel with the critical-line Weil functional — which contains RH and is not claimed. [2026-08-06 round-22 note, D1–D3 unchanged: the intertwiner gains its register chart at the odd places — the affine Gray map, exactly deck-covariant with trivial  $\mu_4$ -cocycle ([verification/v804\_prime\_carrier\_gray.py]) — and its 105 Kraus legs are a complete context protocol over the doily with the exact commutant  $[B_{45}, B_{60}] = 0$  ([verification/v811\_prime\_kraus\_doily.py]).] *Kill:* a finite positive packet whose required GL1 intertwiner fails the exact Choi census structurally. No marker move; no RH statement.

**The Gray register chart (PRIME.CARRIER.GRAY.01).** Registered 2026-08-06 ([verification/v804\_prime\_carrier\_gray.py], 23/23, GRAY-INTERWINER-EXISTS; one declared control repair carried verbatim). *Object:* the affine Gray encoding of  $C_2 \times \mu_4$  into the R-class torsor, exactly respecting the v786 Packet480 kill (a group isomorphism is impossible — exponent 4 vs 2 — and not demanded). *Certified:* deck covariance exact with trivial measured  $\mu_4$ -cocycle; the residual a pure  $\mathbb{Z}_3$  line twist on the  $\sigma$ -fixed classes; the extended Stinespring GL1 channel re-lands on the deployed Weil window at rel  $6.0 \times 10^{-16}$ . *Demands:* D1 odd-place fiber charts; D2 the  $\mathbb{Z}_3$  line twist as crossed-product structure; D3 the infinite-level CP extension (chained to PRIME.CP.INTERWINER.01). *Kill:* a nontrivial cocycle under the frozen chart. No marker move; no RH statement.

**The Kraus context protocol (PRIME.KRAUS.DOILY.01).** Registered 2026-08-06 ([verification/v811\_prime\_kraus\_doily.py], 21/21, KRAUS-DOILY-PROTOCOL; one declared report-text repair carried verbatim). *Certified:* the 105 Kraus legs = 15 diagonal + 90 off-diagonal over the doily  $W(3, 2)$ ; the canonical  $45 + 60 = \text{reversal quotient } (2:1 \text{ over the } 45 \text{ flags}) + \text{fiber}$

expansion ( $15 \times 4 = 60 =$  the **v783** stabilizer rays); the  $Sp(4, 2)$  obstruction certifies no leg-subset partition (the literal partition is spread-relative = the ledger T3 census); the 45 transports quadratic ( $P_x P_y = \pm P_{x+y}$ ), spanning  $\text{End}(\mathbb{C}^4)$  at rank 16; the surprise  $[B_{45}, B_{60}] = 0$  exactly —  $K = \frac{3}{7}E + \frac{4}{7}R$  a direct-sum protocol with  $E$  a conditional expectation. *Measured*: the certified spread is *not*  $\sigma$ -invariant. *Named falsifier*: a  $\sigma$ -invariant spread reproducing the commutant split. [2026-08-06 round-23 addenda ([[verification/v815\\_kraus\\_spread\\_commutant.py](#)]): the falsifier pair is decided —  $\sigma$  acts on the six spreads as two 3-cycles with zero fixed spreads (outer-class forced; no invariant coarsening), but the direct-sum protocol is  $\sigma$ -stable as a family ( $\sum_s B_{45}(s) = 4I + 2B$ ; per-orbit  $A_{\text{orb}}$  spectrum  $\{0^5, 1^2, 4^5, 7^2, 9^1\}$ );  $[E, K] = 0$  extends exactly to channel/KMS-tower/odd places (the unique  $C_2$ -character expectation per place) and the 600-dim register; the  $\{K, \sigma\}$  commutant is 39-dim nonabelian with  $A_{\text{orb}}$  its canonical abelian 5-dim subalgebra (**PRIME.AORB.REFINEMENT.01**), window-silent beyond the  $K$ -spectrum; and the rule-A protocol set is canonically the 105 minimal dual checks of  $[15, 10, 4]$  ([[verification/v820\\_prime\\_kraus\\_rm24.py](#)], **PRIME.KRAUS.RM24.01**).] No marker move; no RH statement.

**The Lorentz spinor coordinates (PRIME.LORENTZ.SPINOR.01).** *Registered* 2026-08-06 ([[verification/v807\\_lorentz\\_nullselector.py](#)], 22 checks with the three honest  $h$ -order FAILs encoded, **LORENTZ-COORDS-REVEAL**). *Exact*:  $M(5, -3, 4) = 2(1, 2)^T(1, 2)$  — the compiler triple is a rank-one spinor square, Euclid  $(2, -1) \mapsto (5, -3, 4)$ . *The reveal*: the locking slope is strictly monotone in the window width  $\alpha$  (66/66, Kendall 1.000, cross-ratio contraction CV 0.004) — the  $h$ -oscillation was  $\alpha$ -jitter; the approach  $2 - r \sim \alpha^{-0.23}$  glacial. *Named candidate*: the 1D monotonicity statement in  $\alpha$ -order (the three  $h$ -order gates stay honest FAILs). *Kill*: a Kendall inversion under the frozen windows. No marker move; no RH statement.

**Mosco selection (PRIME.MOSCO.SELECTION.01).** *Registered* 2026-08-06 ([[verification/v816\\_prime\\_mosco\\_selection.py](#)], 23/23 guards+controls, 6/6 legs, **MOSCO-SELECTS**). *Statement*: selection by Mosco form convergence and Friedrichs minimality on the **v762** dense core — the window-form chain  $q_X$  has exact nested recovery (structural,  $1.2 \times 10^{-12}$ ), vanishing liminf defect (measured; the moving-edge family confirmed a genuine boundary channel, not an escape), Cauchy Friedrichs resolvents and Weyl functions *modulo the typed summable entry-pole ledger*, and cofinal-unique limits ( $2.5 \times 10^{-5}$ ): the unified compactness contract selects its state by form convergence, not by moment stability (moment baseline 22.08 reproduced exactly in-run). *Completed*: the import diagnosis at form level — the raw  $\mu$ -weighted coupling of the core is exactly rung-independent, the whitened control reproduces the 5.8772 delivery failure exactly. *Remaining for theorem grade (typed)*: (i) Mosco-precompactness of  $\{q_X\}$  on  $V_\infty$ , (ii) identification of the limit form domain, (iii) entry-pole ledger summability at infinity. Positivity/sector floor stays a separate contract. No marker move; no RH statement.

**The positive-descent master theorem (TFPT.POSITIVE\_DESCENT.MASTER.01).** *Registered* 2026-08-06 ([[verification/v817\\_positive\\_descent\\_master.py](#)], 14/14 incl. the four must-fire controls, **MASTER-BOTH-INSTANCES**; kernel-checked core `TfptCarrier/PositiveDescentMaster.lean`, the matrix positivity closure in `GramCompactness.lean`, cited). *The theorem (seven frozen hypotheses)*: finite operator systems with unital inclusions, positive states, CP transitions, finite character sectors, *summable* state defects, a sector-adapted closure datum, and a Mosco form family force a unique positive limit state with well-defined positive limit sectors; the core (telescoping + quantitative Cauchy tail; positivity closed; sectors commute) is kernel-checked. *Instance  $G_{\text{net}}$* : hypotheses (1)–(5) measured (defects 0.267/doubling; the expectation CP via Choi;  $\mu_4$  evenness exact); open: normality/split. *Instance prime/package* — *the new number*: state defects summable at 0.701/doubling  $\approx 2^{-1/2}$  (the dilution law), with  $\hat{m}_2(p) = -1/15$  exact for every odd prime; open: the sector floor (**PRIME.FLOOR.RATIO.01**, narrowed in round 24 by the certified floor skeleton **v823/v824** — battery-relative,  $\alpha$ -bounded, still open). *Control  $K_4$* : summable defects with one non-positivity-preserving step give a non-positive limit — positivity is a separate hypothesis. The selection problem is eliminated on both sides. No marker move; no RH statement.

**The floor as one ratio inequality (PRIME.FLOOR.RATIO.01).** *OPEN* (registered 2026-08-06 — the reduced sector-floor contract; evidence [[verification/v818\\_sector\\_floor\\_attack.py](#)], 9/13 + 13/15

with exactly the preregistered honest FAILs, FLOOR-MECHANISM-PARTIAL + CAPTURE-BOUNDED / INGREDIENT-II-THEOREM-SHAPE / AMPLIFIER-MECHANISM). *Statement:* the sector positivity floor  $\inf_X \lambda_{\min}(T_{\text{GLI},X}) \geq 0$  reduces, up to the measured  $O(1)$  interlacing capture (median 1.53, range 1.04–4.36; the 2-mode lock block an  $O(1)$ -faithful witness by Cauchy interlacing), to the *single* ratio inequality  $\rho(X) = \tau(X)/\tau_{\text{pnt}}(X) > 0$  — the prime comb against its own smooth density — with the measured finite lower envelope  $\rho \geq ch^{-3/2}$  (constant  $c \approx 4.85$ , envelope non-decaying, slope +0.14). *Certified/measured ingredients:* the direction  $r(X)$  is owned by the density rotation law  $dr/dt = -\kappa(w \cdot v_1)(w \cdot v_2)(1 + r^2)/(\lambda_1 - \lambda_2)$  (symbolic, sign hypothesis verified finite-level, 100% staircase sign-match, typed density-level by the Epstein control  $s_E = 0.241$ ); the capture is angle-certified frame geometry ( $\cos \theta$  median 0.990, scramble-robust — frame content, not arithmetic); the amplifier  $\rho$  is carried by the smallest prime powers (packet correlation +0.964, partially trend-driven, typed). *Kill:* the envelope failing at depth ( $\rho h^{3/2}$  decaying below every positive constant) or the capture angle collapsing ( $\cos^2 \theta < \frac{1}{2}$  at any rung). Selection and uniqueness are eliminated on both faces (PRIME.MOSCO.SELECTION.01, TFPT.POSITIVE\_DESCENT.MASTER.01) — the floor is the single remaining analytic object of the prime programme. Interlacing runs floor  $\Rightarrow$  margin (necessary, not sufficient); promotion of any  $\rho$  bound is *fenced* on the preregistered envelope + capture gates. [2026-08-06 round-24 narrowing — *what is now closed:* on the deployed 14-rung battery the lock-block floor  $\det \hat{A}_2 = \lambda \tau$  has a complete three-piece uniform certified lower-bound skeleton — the *exact* Lagrange sum-of-squares identity over zero+pole rank-one carriers (machine wards  $\leq 1.4 \times 10^{-9}$ ), the analytic fixed-pair bound  $X_\infty^2(\alpha)(1 - C(\alpha)/h)$  with explicit  $h_0(\alpha) \leq 500$  ( $h$ -uniform at every deployed  $\alpha$  in-domain), the certified top-100 family exhaustion (0.93–0.97 of the floor, all carriers on-line  $\leq 2 \times 10^4$ ), and the psd remainder closed at *citation* grade (product-sup envelope + Abel/RvM explicit constants;  $T_{\text{ver}} = 3 \times 10^{12}$ ) for all  $h$  at the deployed  $\alpha$ , with an  $h$ -free growth law ([verification/v823\_prime\_lagrange\_floor.py], [verification/v824\_prime\_floor\_skeleton.py]; PRIME.FLOOR.LAGRANGE.01 / PRIME.FLOOR.SKELETON.01 below). *What remains open — the contract:* the floor bound is *battery-relative* and  $\alpha$ -bounded — (i) the citation-grade chain re-crosses  $X_\infty^2$  at the validity horizon  $\alpha^* \approx 11.2$  (deeper  $\alpha$  needs verified on-line *depth* beyond  $T_{\text{ver}}$ , never location); (ii) the fixed pair dies at the  $\alpha$ -nodes of  $\sin(\alpha\gamma_1)$ , where the family limit (positive, minimum  $1.62 \times 10^{-4}$ ) must take over — the family version of the analytic bound is not yet written; (iii) the product-sup Lipschitz certificate is a named elementary step, typed, not claimed; (iv) *nothing bounds*  $\inf_X \lambda_{\min}$  on  $V_\infty$  — the  $O(1)$  capture and the  $\rho > 0$  demand at infinity remain exactly this contract. Kill criteria unchanged.] [2026-08-07 round-25 depth-kill + budget note: the contract's *own* kill gates were played at full sieve depth  $X = 18.375$ –25.5 (comb caps to  $1.2 \times 10^{11}$ ; the deep uniform-grid frames decouple depth from dimension) and *survived* — envelope margins  $\rho h^{3/2}/c = 5.76$ –8.92 growing, capture minimum  $\cos^2 \theta = 0.849$ , non-decay trend +0.060 per log-unit: the  $h$ -law wins over any hidden depth law ([verification/v829\_prime\_floor\_depth.py]); the single fixed pair collapses at depth (gap  $40\times \rightarrow 673\times$ , the typed S5.HLAW miss) and the certified top-100 family takes over (0.977–0.981 of the floor,  $\rho_{\text{certfam}}$  within 2% of measured). The round-24  $\alpha^* \approx 11$  citation horizon is *removed*: it was an artifact of the frozen quadratic float convention — under the explicit Higham-linear budget the citation-grade family gate closes 6/6 to  $\alpha = 12.75$ , full sieve depth ([verification/v830\_prime\_float\_budget.py]; the deployed v818 convention stays frozen in the suite, the linear budget is convention for *new* modules only). The envelope constant is now *certified-explicit*:  $\rho \geq \rho_{\text{certfam}} \geq 4.335 h^{-3/2}$  on 73/73 battery points, non-decaying, with the family closed form exact ( $3.6 \times 10^{-11}$ ) and phase conspiracy excluded (dip floor 0.170). The remaining proof blocker is *named*: the alias phase correlations (random-phase predicts  $\tau \sim h^{-1}$ ; the  $D = 1/64$  tower measures  $h^{-2.5}$ ) block the analytic uniformity — the envelope stands certified-explicit, not proven. Kill criteria unchanged; still [O].] [2026-08-07 round-26 closure note: the analytic-envelope thread of this contract is *closed at its typed circularity boundary* ([verification/v831\_prime\_alias\_second\_moment.py], ALIAS-CORRELATION-PAIRCORR, the seven preregistered-honest premise-overturning FAILs pattern-gated). The round-25 blocker resolves into the boundary itself: the premise “the phases carry the gap” is *overturned* — the amplitudes carry it

( $-2.505 = -3.172 + 0.668$  exactly;  $q = S_c/S_2$  sign-changing, deep suppression only  $\times 1.1$ ) — and the exact Guinand split puts the coherent sum in the prime-comb fluctuation at  $0.01 \times$  the Poisson-comb square-root scale (smooth  $-994$  / fluctuation  $+995$ , cancelling; the comb scramble explodes it  $2 \times 10^6$ – $1.7 \times 10^7$ ). By Guinand the comb’s self-cancellation *is* the zero-side floor statement itself — the required control is even deeper than Montgomery-type pair-correlation input. *No further analytic-envelope variants* (stop-list entry per house discipline); the envelope’s constant is arithmetic substance equivalent to the target; the demand moves to PRIME.FLOOR.PAIRCORR.01 below. Kill criteria unchanged; still [O].] [2026-08-07 round-29 margin-law note: the envelope constant is now doubly derived ([verification/v843\_margin\_law\_excess\_lean.py], MARGIN-RECURSES-THE-WALL): the identified-corner (relational) coordinates reproduce this contract’s floor-strand law exactly —  $e_1 = (\tau/\tau_{\text{pnt}}) h^{3/2} \in [4.855, 24.209]$  on all 67 rungs,  $\min \geq 4.335 = c_{\text{cert}}$ , with  $\tau$  agreeing with  $\lambda_{\min}(A_h)$  at  $2.1 \times 10^{-12}$  — two independently derived laws, one object: the consistency itself is a nontrivial ward. The typed law  $\tau(\alpha, h) = e_1(\alpha) h^{-3/2} \tau_{\text{pnt}}(\alpha)$  puts the entire  $\alpha$ -decay in  $\tau_{\text{pnt}}$  and the arithmetic in  $e_1$ ; the fit-free internal  $h$ -exponent measures median  $-1.341$  against the frozen bar edge  $-1.35$  — the  $h^{-3/2}$  law confirmed at  $\sim 1\%$  grade, the 0.009 miss kept as the one preregistered-honest FAIL (not refit). The margin-law attack class (tower recursion / per-prime factorization of the excess) is measured dead and stop-listed per the cell-level self-similarity typed in PRIME.RELATION.SKELETON.01. Kill criteria unchanged; still [O].] No marker move; NO RH claim.

**The floor as an exact sum of squares (PRIME.FLOOR.LAGRANGE.01).** *Registered* 2026-08-06 ([verification/v823\_prime\_lagrange\_floor.py], 7/8 with the one preregistered-honest old-budget FAIL + 9/9, LAGRANGE-CONCENTRATED + PAIR-CERTIFIED; discovery probes prime\_lagrange\_pair\_probe.py and prime\_lagrange\_budget\_probe.py). *Certified (identity-grade wards)*: every per-zero layer of the v692 master identity is exactly rank-one ( $v_\gamma = 2\sqrt{w(\gamma)}(S_1, S_2)$ , closed trigonometric reads, layer deviation  $\leq 1.2 \times 10^{-16}$ ), and  $\det(G_Z + P)$  is an exact sum of squares over the component pairs (Lagrange; ward  $\leq 1.4 \times 10^{-9}$  on all 14 rungs); the pole is the universal non-collinear leg and the moving zero leg sits on the alias comb  $\gamma \approx 2\pi k/D$ . *Measured (per-rung, float-level)*:  $\lambda_{\min}(\hat{A}_2 - G_Z - P) \geq$  margin verified per rung eliminates the tail estimate — budget tightened a median 9.8 orders — and the certified interval  $X^2(\text{pole}, \gamma_1) - \text{BUD}$  is strictly positive on 14/14 rungs, with the floor share growing as  $h^{+1.06}$ . *Honest typing*: the atom-side rank-one framing fails as frozen — the SOS realization lives on the zero+pole side; the psd of the remainder is verified numerically here, its citation-grade upgrade is PRIME.FLOOR.SKELETON.01. *Kill*: any ward failing on the frozen ladder. Feeds PRIME.FLOOR.RATIO.01 (narrowed). No marker move; no RH statement.

**The floor-theorem skeleton (PRIME.FLOOR.SKELETON.01).** *Registered* 2026-08-06 ([verification/v824\_prime\_floor\_skeleton.py], 12/12 + 9/9, FLOOR-PARTIAL + TAIL-CLOSED-ALL-H; discovery probes prime\_floor\_theorem\_probe.py and prime\_tail\_envelope\_probe.py). *Uniform-verified on the deployed battery*: the analytic fixed-pair limit  $X_\infty(\alpha) = 16\alpha\pi^2 \sin(\alpha\gamma_1) \sinh(\alpha/2)$ [bracket] with the explicit bound  $L(h) = X_\infty^2(\alpha)(1 - C(\alpha)/h)$  for  $h > h_0(\alpha) = C(\alpha) \leq 499.6$  (zero sign flips; 25  $\alpha$ -nodes typed —  $h$ -uniform per  $\alpha$ , not ladder-uniform in  $\alpha$ ); the certified top-100 family exhaustion (median 0.952 of the floor, carriers on-line  $\leq 2 \times 10^4$ ; residue 0.015–0.049, non-growing; the analytic family limit positive across all nodes). *The tail closure*: the product-sup per-zero envelope + Abel/RvM with explicit unconditional Trudgian-grade constants close the psd remainder at citation grade on *all 14 rungs* at the fixed verified horizon  $T_{\text{ver}} = 3 \times 10^{12}$  with margins  $10^3$ – $10^6$ ; the growth law is  $h$ -free ( $+1.98 \rightarrow +0.13$ ), the remaining growth  $\alpha$ -only, validity horizon  $\alpha^* \approx 11.2$  (the Vinogradov–Korobov region typed honestly useless). *Net*:  $\lambda\tau \geq X^2(\text{pole}, \gamma_1) - \text{pert}(h) > 0$  at citation grade — no per-rung eigencheck, no zero-location input beyond the critical strip. *Honest boundaries*: battery-relative,  $\alpha$ -bounded, necessary-side; the sup Lipschitz certificate a named elementary step; nothing bounds  $\inf_X \lambda_{\min}$  on  $V_\infty$ . *Kill*: any closure margin failing on the frozen battery at the frozen horizon. Feeds PRIME.FLOOR.RATIO.01 (narrowed). No marker move; no RH statement.

**The comb-native window detector (PRIME.DETECTOR.WINDOW.01).** *OPEN* (registered

2026-08-07 — the window-verification contract of the exclusion-detector strand; evidence [verification/v825\_prime\_exclusion\_ladder.py], [verification/v826\_prime\_exclusion\_battery2.py], [verification/v827\_prime\_zero\_locator.py], [verification/v828\_prime\_comb\_window.py]]. *Statement*: the verified PSD rungs of the GL1 prime-comb tower, inverted through the tent-read Guinand identity on a hash-preregistered battery (no zero datum in any design, warded end to end), form a *window detector*: on a window inside the Nyquist reach  $\pi/D$  the instrument *locates* zeta ordinates as prominent local maxima of the uncensored margin-break profile (out-of-sample validated: 83% detection, 0% false positives, precision 0.086 on the disjoint window (60, 120]), *excludes* all quadruple hypotheses with  $\delta \geq \delta_{\text{mb}}(\gamma)$  (witness-Rayleigh-certified on Cholesky/Higham-certified rungs, deepest  $X = 25.5$ ), and *reconciles* the census against the cached verified count and the rounded RvM main term (COMB-VERIFIES-WINDOW). *Open*: (a) the depth-to-width law  $\Xi_{v2}(X)$  (measured slope  $-1.39$ ) as a theorem rather than a calibration; (b) a completeness mechanism — the locator misses pair-merges below the grid resolution and prominence-limited ordinates (depth law  $54 \rightarrow 62 \rightarrow 83\%$ , measured only); (c) window extension: every verified window lies inside the Platt–Trudgian strip and the Nyquist reach — the contract explicitly does *not* promise new zero-free regions. *Kill*: a validated out-of-sample window with detection  $< 40\%$  at the frozen rule on a PD-certified rung deeper than  $X = 24.81$ ; a census contradiction on any PD-certified window; a false positive surviving the grid-refinement ward. *Mandatory typing (every module of the strand)*: location to  $\pm 0.25$  and exclusion to  $\delta \geq \delta_{\text{mb}}$  are strictly weaker than classical (Turing) verification on every axis; the strand is a necessary-side consistency demonstration of the explicit-formula reading and does not bear on RH. No marker move.

**The depth kill and the budget (PRIME.FLOOR.DEPTHKILL.01 / PRIME.FLOOR.BUDGET.01).** Registered 2026-08-07 ([verification/v829\_prime\_floor\_depth.py], 16/17 with the one preregistered-honest S5.HLAW FAIL + 12/12, ENVELOPE-HOLDS-DEEP + FAMILY-SCALES; [verification/v830\_prime\_float\_budget.py], 10/10 + 19 checks with the five preregistered-honest asymptotic FAILs, BUDGET-LINEAR-CLOSES-ALL + ENVELOPE-ASYMPTOTIC-PARTIAL). The PRIME.FLOOR.RATIO.01 kill gates survived at full sieve depth with growing margins; the certified family replaced the collapsing single pair; the Higham-linear budget removed the  $\alpha^* \approx 11$  horizon (citation grade to  $\alpha = 12.75$ ; the deployed v818 convention stays frozen, the linear budget is convention for new modules only); and the envelope is certified-explicit,  $\rho \geq 4.335 h^{-3/2}$  on 73/73, with the proof blocker named (alias phase correlations). Details in the dated note on PRIME.FLOOR.RATIO.01 above. No marker move; no RH statement. [2026-08-07, round 26: the named blocker is now *typed*, not merely named — see PRIME.FLOOR.PAIRCORR.01 below: the envelope-proof flank of the ratio contract is pair-correlation-equivalent substance, and the analytic-envelope route is closed at that boundary.]

**The floor-pair-correlation bridge (PRIME.FLOOR.PAIRCORR.01).** OPEN (registered 2026-08-07 — the honest typed closure of the analytic-envelope route; evidence [verification/v831\_prime\_alias\_second\_moment.py], ALIAS-CORRELATION-PAIRCORR, 15 checks with the seven preregistered-honest premise-overtaking FAILs, pattern-gated). *Statement (the bridge)*: the floor’s envelope constant is *Guinand-equivalent to the zero-side floor statement*. In the pole frame the correlation-corrected alias second moment is exact (sympy-proven):  $\tau = (S_2 - S_c)/2 - S_c^2/(4P_0) + O(P_0^{-2})$  with the coherent phase sum  $S_c$  multiplying the leading term; the measured facts are that the *amplitudes*, not the phases, carry the  $h^{-2.5}$  tower gap ( $-2.505 = -3.172 + 0.668$  exactly), and that  $S_c$  is carried by the prime-comb fluctuation  $S_p - S$  sitting at the Poisson-comb square-root scale ( $0.01 \times \sigma_{\text{sqr}}$ ; smooth and fluctuation shares  $-994/+995$ , cancelling; the comb scramble explodes the coherent sum  $10^6 \times$  — the cancellation is comb-specific arithmetic). Proving the tower  $\tau$ -law is therefore variance-type control of an oscillatory prime sum at its square-root scale — pair-correlation-equivalent substance on the zero side by the explicit formula — and by Guinand the comb’s self-cancellation *is* the zero-side floor statement itself. *Demands (both directions typed)*: (i) *forward* — an unconditional variance-type bound  $|\text{aniso}(S_p - S)| \leq C \sigma_{\text{sqr}}$  on the deployed frames, uniform in  $h$ , from prime-side input alone, with the pair-correlation-grade entry point *declared* (no hidden circularity); (ii) *backward* — the

calibration direction: convert measured floor margins into effective pair-correlation statements on the deployed battery (an instrument reading, never a theorem about zeros). *Kill*: a measured frame family where the suppression persists while the fluctuation share collapses below 0.6, or the scale ratio rising systematically above 3 (one-point regime — the boundary was mis-typed and the analytic-envelope route reopens on PNT-grade input). *Close*: the forward demand delivered with its entry point named and typed. *Stop list (house discipline)*: no further analytic-envelope variants — the envelope constant is arithmetic substance equivalent to the target. The contract explicitly does *not* promise an unconditional proof; deciding the bridge in the proving direction contains pair-correlation-grade arithmetic (fenced). A boundary, not a theorem; no marker move; no RH statement. [2026-08-07, round 27: both corner-era routes at this wall are now decided (see the two entries below). The corner route's *structural* half is closed weight-generically — the character-corner identity is polynomial algebra in free weights and the Hjelmslev CP tower carries it verbatim through five levels at certified defect zero — but identity and state corner are comb-blind and the level- $m$  motion is the pure register-dilution law  $16^{1-m}$ : the wall *relocates into the identification step* and stays this contract. The commutant SOS route is closed definitively by exact certificates (unique trivial feasible point; forced  $G = \text{diag}(1, -1, -1)$  with dual point  $q(0, 1, 0) = -1$ ). Both routes stop-listed; register-side Kraus dressings provably cannot carry the identification (POSITION-KRAUS-TRADEOFF, probe-level) — the one live tower door is position-dependent Kraus data on the level- $m$  flags. Demands, kill and close conditions unchanged.] [2026-08-07, round 28: the wall acquires its full coordinate system. (i) *The four closure coordinates are complete* ([`verification/v837_corner_closure_quantifier.py`], [`verification/v838_corner_expectation_position.py`]): register closed (the 90-cell class map has no identity $\wedge$ visibility $\wedge$ placement cell), tower level-exact ( $m \leq 3$ ; jet characters position-blind), compressions closed (all 5276 subgroup expectations, the full pinching, Stinespring-side compressions — identity and new placement mutually exclusive cell by cell), and carrier-side position dependence *extremally pinned* ( $\hat{c}_{\text{GL1}} = -1$  locks the entire  $C_2$  mass; identity-true carriers read zero on every self-consistent comb) — stop-list entries PRIME.CORNER.OPENDOORS.01 / PRIME.CORNER.LEVEL2.01 / PRIME.CORNER.EXPECTATION.01 / PRIME.CARRIER.POSITION.01. (ii) *The zero side is quantified* ([`verification/v839_paircorr_bridge_saturation.py`], [`verification/v840_gue_ablation_loopgain.py`]): the demand in Montgomery  $F(\alpha)$ -form with the named minimal missing input (unconditional  $F(\alpha, T) \leq C_F$  out to  $\alpha \approx 3$ ); the demand/GUE ratio saturates structurally at  $R_\infty = 1.11$  (tight band unfolded  $\alpha \in 1..2$ , pinned across the admissible family); the bootstrap loop is short ( $g = 1/(k^2 R^2 c_{\text{sup}})$ ; even ideal supply gives  $0.655 < 1$ ; Poisson ward passed) — the wall's self-conservation *is* the saturation, and the certified ladder doubles as a calibrated form-factor instrument ( $\hat{F}(\text{win}) = 0.44 \pm 0.14$ ). (iii) *The relation route opens* (the two registrations below): the identification carrier exists with the relational/multiplicative input, the identified corner's excess is positive on all 67 rungs, and the certified skeleton encloses  $\tau_X$  strictly positively — the wall in its sharpest form,  $\tau_X = \lambda_{\min}(\text{structural}) + \text{EXCESS}$ , with the SHARPENED DEMAND: a uniform lower bound on the excess margin (a finer-than-statistical datum that no finite table settles). Forward/backward demands, kill and close conditions of this contract unchanged.] [2026-08-07 round-31 transport note ([`verification/v855_invariance_atlas.py`]): the pair-type object this registration names — the alias second moment / the form-factor band  $\alpha \in 1..2$  — is *exactly* the bilinear chain fluctuation  $d\text{Chain}_2$  in Selberg-hierarchy coordinates (the demand concentrates there; the level-2 head is elementary-controlled by Selberg symmetry with on-range-verified constants) and the comb block's own soft direction at compression level (the minimizer tracking law, 0.0056 rad); the sieve input class reaches the band in *scope* (window-free, deterministic) but misses the *grade* by  $205 \times -1.1 \times 10^4 \times$  — input-class-invariant, the same wall; still [O].] [2026-08-08 round-32 note ([`verification/v856_connected_current_lean.py`], CONNECTED-COVARIANCE-PARTIAL): the pair-type object now has an explicit *zero-free construction* — the centered KMS correlator of the Möbius commutator current is the pair-correlation form (its stationary part carries the deployed pair weights  $\lambda_m^2 = \Lambda(m)^2/m$ ; the centered main term *is* the deployed smooth half-line subtraction, corr 0.99; the state is anchored to the critical line at  $1/2$ , the wrong half-weight losing the pair-weight

carrier), it *evades* the nuclear tax (relapse ratio  $\approx 0.3$  — genuinely different bookkeeping from the event-local completions), and the gap to the deployed form is the *grade gap* itself: the window form is linear in the comb and routing it through the correlator costs the variance,  $2.5\text{--}4.1 \times 10^5 \tau$  — the alias second moment is reachable by construction, the floor demand is not payable from it. The instrument (zero-free, state-anchored, exact main subtraction) is typed for measurement use; still [O.] [2026-08-08 round-33 note ([[verification/v861\\_krein\\_normalform\\_lean.py](#)] — [[verification/v866\\_source\\_contractor\\_formula.py](#)]): the floor demand of this contract now reads, in Krein coordinates,  $\|C\| \leq 1$  for the *explicit* closed-form source contractor  $C = W_- F G_+^{-1} F^H W_+$  (Loewner form  $G_- \preceq G_+$ ; the  $\tau$ -margin is the *single* soft defect mode's contraction margin — DEFECT-ONE), the scalar shadow is the pole-port impedance with the closed-form Poisson/-Cauchy  $a$ -term at  $s = 1/2$  and  $\kappa - 1 \approx 6e^{-\alpha/2} \rightarrow 1$ , and the absolute-bound route is *provably* closed (Perron  $\rho(|C|) > 1$  everywhere): the bridge input this contract demands must control the *phases* of the pair kernel, not its magnitudes — the pair-correlation substance in its sharpest operator clothing (PRIME.KREIN.CONTRACTOR.01); still [O.] [2026-08-08 round-34 note ([[verification/v869\\_cdcore\\_clarkphase.py](#)] — [[verification/v872\\_damping\\_compensation.py](#)]): the wall now reads as *one* two-term inequality at *one* named direction, both terms source-built — the exact split  $I - C^*C = T_1 + T_2$  (the indefinite cross-measure Jacobi geometry plus the PSD Christoffel damping, warded at  $10^{-12}$  on all 42 rungs) makes the floor demand  $v^H T_2 v > |v^H T_1 v|$  at the soft direction (kz 9:  $-0.0328 + 0.0330 = 1.68 \times 10^{-4} = \tau$ ;  $s = \kappa\tau$  with the certified bulk floor at 98–100% of  $\lambda_2$ ), and the kernel behind both terms is *named*: the Christoffel–Darboux kernel of the explicit source measure — one Jacobi chain plus arm weights, the controls' chains  $O(1)$  different. The pair-correlation substance this contract types is now the deviation of the source Jacobi chain and Christoffel weights from their free counterparts, read at a single named direction; the classical escapes are stop-listed (pole transformations at  $x_{\text{pole}}$ , finite-rank cross-defect corrections,  $\mu_4$ -offset Clark kernels); the preregistered next test is PRIME.PRUEFER.COMPENSATION.01 (see the contractor block); still [O.] [2026-08-08 round-35 closing note ([[verification/v873\\_pruefer\\_cotlar\\_decision.py](#)], the *executed* preregistration): the Cotlar decision is in — the round-34 frozen contract was executed under the full provenance chain (committed pre-execution in 526ca3eb; two documented mechanical fixes with the spec literal SHA byte-identical throughout; the honest FAILs pattern-gated, nothing refit) and returns COTLAR-GROWING by its own rules: the blockwise Cotlar–Stein / phase-cell class on source-only Prüfer phases is *closed* and stop-listed — the battery is flat ( $S_U$  5.55–6.30) but the frozen deep holdouts explode (kz 142:  $S_U$  32.07/ $S_C$  58.78; slopes 0.209/0.408  $\geq 0.20$ , both channels GROWING; kills 2 + 5 fire). The demand narrows: what survives is a *non-stationary* phase bound on the explicit kernel — the predeclared  $\pi$ -resonance cells  $d = 7\text{--}8$  are the known critical structure (the envelope is bimodal exactly there), and the typed growth law constrains every candidate: the explosion tracks *kz-depth* ( $\alpha$ ), not window size  $h$  (kz 90/116 at  $h = 1430/1433$  stay in band; kz 177 at smaller  $h = 1219$  explodes) — a uniform certificate must be  $\alpha$ -aware. The carry-forward assets: the danger geometry is real (source-only cells capture 0.98–0.61 of  $\lambda_{\min}(T_1)$  without the soft eigenvector) and the discrimination stands (Epstein 3927%, scramble 314462%); still [O.] [2026-08-08 round-36 correction note ([[verification/v877\\_complete\\_comb\\_reversal.py](#)]): the round-35 note's holdout data was comb-truncated (ATOM\_MAX artifact; complete-comb floors strictly positive) — on complete combs the identically frozen rule returns BOUNDED, the class closure and stop-list entry are amended, the  $\alpha$ -not- $h$  law is retired, and the route reopens (PRIME.PRUEFER.COMPENSATION.02); the floor demand in embedding coordinates is now the registered target PRIME.CARLESON.PRIME.01 [O] (the Carleson embedding of  $\tilde{\nu}_-$  into  $\tilde{\mu}_+$  at degree  $h-1$  with constant  $1 - \tau_h$ , machine-exact reduction [[verification/v876\\_carleson\\_polygon\\_wave.py](#)]); still [O].]

**The character corner and the Hjelmslev tower (PRIME.CORNER.CHARACTER.01 / PRIME.HJELMSLEV.CPTOWER.01).** Registered 2026-08-07 ([[verification/v835\\_corner\\_hjelmslev\\_tower.py](#)], 21/21 + 20/20, CORNER-IDENTITY-SYMBOLIC + HJELMSLEV-STRUCTURE-ONLY). *Certified/exact*: the corner identity  $T_{\text{GL1},X} = V^* e_{\text{GL1}} \pi(P_X) e_{\text{GL1}} V$  as polynomial algebra in free event weights (symbolic toy leg exact; all three deployed windows on two independent code paths);  $\hat{c}_j = -1$

exactly for all 136 events, consuming only the glue identity  $\sum_m \Theta_m = 240 \sigma_3$  and the uniform  $H^*$  mass;  $e_{GL1}$  a rank-one projector, corner compression single-Kraus CP; the chain-ring ladder  $\mathbb{Z}[i]/(1+i)^m$ ,  $m = 1..5$ , with exact counts  $15 \cdot 8^{m-1}/35 \cdot 16^{m-1}/105 \cdot 32^{m-1}$ , strictly projective CP channels (certified cb defect identically zero) and level-exact corner coefficients. *Typed (the decision)*: identity and state corner are comb-blind (scramble control); the level- $m$  corner motion is the pure register-dilution law  $16^{1-m}$  — no register/geometry lift carries the prime placements; the arithmetic wall relocates into the identification step and stays PRIME.FLOOR.PAIRCORR.01. Register-only tower lifts are stop-listed; the live door is position-dependent Kraus data on the level- $m$  flags (register-side dressings provably cannot carry it: POSITION-KRAUS-TRADEOFF, 13/13, probe-level). *Kill*: any level breaking the exact counts/defect-zero wards, or a register-only lift moving the comb-specific gap. No marker move; no RH statement.

**The commutant SOS closure (PRIME.COMMUTANT.SOS.01).** *Registered* 2026-08-07, closed at registration ([verification/v836\_commutant\_sos\_closure.py], 29/29, COMMUTANT-SOS-INFEASIBLE; the declared SPEC v1→v2 witness-triple repair documented, intent/bars/verdict rule unchanged). *Certified/exact*: the five commutant coordinates  $\Pi_A(l)$ ,  $l \in \{0, 1, 4, 7, 9\}$  (exact Lagrange projections, traces  $\{5, 2, 5, 2, 1\}$ , structure constants  $\Pi_a \Pi_b = \delta_{ab} \Pi_a$  on all 25 pairs,  $\dim C(\{K, \sigma\}) = 39$  exact); the reduction lemma (abelian conjugation reads only Gram diagonals), collapsing the SDP to a rational LP with the *trivial diagonal rewriting* as unique feasible point (Vandermonde kernel  $\{0\}$ ) — no cross-sector positive transfer; the flag-space Gram lift forced to  $G = \text{diag}(1, -1, -1)$  (not PSD) with the rational dual certificate  $q(0, 1, 0) = -1 < 0$  and achievable  $q$ -negative directions (including 832/2799 nonneg-cone truncation cuts). *Typed*: FENCE-CLEAN on the algebra layer, FENCE-HIT on the positivity branch — the nontrivial version of the route is pair-correlation substance (PRIME.FLOOR.PAIRCORR.01). Route closed; stop-list entry. No marker move; no RH statement.

**The corner-closure quantifier (PRIME.CORNER.OPENDOORS.01 / PRIME.CORNER.LEVEL2.01 / PRIME.CORNER.EXPECTATION.01 / PRIME.CARRIER.POSITION.01).** *Registered* 2026-08-07, closed at registration (stop-list entries; [verification/v837\_corner\_closure\_quantifier.py], 19/19+17/17+3, DOORS-CLOSED + LEVEL2-CLOSED; [verification/v838\_corner\_expectation\_position.py], 13/13+13/13, EXPECTATION-CLOSED + POSITION-CARRIER-TRADEOFF). *Certified/exact*: the 18-class collapse of all 128 characters (product law verified on the full carrier Fourier, exact Fractions); the generalized mass ward (the character corner reads one position-blind carrier mode per event); the level- $m$  duals and polar annihilators ( $m \leq 3$ ; identity characters exactly  $\text{Ann}(\langle S_m \rangle)$ , Vfac spectra  $\{1, 0, -1/7\}$ ); the full subgroup-lattice census (5276 subgroups, 74 259 components; the pinning criterion  $\text{Ann}(H) \subseteq \{1, \psi_0\}$  exact); the pinching completeness; the exact tower address (period 4 on rational integers). *Typed (the closure)*: identity  $\Rightarrow$  not visible in every cell of every class map; the only position-reading slot is exactly the slot the identity freezes; the deployed-form identity is an *extremal* state condition (EXTREMAL PINNING) — no state-preserving, identity-true carrier map  $F(n, u)$  separates self-consistent combs. The structural half kernel-checked (TfptCarrier/SectorPositiveDescent.lean; IdentificationPositivity the single named hypothesis). *Kill (of the closure)*: a cell of any measured class map carrying identity, visibility and placement together, or an identity-true state carrier separating the self-consistent null. The wall stays PRIME.FLOOR.PAIRCORR.01. No marker move; no RH statement.

**The bridge map, the saturation and the loop (PRIME.FLOOR.BRIDGEMAP.01 / PRIME.FLOOR.GUESAT.01 / PRIME.FLOOR.GUEABLATION.01 / PRIME.FLOOR.LOOPGAIN.01).** *Registered* 2026-08-07, evidence lines feeding PRIME.FLOOR.PAIRCORR.01 ([verification/v839\_paircorr\_bridge\_saturation.py], 9/9+7/7, BRIDGE-FULLY-BEYOND + GUE-SATURATING; [verification/v840\_gue\_ablation\_loopgain.py], 10/10+8/8, SATURATION-STRUCTURAL + LOOP-SHORT; scholarship correction throughout: Montgomery's  $F(\alpha)$  on  $|\alpha| < 1$  is RH-conditional, every supply row tagged). *Measured*: the demand ledger in Montgomery  $F(\alpha)$ -form on all 73 frames; the two-tier gap (unconditional supply misses by 52×; conditional window covers 43–86%; full-GUE rms closes 39/73); the named minimal missing input ( $F(\alpha, T) \leq 0.33$ –1.7 out to  $\alpha \approx 3$ ); the saturation plateau  $R_\infty = 1.11$  inside the frozen band, structural across the

Guinand-admissible family and pinned in the *unfolded* coordinate; the calibrated form-factor instrument  $\hat{F}(\text{win}) = 0.44 \pm 0.14$  at  $T \leq 2 \times 10^4$ ; the loop-shortfall law  $g = 1/(k^2 R^2 c_{\text{sup}})$  with headline  $g(k=1) = 0.529$  and the ideal-supply bound  $1/R^2 = 0.655 < 1$ ; the Poisson circularity ward. *Typed*: necessary-side consistency and instrument calibration only — the saturation is the strongest form of comb-zero bridge consistency a necessary-side battery can deliver, the loop is *not* a proof architecture, and the measured positivity everywhere is finer-than-statistical information. *Kill*: a Guinand-admissible variant landing off-plateau with positivity intact (the structural reading fails), or the instrument's Poisson control drifting from 1. No marker move; no RH statement.

**The relation carrier (PRIME.RELATION.MULT.01 [O], carrying PRIME.RELATION.LADDER.01).** *Registered* 2026-08-07, OPEN ([verification/v841\_relation\_carrier\_ladder.py], 18/18 + 13/13, RELATION-CARRIER-EXISTS + EXCESS-NONNEGATIVE). *The finding*: the identification carrier *exists* with the relational input — the first structure in the program to pass all four gates (state, identity at truth by degeneration, placement, and the self-consistency null that killed every compression and every single-event carrier). *Certified/exact*:  $\mu * 1 = \delta$  and  $\Lambda = \mu * \log$  as exact structure in the truncated convolution frame (the relational input *is* the Euler-product datum); the four-comb discriminator (true comb 0 exactly;  $\zeta_{\mathbb{Q}(i)}$  and  $\chi_4$  blind — Selberg-class-correct; the  $h = 2$  Epstein comb  $x^2 + 5y^2$  separates at 200.0 on the class-group obstruction sites); the amplitude wiring ( $L$ -safety automatic); support and identity wards on all 67 reachable rungs. *Measured*: the identified corner's excess over its comb-blind structural skeleton is positive on all 67 rungs (+2.29 to +3.70) while  $\tau_X$  decays to  $1.7 \times 10^{-5}$  — the sharpest wall form  $\tau_X = \lambda_{\min}(\text{structural}) + \text{EXCESS}$ . *Demand (sharpened)*: a *uniform* lower bound on the excess margin along the ladder, with its arithmetic entry point declared (PRIME.FLOOR.PAIRCORR.01 discipline binding — no hidden circularity); interim: amplitude-grade wiring for combs with claimed coefficient data beyond the four frozen test combs. *Kill*: a reachable rung with certified-negative excess at the true comb, or an  $L$ -function comb falsely flagged by the amplitude reading. *Close*: the uniform excess bound delivered and typed. Explicitly *not* promised: positivity along the ladder, or any RH-relevant bound. No marker move; no RH statement.

**The certified skeleton (PRIME.RELATION.SKELETON.01 [O]).** *Registered* 2026-08-07, OPEN ([verification/v842\_excess\_certified\_skeleton.py], 9/9, SKELETON-CERTIFIED). *Certified*: strictly positive interval enclosures of  $\tau_X$  on *all* 67 reachable rungs ( $\alpha = 2.773\text{--}6.304$ , masses to  $3 \times 10^5$ ; widths  $5.4 \times 10^{-11}\text{--}6.6 \times 10^{-9}$ , three to five orders below the margin; chol\_cert discipline above the deployed float data; certified midpoints reproduce the float ladder at  $4.9 \times 10^{-13}$ ); the certified discriminator (the  $h = 2$  fake fails its margin certificate on every anchor, routed-corner enclosure strictly negative, exact rational  $\Lambda_F$  routing). *Typed (the scaling law)*: certifiability degrades faster than the margin decays (width slope +0.62 vs margin slope  $-0.52$  per  $\alpha$ ), so the coordinate system has a finite certification horizon — with large headroom:  $\alpha^* \approx 9.1$  (mass  $\sim 7.5 \times 10^7$ ), far beyond the deployed ladder end; certifiability is not the binding constraint on any reachable rung. *What stays open (the wall in its sharpest coordinates)*: the infinite quantifier over the sequence of strictly-positive certified enclosures — the sharpened demand of the relation route: a *uniform* lower bound on the excess margin, a finer-than-statistical datum by the GUE-side findings; no finite table, certified or not, settles it. *Kill*: a reachable rung whose certified enclosure at the true comb loses strict positivity under the frozen budget discipline, or a certified-positive enclosure for the  $h = 2$  fake. *Close*: the uniform excess bound delivered with its entry point declared. Explicitly *not* promised: an unconditional proof. [2026-08-07 round-29 closing note: the quantifier is in its final coordinates ([verification/v843\_margin\_law\_excess\_lean.py]): the demanded object sits on the one scalar series  $e_1(m) = (\tau/\tau_{\text{pnt}}) h^{3/2}$ , certified  $\geq 4.335$  on every reachable rung and doubly derived (the consistency ward); the wall is *self-similar at cell level* — the excess is itself a growing cancellation (CI median 15.6, rising  $5.65 \rightarrow 29.46$ ; 22 287 of 351 122 per-prime aggregates nonpositive — no per-prime positive control); the tower gives scale, not recursion; the compiler connection is null (0/96). The finite-table limit is now *kernel-checked*: TfptCarrier/ExcessSkeleton.lean proves pointwise\_pos\_not\_uniform (per-rung positivity at every rung does not yield a uniform bound) and the bridge theorem excess\_floor consumes UniformMarginBound as the single named

hypothesis — this registration’s “no finite table settles it” is a theorem, not a typing. Same-class margin-law re-attacks are stop-listed per the measured self-similarity. Demand, kill and close conditions unchanged; still [O.] [2026-08-07 round-30 closing note (the three-architecture decision)]: the three v5.4 proof-architecture candidates aimed at this registration’s demand are measured and typed *closed* at three walls that are one wall ([`verification/v846_schur_spectral_mother.py`], [`verification/v847_wedge_cellcone_transport.py`]): (i) the relational Schur–Gram route dies structurally at the forced Cauchy–Schwarz diagonal price (price/ $\tau \sim 3.1 \times 10^4$ – $4.5 \times 10^4$  at the density-anchored floor; trace *and* form; the untouched-projection deviation *is* the price), and the price is *geometry-independent* — the unitary spectral mother ( $J^2 = I$ ;  $U_d U_m = U_{dm}$  exact) gains 20–33 $\times$  at symbol grade but never at Gram grade (the modulus bound); harvesting interference is Fejér–Riesz of the total symbol = the positivity itself; (ii) the sign-register wedge lift *exists* as exact algebra (nonnegative wedge weights per frame, where the naive signed Lagrange provably cannot) but the frame-uniform law fails at  $\sim 15\%$  against the  $4.4 \times 10^{-5}$  floor need — the commutant uniformity wall transported; (iii) the completed-cell Lorentz transport breaks the cone at the  $n = 2$  cell on 67/67 rungs under both predeclared groupings (endpoint back in-cone = the certified  $\tau > 0$ ; ray cascade to  $(5, -3, 4)$  confirmed, Kendall  $\pm 1.000$ ), with the inverted Epstein signature kept as its frozen FAIL: the fake stays in-cone — the violations are the true comb’s own  $\psi(x) - x$  fluctuation; the arithmetic sits *in* the violations. *Stop list (extended)*: diagonal relational mothers and manifest-Gram interference harvesting; constant (frame-uniform) wedge weight laws (the named next object is an  $h$ -dependent, analytically controlled law); unrenormalized small- $n$  completed-cell transports (the named next object is a renormalized cell map absorbing the  $n = 2$  fluctuation). The extraction side is now separated and complete (PRIME.EXTRACTION.CHAIN.01 below): the wall is exactly hypothesis (H) on a ladder cofinal in the mesh-refinement order. Demand, kill and close conditions unchanged; still [O.] [2026-08-07 round-31 closing note (the invariance atlas)]: the wall this registration names is now proven *invariant* across four axes ([`verification/v850_completion_tax_flow.py`], [`verification/v851_cluster_kernel_field.py`], [`verification/v855_invariance_atlas.py`]): (i) *compressions* — the nuclear-norm tax theorem closes the entire positive-dilation completion class with fixed comb coupling (exact optimum  $\geq 4057.6 \times \tau$  on every rung, diverging  $\sim h^{0.24}$  against  $h^{-3/2}$ ; zero-gap dual certificates; the measured v846 designs were suboptimal by only 13.7–20.4 $\times$  — optimality could not have saved them), the finite-dimensional spectral-flow index is a reformulation of the endpoint pivot signs (distance-to-crossing = metric margin / flow velocity, ratio 1.00), and the 2-dim corner cluster expansion is blind to multiplicativity (position-geometric weights); (ii) *geometries* — the tax is basis-invariant ( $\leq 1.1 \times 10^{-15}$ ): the round-30 TV-price geometry-independence is now a corollary; (iii) *criteria* — Weil-window = Li = Nyman–Beurling (BD-constant =  $2\lambda_1$  at  $4.9 \times 10^{-17}$ ; the NB span *is* the spectral-mother geometry; the Li transfer fails by cone geometry: certified partial data everywhere, one uniform statement missing everywhere); (iv) *input classes* — the sieve toolbox reaches the  $\alpha \in 1..2$  band in scope but misses the grade by  $205 \times$ – $1.1 \times 10^4 \times$ , and the demand at truth exceeds the factor-2 gate on every rung: unclosable by *any* upper-bound input class. *The one object in four languages*: the uniform  $e_1$  bound = the form-factor band  $\alpha \in 1..2$  = the bilinear chain fluctuation  $d\text{Chain}_2$  = the comb block’s soft direction (the minimizer tracking law, 0.0056 rad). *Stop list (extended)*: positive-dilation completions with fixed comb coupling, in any basis; finite-dimensional index/spectral-flow reformulations; 2-dim corner cluster expansions; elementary upper-bound supply rows against the factor-2 gate. The graded kernel law (PRIME.CELLCONE.GRADEDKERNEL.01 below) is the round’s positive core; the minimal hypothesis is now ( $H_{\text{cof}}$ ) (FORM.PRIME.COFINAL.WEIL.01 below — UniformMarginBound demoted to an over-strong sufficient lemma). Demand, kill and close conditions unchanged; still [O.] [2026-08-08 round-32 note (the grade law)]: the graveyard this registration accumulated is now *one* kernel-checked homogeneity theorem (TfptCarrier/GradeNoGo.lean, FORM.PRIME.GRADE.NO\_GO.01 [F]; numeric mirror [`verification/v859_grade_no_go_elevator.py`]): free positivity lives at *quadratic* grade (Gram squares of linear maps are 2-homogeneous), the deployed target at *linear* grade, and no exact match exists on any scaling ray — with the affine escape priced exactly (the representation error is the sum of two PSD prices) and the background killed too. The measured exchange rate is  $10^4$ – $10^5$  across

four instances: the Schur–Gram diagonal price ( $3.1\text{--}4.5 \times 10^4 \tau$ ), the class optimum ( $\geq 4057.6 \tau$ ), the covariance variance price ( $2.5\text{--}4.1 \times 10^5 \tau$ , [verification/v856\_connected\_current\_lean.py] — the first design that *evades* the nuclear tax, dying on the grade gap instead), and the wedge uniformity gap. The only licensed grade-1 mechanism is the *rates elevator* (`tangent_psd_on_kernel_null`, kernel-checked), and its deployed instantiation is verified real but the Lévy/conditional-positivity class is *closed on both deployed sides* by signed spectral densities (the comb at exactly 50% negative mass; the arch on the digamma band ending at  $\tau^* = 6.27$  where  $\operatorname{Re} \psi(\frac{1}{4} + i\tau/2) = \log \pi$ ). The finite-head phase-rescue family is closed by *theorem* ([verification/v860\_phase\_lever\_closure.py]: private-best-phase budget margins  $\geq 1.3894 / 0.6296$  on all 67 rungs; the D-frontier saturates at  $0.73 < 1$ ). *Stop list (extended)*: finite-head phase rescues on the deployed frame family; amplitude-Gram constructions without grade accounting (exact matching of the linear window read by any Gram square, affine backgrounds included); Lévy/conditional-positivity elevators on the deployed sides. The rates door itself stays licensed and typed; demand, kill and close conditions unchanged; still [O].] [2026-08-08 round-33 note (the wall’s operator form): the demand now has exact operator coordinates ([verification/v861\_krein\_normalform\_lean.py] – [verification/v868\_divisor210\_audits.py]): the deployed window form is the Krein normal form  $Q = \|B_+t\|^2 - \|B_-t\|^2$  with the Douglas pencil identity  $1 - \|C\|^2 =$  the  $\tau$ -margin, the defect has rank *one* ( $\lambda_1 =$  the  $\tau$ -margin exactly — one arithmetic-carrying soft mode), the contractor factors *exactly* as the target-free closed-form source expression  $C = W_-FG_+^{-1}F^HW_+$  (blind holdouts, certified defect transfer), the norm statement is the exact reformulation  $\|C\| \leq 1 \Leftrightarrow G_- \preceq G_+ \Leftrightarrow K \succeq 0$ , and the theorem frame is kernel-checked (FORM.PRIME.KREIN.DEFECT.01 [F]: `SourceContractor` the named hypothesis, composed to Weil nonnegativity). What the campaign *proves* about the tool space: the absolute half of the classical toolbox is closed (the Perron value  $\rho(|C|) = 2.4\text{--}3.1 > 1$  everywhere — no absolute-value inequality can ever certify: the certification must be *phase-aware*); quadrature does certify  $s > 0$  on every reachable rung (the Neumann frontier killed) but the depth law  $m^* \sim \sqrt{\operatorname{cond}(G)}$  is *intrinsic* (no source preconditioner waives it — the soft end is a continuum edge of distributed arithmetic structure); the scalar shadow is the pole-port impedance with the closed-form  $a$ -term and  $\kappa \rightarrow 1$ ; the register is canonical up to *proven* gauge. *Stop list (extended)*: source-word searches at all lengths (the monomial closure theorem: every word of any length misses by  $\geq 0.87$ ); absolute-value bounds on the weighted frame kernel (the Perron closure); fixed-depth quadrature certificate families without a conditioning bound; chirality deciders through quadratic readouts. The named demand is sharpened to PRIME.KREIN.CONTRACTOR.01 [O] (below): one phase-aware bound on the closed-form kernel; still [O].] No marker move; no RH statement.

**The margin law and the kernel-checked wall (PRIME.MARGIN.LAW.01 + FORM.PRIME.EXCESS.SKELETON.01).** *Registered* 2026-08-07 ([verification/v843\_margin\_law\_excess\_lean.py], 10 checks with the one preregistered-honest S2.2 FAIL kept and pattern-gated, +4 Lean-mirror checks, MARGIN-RECURSES-THE-WALL; Lean core `TfptCarrier/ExcessSkeleton.lean`, lake build green, no *sorry*, carried in the shipped Lean manifest). *Measured (the law)*:  $\tau = e_1(\alpha) h^{-3/2} \tau_{\text{pnt}}(\alpha)$  with the envelope constant doubly derived ( $e_1 \in [4.855, 24.209]$  on all 67 rungs,  $\min \geq 4.335$ );  $\tau_{\text{pnt}}$  carries the entire  $\alpha$ -decay,  $e_1$  the arithmetic; the fit-free  $h$ -exponent median  $-1.341$  vs bar edge  $-1.35$  (the one honest FAIL, not refit). *Measured (the negatives, typed)*: the tower gives scale, not recursion (dyadic-level ratio defects non-decaying); the excess is a growing cancellation — the wall self-similar at cell level (no per-prime positive control, no Euler-factorized positivity); the compiler connection is null (0/96, Bonferroni-honest). *Kernel-checked [F]*: the two-giants decomposition, the interval-certificate and certified-negative discriminator lemmas, the finite-ladder quantifier, `excess_floor` with `UniformMarginBound` the single named hypothesis (implying the per-rung `IdentificationPositivity` instances; the converse fails), and `pointwise_pos_not_uniform` — no finite certificate table discharges the uniform bound. *Stop list*: same-class margin-law re-attacks (tower-recursion / cell-factorization of the excess). *Demand (unchanged, final coordinates)*: a *uniform* lower bound on  $e_1$ , with its arithmetic entry point declared (PRIME.FLOOR.PAIRCORR.01 discipline binding). *Kill*: a reachable rung with  $e_1 < 4.335 \times 0.999$ , or a kernel-checked discharge of `UniformMarginBound` from structural input alone (which would contradict the typed bound-

ary and force a re-audit). [2026-08-07 round-31 demotion note (FORM.PRIME.COFINAL.WEIL.01): UniformMarginBound is from this round *formally demoted to an over-strong sufficient lemma* in the load-bearing chain — the minimal H theorem (TfptCarrier/CofinalWeil.lean, numeric mirror [verification/v849\_cfin\_unique\_cofinal\_lean.py]) proves uniform  $\subsetneq$  pointwise  $\subsetneq$  cofinal kernel-checked in both directions, and the extraction chain consumes exactly the strictly weaker ( $H_{\text{cof}}$ ): PSD at the rungs of one pre-fixed cofinal ladder, no uniform  $\delta$ . This block’s lemmas and excess\_floor remain kernel-checked and load-bearing as the *sufficient* side; the named wall moves to ( $H_{\text{cof}}$ ). Still [F]/[O] as stated.] No marker move; no RH statement.

**The continuum extraction chain (PRIME.EXTRACTION.CHAIN.01 [O]).** Registered 2026-08-07, OPEN ([verification/v848\_extraction\_chain.py], 19/19, EXTRACTION-CHAIN-COMPLETE; hypothesis (H) is *never evaluated* — no positivity statement about any actual tower form or the actual Weil functional is made). *Measured:* on the canonical  $(D, X)$  ladder, cofinal in the *mesh-refinement* order ( $D_j = 2^{-j}$ ,  $j = 4..11$ , 68–8704 lags, the full pinned prime-power table) every element of the canonical dense family identifies its independently computed Weil value  $Q_W = \text{POLE} + \text{ARCH} + \text{ATOM}$  (exact piecewise-cubic autocorrelations in  $\mathbb{Q}$ , fifth-point certificates; Suzuki normalization, W1-certified layers) with rates  $\sim 4^{-j}$  (median  $-1.818$  per level), a geometric ledger (median ratio 0.302, decay law  $-1.752$ ), and the  $X$ -direction *exactly* stable once the window covers the element (44 cutoff pairs at float zero); the pure-box class carries the *exact* Weil value at every finite level (maximal error  $5.3 \times 10^{-14}$  — the sharpest dyadic-inheritance statement); recovery is constructive and Cauchy (the exact strong- $L^2$  law in  $\mathbb{Q}$ ); the gated far-spike liminf family has strictly positive margins, while the band directions are typed and *never* gated — wall content under (H), refused by design. Both controls fire (the density-only wrong limit stalls at ratio  $1.6 \times 10^6$ ; the scrambled tower inflates the ledger by  $5.5 \times 10^5$ ). *The deliverable (the quantifier reduction):* under (H) — finite positivity of the tower forms along a ladder cofinal in that same mesh-refinement order — every ladder value is nonnegative and the limit passage is arithmetic-free: *no* Mosco compactness, *no* uniform  $\delta > 0$ , *no* diagonal argument enters the implication; per-element convergence is the only analytic input, and the three infinite-level ingredients of PRIME.MOSCO.SELECTION.01 migrate to the operator-level selection programme, where they stay typed. With the classical legs cited (density; the continuity of  $Q_W$  at fixed support — *corrected* in the round-CCCLXV note below, the pure sup-norm  $C^0$  form of this leg being false; Weil 1952, Bombieri 2000, Suzuki arXiv:2606.09096, Iwaniec–Kowalski Thm 5.12), the chain *cofinal finite positivity*  $\Rightarrow$  *Weil positivity*  $\Rightarrow$  (via Weil’s criterion) *the target* is theorem-grade modulo the named citations. *What stays open (by design, the whole wall):* hypothesis (H) — equivalently the uniform demand of PRIME.RELATION.SKELETON.01 restricted to a ladder cofinal in the mesh-refinement order — and *nothing else*. *Kill:* an element of the dense family whose ladder values converge to a value different from the independently computed  $Q_W$ , a non-summable ledger on the true tower, or a failure of the  $X$ -stabilization on a covered element. *Close:* any discharge of (H), which via this chain yields Weil positivity and, via the criterion, the target. Explicitly *not* promised: an unconditional proof. [2026-08-07 round-31 note: step 2 of the chain — the arithmetic-free limit passage under (H) — is now *kernel-checked* (TfptCarrier/CofinalWeil.lean: the  $\varepsilon/2$  argument, the no-diagonal dense-family form, the assembly; axioms clean), and hypothesis (H) is formalized as CofinalHypothesis ( $H_{\text{cof}}$ ), preregistration-shaped (the ladder is data, never mined from measured signs) and proven strictly weaker than UniformMarginBound (FORM.PRIME.COFINAL.WEIL.01 below). The genuinely open item is unchanged: ( $H_{\text{cof}}$ ) itself; still [O].] [2026-08-13 round-CCCLXV note ([verification/v912\_form\_convergence\_theorem.py], 35/35, PRIME.FORM.CONVERGENCE.THEOREM.01): *two of the inputs above change type*. (i) Per-element form convergence — the “only analytic input” of the quantifier reduction — is no longer measured but *proven unconditionally* on the deployed dyadic cellwise-affine family, with the explicit rate  $O(D^2 \log(1/D)) = O(2^{-2j}j)$  and explicit constants ( $\Theta_0 = 3 \log 2 + \pi/2 + \gamma + \log \pi$ ; log coefficient exactly  $7/24$  for continuous  $f$ ); the load-bearing step is the exact defect lemma  $k_d - K(dD) = -(D^2/12)G(dD)$ , verified with zero deviation in exact rationals, and the only cited arithmetic input is Rosser–Schoenfeld  $\psi(x) < 1.03883x$  for support uniformity. (ii) The continuity leg as cited above is *false* in the pure sup norm: the explicit even Lipschitz family

$e_n(w) = \frac{1}{n} \min(1, we^{n^2})(1 - w/2)_+$  has  $\|e_n\|_\infty \rightarrow 0$  while  $|A[e_n]|$  grows linearly ( $2.57 \rightarrow 12.09$ ). The correct hypothesis — uniform convergence plus an equi-Lipschitz (Dini) condition at the origin, supplied by the admissible even compactly supported BV class — replaces it here and wherever the contract cited it; the *density* leg in that topology stays cited. This is a correction of a wrong cited statement, not a downgrade. What stays open is unchanged: ( $H_{\text{cof}}$ ), still [O], still of RH strength.] No marker move; no RH statement.

**The minimal H theorem (FORM.PRIME.COFINAL.WEIL.01 [F]).** *Registered* 2026-08-07, FORMAL (TfptCarrier/CofinalWeil.lean, 11 declarations, kernel-checked — lake build green at 3415 jobs, no sorry, no native\_decide, axioms propext/Classical.choice/Quot.sound only, verified by #print axioms on all eight theorems; numeric mirror [verification/v849\_cfin\_unique\_cofinal\_lean.py] part B; carried in the shipped Lean manifest at 87 files). *Proven:* the cofinal hypothesis ( $H_{\text{cof}}$ ) — a *pre-fixed* strictly monotone ladder, cofinal in the *mesh-refinement* order in which the form convergence holds and not in the window or cap parameter alone (at a fixed mesh the read is cap-independent, hence eventually constant, and its limit carries the frozen quadrature defect of that mesh), along which the tower matrices are PSD; preregistration-shaped: the ladder is *data* of the structure, supplied before any certificate, never mined from measured signs — and the minimal implication: `limit_nonneg_of_cofinal_seq` (the  $\varepsilon/2$  argument at sequence level) and `weil_nonneg_of_cofinal` (per-element form convergence  $+(H_{\text{cof}}) \Rightarrow$  the limit functional nonnegative on the *whole* dense family, with *no* diagonal argument — one PSD certificate per rung covers every element simultaneously); the strict hierarchy  $\text{uniform} \subsetneq \text{pointwise} \subsetneq \text{cofinal}$ , kernel-checked in both directions (witnesses  $1/(m+1)$  — reusing `pointwise_pos_not_uniform` — and  $\pm 1$  on the even ladder); and the assembly `cofinal_weil`. *Absent by design:* no uniform margin, no inverse/resolvent, no Mosco coercivity, no limit operator, no diagonal argument. *The formal demotion (the round's status statement):* the load-bearing chain of PRIME.EXTRACTION.CHAIN.01 consumes *exactly* ( $H_{\text{cof}}$ ); `UniformMarginBound` implies ( $H_{\text{cof}}$ ) and the converse fails kernel-checked — `UniformMarginBound` is from this round an over-strong *sufficient lemma*, no longer the named wall; the wall content has strictly decreased to cofinal PSD on one pre-fixed mesh-refining ladder. *The honest boundary:* ( $H_{\text{cof}}$ ) stays a hypothesis everywhere; nothing proves, gates or evaluates positivity of any actual tower form; no continuum theorem, no RH statement, no arithmetic identification is formalized or claimed.

**The graded kernel + tail theorem candidate (PRIME.CELLCONE.GRADEDKERNEL.01 [O]).** *Registered* 2026-08-07, OPEN ([verification/v851\_cluster\_kernel\_field.py] part C, 7/7, verdict KERNEL-GRADED). *The candidate statement* (measured on the 67 deployed rungs under the G1/Stieltjes completion; *not* proven, *not* uniform in  $\varepsilon$ , not an RH claim): for every  $\varepsilon > 0$  there exist finite  $P_0(\varepsilon)$  and  $X_0(\varepsilon)$  such that for every window  $X \geq X_0(\varepsilon)$ , every state of the completed-cell flow beyond  $n = P_0(\varepsilon)$  has relative depth  $-\lambda_{\min}/\lambda_{\max} \leq \varepsilon$ . *Measured witnesses:*  $P_0(10^{-1}) = 0$  on all rungs;  $P_0(10^{-2}) \leq 73$ , collapsing to  $\{2, 3\}$  above  $\alpha \approx 3.9$ ; the skin-depth law  $D(X) = 0.754 X^{-0.548}$  ( $R^2 = 0.99$ , Kendall  $-1.00$ ) predicts  $X_0(\varepsilon) = (\varepsilon/0.754)^{-1/0.548}$  ( $X_0(10^{-3}) \sim 1.8 \times 10^5$ ); the strict ( $\varepsilon \rightarrow 0$ ) census wanders with the horizon ( $P_0 \sim X^{1.03}$  at relative position 0.98) — the wandering is vanishing-depth skin, *not* kernel growth; the G2 mechanism typed by contrast (centroid offset  $1653\times$ , running imbalance  $39\times$  — delayed compensation). *Honest scope:* “resolved” is an in-range statement; the  $\varepsilon < 10^{-2}$  rows are extrapolated only through the fitted law. *Kill:* a rung family on which the depth- $10^{-2}$  kernel grows without bound, a measured break of the skin-depth power law, or a G1 state beyond the predicted  $P_0(\varepsilon)$  with depth  $> \varepsilon$  inside the verified range. *Close:* a proof of the candidate statement for the canonical completed-cell flow — the graded weakening of the cone-transport demand; note it does *not* by itself give ( $H_{\text{cof}}$ ): the graded statement bounds depth, not sign. No marker move; no RH statement.

**The code-layer compressions of round 31 (CFIN.UNIQUE.01, ARF.VACUUM.SYNDROME.01, OVOID.DECODER.01, ARF.BENT.CSS.01, PRIME.RELATION.MANGOLDT.01).** *Recorded* 2026-08-07 (evidence lines, not contracts; exact identity grade; no marker moves). The uniqueness theorem: every admissible compiler tuple is isomorphic to  $C_{\text{fin}}$  (census 14400, one  $Sp(4, 2) \times S_5$  orbit,  $\text{Aut}(C_{\text{fin}}) \cong C_6$ ; the strict terminal-object reading fails honestly — unique up to *non-unique* isomor-

phism; P2 stays narrowed to the v799 residual R1') ([verification/v849\_cfin\_unique\_cofinal\_lean.py], 22/22). The Arf–vacuum syndrome theorem and the ovoid decoder: the two-bit syndrome table 15/20/45/60 on the Hamming code with the local Hecke theorem  $12/16 \rightarrow 28/-4$  at all 15 points, the code pencil with the canonicity lemma and the vacuum bit as the  $\text{RM}(2, 4)$  parity bit;  $\text{SNF}(N) = \text{diag}(1^{10}, 0^5)$  torsion-free with the integral decoder  $N^+ = N^T/4 - J/36$  and the  $S_6$  standard representation; the three-ring correspondence ([verification/v852\_arf\_vacuum\_ovoid.py], 24/24 + 23/23). The bent CSS pair and the commutator theorem:  $\hat{s}_q = -4s_q$ , the  $(16, 6, 2)$  difference set,  $[[15, 5, 3]] \rightarrow [[16, 4, 4]]$  with the vacuum transformation typed, the  $-4$  triple point (Hecke = Gauss = Walsh =  $a_3$ ); and  $\mathcal{L} = -M[D, Z] = -[D, \log Z] = T(\Lambda)$  exact with the four-comb discipline at operator level and the carrier bridge (typed honestly as classical incidence algebra) ([verification/v853\_bent\_css\_mangoldt.py], 22/22 + 30/30). Full statements and keyboxes in Paper 1; audit cards in Paper 3.

**The grade barrier as a class closure (FORM.PRIME.GRADE.NO\_GO.01 [F]).** *Registered* 2026-08-08, FORMAL (TfptCarrier/GradeNoGo.lean, 13 declarations (10 theorems), kernel-checked — lake build green at 3415 jobs, no sorry, no native\_decide, axioms propext/Classical.choice/Quot.sound only, verified by #print axioms on all ten theorems; numeric mirror [verification/v859\_grade\_no\_go\_elevator.py] part A; carried in the shipped Lean manifest at 87 files). *Proven:* the core no-go `grade_no_go` — a 1-homogeneous map and a 2-homogeneous map that agree on a scaling ray agree with 0 at the base point (evaluate at  $t = 1$  and  $t = 2$ :  $2A(m) = 4A(m)$ ; two lines — a *class* closure, not a construction failure); `gram_two_homogeneous` and `grade_no_go_gram` (Gram squares of linear maps are 2-homogeneous, so exact matching annihilates target *and* factor: grade 1 cannot live inside grade 2); the affine trilemma — `affine_gram_expansion` (the exact split into constant + cross + quadratic components, the cross term the *only* grade-1 slot a Gram square owns), `affine_gram_tax` (the representation error against the pure cross term is *exactly* the sum of two PSD prices: the tax), and `affine_grade_no_go` (exact affine matching forces target, background *and* factor to vanish); and the elevator lemma `tangent_psd_on_kernel_null` — derivatives of PSD families at kernel-null directions are positively signed: *rates, not amplitudes*, the only grade-1 mechanism the no-go does not exclude, proven once, abstractly. *The unifying diagnosis:* the four graveyard instances (the Schur–Gram tax, the covariance grade gap, the wedge uniformity failure, the phase-lever closure) are instances or relatives of this one law; the measured exchange rate is  $10^4$ – $10^5$ . *Honesty notes:* the no-go excludes exact Gram matching on scaling rays and nothing more — denominator representations died separately by measurement (cited); non-Gram mechanisms are outside the class closure; the trilemma exhaustiveness outside the affine family stays an informal doc-level reading, typed. No continuum theorem, no RH statement, no arithmetic identification.

**The conditional route to the cofinal hypothesis (FORM.PRIME.COFINAL\_CURRENT.01 [F]).** *Registered* 2026-08-08, FORMAL (TfptCarrier/CofinalCurrent.lean, 13 declarations (7 theorems), kernel-checked — lake build green at 3415 jobs, no sorry, no native\_decide, axioms propext/Classical.choice/Quot.sound only, verified by #print axioms on all seven theorems; numeric mirror [verification/v856\_connected\_current\_lean.py] part C; carried in the shipped Lean manifest at 87 files). *Proven:* the conditional theorem `cofinal_current` — *ConnectedTail* (every tail block PSD) + *PositiveHeadCell* (an open phase cell with a uniform Rayleigh floor) + *PhaseRecurrence* (cofinal visits)  $\Rightarrow (H_{\text{cof}})$ , with the index sequence *extracted* from the recurrence; the composed corollary `cofinal_current_weil` through the minimal H theorem: the three hypotheses + per-element form convergence force the limit functional nonnegative on the whole dense family — no hidden wall behind the two analytic lemmas; the  $\delta = 0$  relaxation (the chain consumes only PSD on the cell; the strict floor is a stronger input than required, typed) and the Rayleigh reading; and the counterexample `connectedTail_not_sufficient` (a PSD-tail split with everywhere-negative head admits no cofinal ladder — the hypotheses are jointly necessary). *Provenance typed, not proven:* hypothesis 1 is what the connected-covariance construction delivers structurally (PRIME.RELATION.CONNECTED\_COVARIANCE.01); hypothesis 2 is what the phase-cell census *measured empty at the deployed frame* (PRIME.COFINAL.PHASE\_CELL.01 — the frame awaits

a head with a nonempty cell); hypothesis 3 is the classical Kronecker/Weyl citation (arithmetic input: log-prime independence, the two-line UFD lemma warded in the probe). *The honest boundary*: nothing proves any of the three hypotheses; a continuation contract, not a discharge; no continuum theorem, no RH statement, no arithmetic identification.

**The Krein-defect theorem frame (FORM.PRIME.KREIN.DEFECT.01 [F]).** *Registered* 2026-08-08, FORMAL (TfptCarrier/KreinDefect.lean, 14 declarations (10 theorems), kernel-checked — lake build green at 3416 jobs, no sorry, no native\_decide, axioms propext/Classical.choice/Quot.sound only, verified by #print axioms on all ten theorems; numeric mirror [verification/v861\_krein\_normalform\_lean.py] part C; carried in the shipped Lean manifest at 88 files). *Proven*: the defect representation `defect / formAt_defect` — the quadratic form of  $B_+^H B_+ - B_-^H B_-$  is  $\|B_+ f\|^2 - \|B_- f\|^2$  exactly; the main theorem `defect_psd_of_contraction` —  $B_- = C B_+$  with the Loewner certificate  $(1 - C^H C) \succeq 0$  forces the defect PSD, via the exact factorization  $B_+^H (1 - C^H C) B_+$  and congruence; the converse `douglas_contraction_of_defect_psd` (full-rank case proven,  $C := B_- B_+^{-1}$ ; the singular case deliberately the classical Douglas 1966 citation — the range bookkeeping not formalized, typed); `SourceContractor` — the *named hypothesis*: contractor + factorization + *independent* norm certificate as data a consumer must supply from the source side, with the circularity warning typed (a target-computed  $C$  — eigendata or the converse's pseudoinverse — is a reformulation of the positivity, not a proof); and the cofinal composition `krein_cofinal / krein_cofinal_weil` — contractors on a cofinal ladder yield  $(H_{\text{cof}})$  (index extraction from the recurrence) and compose through the minimal H theorem to Weil nonnegativity on the whole dense family: no hidden wall behind the contractor certificates. *Provenance typed, not proven*: the contractor hypothesis is exactly what the Krein source-algebra search must deliver; [verification/v866\_source\_contractor\_formula.py] delivers its closed *formula* and proves the absolute half of the toolbox cannot certify it — the certificate demand is the registration below. *The honest boundary*: nothing here proves the contractor exists; no continuum theorem, no RH statement, no arithmetic identification.

**The phase-aware contractor certificate (PRIME.KREIN.CONTRACTOR.01 [O]).** *Registered* 2026-08-08, OPEN ([verification/v866\_source\_contractor\_formula.py], 10 + 8/8 with the one frozen-honest FAIL kept and pattern-gated, SYLVESTER12-ANGLES-FAIL + NORM-REFORMULATION; with [verification/v861\_krein\_normalform\_lean.py], [verification/v862\_defect\_polar\_weld.py], [verification/v864\_softport\_kappa\_law.py], [verification/v867\_radau\_conditioning.py]). *The demand, verbatim*: supply a *target-free* certificate that the explicit closed-form source contractor  $C = W_- F G_+^{-1} F^H W_+$  (the factorization warded at  $1.4 \times 10^{-14}$  incl. blind holdouts; density roots and window geometry only) is a contraction,  $\|C\| \leq 1$ , uniformly along the deployed cofinal ladder — equivalently the Loewner comparison  $G_- \preceq G_+$  of the two weighted frame Grams, equivalently  $K \succeq 0$  (the Douglas equivalence, machine-verified) — with the certificate's constants independent of the measured target signs. *What is proven about the tool space (the registration's teeth)*: (i) *no absolute-value bound can succeed* — the Perron value  $\rho(|C|)$  is the infimum over all absolute Schur/Hilbert-type tests and it exceeds 1 on every heavy rung ( $2.363 \rightarrow 3.103$ , growing with depth): any valid certificate must use the *phases* of the kernel; (ii) quadrature certifies  $s > 0$  pointwise on every reachable rung but the depth grows as  $m^* \sim \sqrt{\text{cond}(G)}$  and the conditioning is *intrinsic* — a uniform certificate cannot be a fixed-depth quadrature family; (iii) no source word of *any* length reaches  $C$  (the monomial closure theorem) — the certificate cannot be a word bound; (iv) the discharge target is kernel-checked: per-rung `SourceContractor` instances on a cofinal ladder compose through `krein_cofinal_weil` to Weil nonnegativity. *The named candidate shape*: an oscillatory-integral / stationary-phase bound on the entries of the weighted frame kernel, strong enough to beat 1 by the measured  $e^{\text{slope} \alpha}$  margin (fit-free defect law, slope  $-2.40$ ). *Kill*: a rung with  $\rho(|C|) \leq 1$ ; a proof that every phase-aware bound of the named shape degrades to an absolute bound on the deployed kernel class; a measured break of the factorization on any new rung. *Close*: the certificate itself — it would discharge the `SourceContractor` hypothesis and with it  $(H_{\text{cof}})$  for the defect ladder. [2026-08-08 round-34 note ([verification/v869\_cdcore\_clarkphase.py] — [verification/v872\_damping\_compensation.py] + the Prüfer preregistration): the kernel is named —

the closed-form contractor's reproducing kernel  $K_+$  is the Christoffel–Darboux kernel of the explicit source measure  $\mu = \sum [2 \sin^2(\varphi/2) d_+(\varphi)/L] \delta_{2 \cos(D\tau)}$  (rank-2 displacement collapse with a  $10^{13}$  gap on all rungs incl. blind holdouts; firewalled chain-only reconstruction at  $\sim 10^{-13}$ ): the demand now reads — control *one* source Jacobi chain and the arm weights. *The compensation identity*:  $I - C^*C = T_1 + T_2$  exactly ( $T_1 = I - U^*U$  indefinite Jacobi cross-geometry;  $T_2 = U^*(I - D_-^2)U$  PSD Christoffel damping) with the exact global minimizer equal to the soft direction and the isolated scalar  $s = \kappa\tau$  — the demand localized to the two-term inequality at one named direction, both terms source-built. *The closures (stop-list entries)*: the Gauss-node unitary core does not exist (co-isometry defect  $O(1)$ ; (H-U) the open hypothesis,  $D_-^G \leq 0.997041$  the surviving asset — [verification/v870\_christoffel\_gauss\_frame.py]); classical pole transformations at  $x_{\text{pole}} = 2 \cosh(D/2)$  are dead (typed residue: small-negative-mass Geronimus at 3.5–8.1%, no  $\lambda(D)$  law); finite-rank cross-defect corrections are dead (rank grows linearly with  $h$ , Pearson +0.996 — [verification/v871\_pole\_uvarov\_crossdefect.py]);  $\mu_4$ -offset Clark kernels are dead (the equal-phase law). *The preregistration (the contract's next test, preregistered pre-execution)*: PRIME.PRUEFER.COMPENSATION.01 — source-only Prüfer phases of the two arm chains, exactly 16 fixed phase cells (boundaries  $k\pi/8$ , entrywise pairing primary, danger families predeclared), Cotlar–Stein decision rules and all success/kill criteria frozen *before* any deployed-data run, deep holdouts kz 90/116/142/177/243 never used by the CD wave, the decisive analysis locked behind the operational unfreeze flag, 8/8 synthetic freeze wards green; committed byte-exact with round 34: SPEC SHA-256 4621b89958531d6d3baae9fe762dbd10d23dc22eed77a57aa1c6b179a7440811, FILE SHA-256 93208a1b2edebde31388a87192055c9bf3e5a3de29af9fb3e34ddc904825d9fd (experiments/tfpt-discovery/pruefer\_compensation\_probe.py; any modification invalidates the preregistration). Demand, kill and close conditions unchanged; still [O.] [2026-08-08 round-35 note ([verification/v873\_pruefer\_cotlar\_decision.py], the executed preregistration): the contract was executed as committed — the unfreeze flag created *after* the round-34 commit 526ca3eb (audit line committed), the two mechanical fixes documented verbatim (the  $h \leq 900$  battery filter the spec's own text mandates; the overflow renormalization, atan2-invariant, validated at  $\leq 6.6 \times 10^{-4}$  against a 60-digit mpmath reference), the spec literal SHA byte-identical throughout, the executed FILE SHA-256 15d35b26c652ef4144df713cafba7e049e2dcb9bb7a168462cd65735f5ffb4cc — and it returns COTLAR-GROWING by its own frozen rules: danger geometry real (capture 0.98–0.61 of  $\lambda_{\min}(T_1)$ , source-only), local dominance and envelope decay FAIL (the envelope bimodal at the predeclared  $\pi$ -resonance cells), the battery flat but the frozen deep holdouts exploding (slopes 0.209/0.408  $\geq 0.20$ , both channels; kz 142:  $S_U$  32.07/  $S_C$  58.78), the contrast bars passing, kills 2 + 5 firing, RUN 3 correctly withheld. The blockwise Cotlar–Stein phase-cell route is closed and stop-listed; the surviving demand is a non-stationary phase bound with the  $\pi$ -resonance as the known critical structure and the  $\alpha$ -not- $h$  growth law as the typed constraint (the explosion tracks kz-depth, not window size). The first fully preregistered uniformity test of the program, executed and booked exactly as frozen; still [O.] [2026-08-08 round-36 CORRECTION note ([verification/v877\_complete\_comb\_reversal.py], the complete-comb re-decision): the round-35 note above is corrected, dated — the RUN-2 deep-holdout data was *comb-truncated* (core.build\_window silently caps the atom table at  $\text{ATOM\_MAX} = 4 \times 10^5$ ; kz 142/177/243 demand support to  $4.4 \times 10^5/8.0 \times 10^5/1.8 \times 10^6$ ; PRIME.DEEPALPHA.SIGN.01 certifies the complete-comb floors strictly positive). The v1 verdict was correct *on the data as measured* and v873's booking stands as the record of the preregistered v1 execution (spec SHA 4621b899..., executed file SHA 15d35b26...); but on complete combs the *identically* frozen v1 rule returns BOUNDED in both channels (v2 spec SHA 5fd6bf61...; kz 142:  $S_U$  32.07  $\rightarrow$  5.65,  $S_C$  58.78  $\rightarrow$  4.10; slopes 0.006/0.007) and RUN 3 passes ( $K_1/K_2$   $h$ -stable on all five holdouts). *The stop-list entry is amended*: “blockwise Cotlar–Stein on Prüfer phase cells” is removed — the class is *not* closed; the route reopens as PRIME.PRUEFER.COMPENSATION.02 with the remaining gap typed (the Cotlar constant  $B_C \sim 3.9$ –4.3 vs the needed 1; the cell-local defects  $\sum_r \varepsilon_r \sim 7$ –10 — the RUN-1 anatomy was in-cap and stands); the  $\alpha$ -not- $h$  law is retired (it was the truncation's signature). The  $\pi$ -resonance critical structure and the discrimination stand; the surviving demand is now carried in

embedding coordinates by PRIME.CARLESON.PRIME.01 [O] (registered below). Demand, kill and close conditions otherwise unchanged; still [O].] Existence is *not* claimed; no RH statement.

**The 210 canonicity audits of round 33 (E8.DIVISOR210.CANONICITY.01).** *Recorded* 2026-08-08 (evidence lines, not contracts; no marker moves). The label register  $\mathbb{F}_2^4$  is the divisor lattice of  $210 = 2 \cdot 3 \cdot 5 \cdot 7$  (Walsh–Hadamard carries the lattice Möbius function; the vacuum column is the rank-one Redheffer closure with  $\det = \prod_p (1 - 1/p) = 8/35$  exact, [verification/v863\_redheffer\_colligation.py]); the canonicity guard quantifies the selection ([verification/v868\_divisor210\_audits.py],  $36/36 + 12/12 + 15/15$ ): the Boolean layer is measured generic (210/210), the two Euler determinants pin  $\{2, 3, 5, 7\}$  uniquely, the anchor prime 2 is forced by ramification, and after the crossref exactly two gauge classes remain — the family-cycle chirality, then *proven gauge* by two exact no-go wards (the orientation functional vanishes identically via the Möbius complement  $d \mapsto 210/d$ ; quadratic readouts provably blind by transposition); the  $17 = 12+3+1+1$  moving-sector identification honestly buried at space level (principal angles  $\sim 90^\circ$ ); the 6-vs-7 honesty note stands untouched. Full statements and keyboxes in Paper 1; audit cards in Paper 3.

**The Prime Carleson embedding contract (PRIME.CARLESON.PRIME.01 [O]).** *Registered* 2026-08-08, round 36 ([verification/v876\_carleson\_polygon\_wave.py]; the ward making the reduction machine-exact at rel  $1.5 \times 10^{-13}/1.3 \times 10^{-15}$  on kz 9/13). *The chain:* for the canonical window family, the Krein floor  $\tau_h > 0$  is *equivalent* to the Carleson embedding of the negative arm measure into the positive arm measure at polynomial degree  $h - 1$  —  $\|C_h\| \leq 1 \Leftrightarrow G_- \leq G_+$  (exact, PRIME.SOURCECONTRACTOR.NORM.01)  $\Leftrightarrow \int |P|^2 d\tilde{\nu}_- \leq \int |P|^2 d\tilde{\mu}_+$  for all  $P \in \mathcal{P}_{h-1} \Leftrightarrow M(\tilde{\mu}_+ - \tilde{\nu}_-) \geq 0$ , with the embedding constant  $\lambda_{\max}(M_-) = \|C_h\|^2 = 1 - \tau_h$ . *The demand:* a proof of the embedding uniformly along a window subfamily cofinal in the mesh-refinement order would discharge SourceContractor (FORM.PRIME.KREIN.DEFECT.01) and, through krein\_cofinal\_weil, hypothesis ( $H_{\text{cof}}$ ). *The typed constraints a candidate must respect:* the certificate must be phase-aware (the Perron closure of v866); the  $\pi$ -resonance is the known critical structure and it is *thick* — not removable by excision (PRIME.PI.RESONANCE.ANATOMY.01); the polygon/two-gauge compression is the *wrong* compression of the Loewner order (PRIME.PHASE.POLYGON.01); the constructive side is the residual quadrature at full degree (PRIME.RESIDUAL.QUADRATURE.01) under the complete-comb constraint (PRIME.TRUNCATION.AUDIT.01). Arrow (iv) (a prefix-local scalar reduction) is typed NEEDS PROOF. *Kill/close:* closes if the embedding is proven on a cofinal family, or if a certified negative canonical window family is exhibited (none exists: PRIME.DEEPALPHA.SIGN.01). Existence is *not* claimed; no RH statement. [2026-08-09 round-37 note ([verification/v879\_envelope\_partition\_closure.py]): the constraint set is extended, dated — the partition class is closed entirely (PRIME.PARTITION.CLOSURE.01: Prüfer cells at every tested resolution with no predeclared accounting below 3.568 vs the bar 1; honest tight wave-packet frames on both natural phase spaces with diagonal mass  $\leq 18\%$  under any pairing and a growing Schur constant 4.9–9.2), so no partition-based accounting can certify  $\|C\| \leq 1$ : the embedding certificate must be non-decompositional — eigen-enclosure, Loewner-order or passivity arguments that treat  $M(\tilde{\mu}_+ - \tilde{\nu}_-)$  whole, all three measured territory. The new carried assets: the exact tridiagonal multiplication identity on the odd frame (CD telescoping at  $1.3 \times 10^{-13}$ , rank-2 boundary defect at  $k = 0$ ), the derived entrywise bound with  $h$ -uniform source constants ( $C_{\text{int}} \in [1.232, 2.328]$  from  $h = 168$  to 1445 — the route’s first analytic uniformity statement), and their source-law persistence as  $h \rightarrow \infty$  as the named open object. Demand, kill and close conditions otherwise unchanged; still [O]; no RH statement.]

**The Carleson wave of round 36 (PRIME.PHASE.POLYGON.01, PRIME.RESIDUAL.QUADRATURE.01, PRIME.PI.RESONANCE.ANATOMY.01).** *Recorded* 2026-08-08 (evidence lines with typed honest FAILs; no marker moves; [verification/v876\_carleson\_polygon\_wave.py],  $6/8 + 7/8 + 5/5$ , the three frozen-honest FAILs pattern-gated). The polygon: Lemma A exact (the difference of the two source canonical systems is a signed rank-one phase sum; round-trip  $2.4 \times 10^{-12}$ ), but  $S_0(m) \geq |S_2(m)|$  breaks on all 47 rungs at  $2\theta \approx 1.5\pi$  — the two-gauge compression is the *wrong* compression of the Loewner order, typed with both readings; holonomy null with inverted contrast; the deep-holdout

rows retype NUMERIC-VOID per the round-36 truncation audit (battery verdict unaffected). The residual horizon law: the difference Jacobi chain (row-scaled Wheeler + mp escalation past the 53-bit noise wall) realizes the positive representing measure  $\rho_h$  at *full degree*  $h$  on all 44 floor-positive rungs ( $k^*/h$  median 1.000, slope 0.974,  $r_h = 0.998$ ); a single atom sits just beyond the pole port (mass beyond the fold is *required*: the Radau-anchored functional breaks early, genuine); discrimination 7–9×; the three sub- $h$  stalls were exactly the truncated windows — a validated sign detector. The resonance anatomy: the resonant set is *thick* (a constant  $\sim 1/5$  of the minus nodes; excision gains nothing; the soft direction disjoint) — the wall is phase-distributed at this partition resolution; the  $\alpha$ -law addendum confounded by truncation, typed. Full narrative in the prime-front paper; audit card in Paper 3.

**The artifact resolution and the Cotlar reversal of round 36 (PRIME.DEEPALPHA.SIGN.01, PRIME.PRUEFER.COMPENSATION.02, PRIME.TRUNCATION.AUDIT.01).** *Recorded* 2026-08-08 (the round's headline; [verification/v877\_complete\_comb\_reversal.py], 8/8 + 9/9). The deep-alpha negative floors at kz 142/177/243 were *comb-truncation artifacts* of the deployed atom table (ATOM\_MAX =  $4 \times 10^5$ , silent cap; the windows demand support to  $1.8 \times 10^6$ ): with the complete canonical comb the certified corner enclosures are strictly positive ( $+1.06 \times 10^{-5} / +1.52 \times 10^{-5} / +1.20 \times 10^{-5}$ ), no negative canonical window exists, the certified ladder's end at kz 121 was the atom-table *data* boundary, and ( $H_{\text{cof}}$ ) is unaffected; the silent truncation is typed an infrastructure bug with the registered follow-up (the builder must fail loudly beyond its table). Under the *identically* frozen v1 bars the complete-comb deep sums collapse into the battery band and both channels return BOUNDED (slopes 0.006/0.007); RUN 3 unlocks and passes ( $K_1/K_2$   $h$ -stable on all five holdouts) — the route reopens, with the remaining gap typed ( $B_C \sim 4$  vs 1; local defects 7–10). The truncation audit re-types every affected finding, dated (pruefer v1 confounded  $\rightarrow$  re-decided; residual strengthened; resonance  $\alpha$ -law void; polygon holdouts numeric-void; the other waves unaffected). The dated corrections live on the PRIME.KREIN.CONTRACTOR.01 and PRIME.FLOOR.PAIRCORR.01 blocks above; v873's booking stands as the record of the preregistered v1 execution. No RH statement.

**The partition closure of round 37 (PRIME.PARTITION.CLOSURE.01; dated notes on PRIME.CARLESON.PRIME.01 and PRIME.PRUEFER.COMPENSATION.02).** *Recorded* 2026-08-09 (the compact closing round of the reopened Cotlar route; [verification/v879\_envelope\_partition\_closure.py], 7 checks with the two frozen-honest FAILs T1.4/C.1 pattern-gated at exit 1 + 6/6; the contract chain 4621b899...  $\rightarrow$  5fd6bf61...  $\rightarrow$  27d9f0a3...  $\rightarrow$  2199e9b1... machine-verified at run time). The analytic core, saved: multiplication by  $\cos \theta$  is *exactly* tridiagonal on the odd frame (mirror-closed boundary, single rank-2 defect at  $k = 0$ ; the CD telescoping identity at  $1.3 \times 10^{-13}$  on every rung incl. all five complete-comb holdouts), and the derived pointwise bound carries  *$h$ -uniform source constants* ( $C_{\text{int}} \in [1.232, 2.328]$  from  $h = 168$  to 1445, slope +0.106;  $\sqrt{(G_+^{-1})_{00}} \in [0.661, 0.899]$ ) — the route's first analytic uniformity statement, at entrywise grade. The structural closure (the honest FAILs are the finding): the cell-envelope decay is *cancellation-carried*, provably not entrywise-derivable (every Prüfer cell contains near-diagonal pairs; the analytic envelopes blow up as  $h^2$ ); no predeclared accounting below 3.568 (bar 1);  $\sum_r \varepsilon_r = 6.9$ –12.9  $\alpha$ -uniform on clean complete-comb data; and the honest tight wave-packet frames (both natural phase spaces;  $\|B\| = \|C^G\|$  exact) leave diagonal mass  $\leq 18\%$  under any pairing with a Schur constant growing  $4.985 \rightarrow 9.222$  (slope +0.295), while the Gabor envelope does decay in phase space (the first partition to show decay) and the discrimination is massive (truth 4.985 vs Epstein 128.7 vs scramble 9006.6). *The consequence*: the cancellation is *atomically global* — no partition-based accounting (cells, packets, any frame decomposition of this family) can certify  $\|C\| \leq 1$ ; the certificate demand of PRIME.CARLESON.PRIME.01 is retyped *non-decompositional* (eigen-enclosure / Loewner-order / passivity — all three measured territory); wave-packet partitions and entrywise envelope bounds are stop-listed on the PRIME.PRUEFER.COMPENSATION.02 block. The surviving assets: the tridiagonal identity, the  $h$ -uniform derived constants, their source-law persistence (the named open object). Full narrative in the prime-front paper; audit card in Paper 3. No RH statement.

**The rounds-38/39 sandbox campaign (the operator geometry and the source law of the wall).** *Recorded* 2026-08-09 (a campaign note; 25 frozen sandbox probes, every spec amendment fail-first on record; the promoted verification modules `[verification/v880_finite_anchor_theorems.py]`, `[verification/v881_carleson_port_geometry.py]` and `[verification/v882_port_source_mellin.py]` embed their discovery probes byte-exact). The campaign gives the Carleson wall its exact operator geometry (rank-2 Jacobi displacement, gauge-exact Carleson Gram with the classical testing diagonal, exact port Schur reduction — the port *is* the wall — and the preserved IKS integrable class), its universal source law (the  $\sqrt{n/X}$ -uniformized prime measure, the PNT edge law  $1 - e^{-1/2}$  as the  $\sim 39\%$  port mass, the Mellin–Cauchy port symbol, the closed criticality budget), and the hardness decision (`PRIME.ERRORTERM.SCALE.01` below); the six ledger-only items of the campaign are registered as the blocks that follow. Full narrative in the prime-front paper; audit cards in Paper 3. No marker moves; no RH statement.

**The error-term scale decision (`PRIME.ERRORTERM.SCALE.01`, decided).** *Recorded* 2026-08-09, DECIDED (sandbox probe `errorterm_demand_curve_probe.py`; the demand-curve measurement behind the round-39 hardness statement). *The finding:* the wall margin equals injected off-critical perturbation *energy* on the  $\tau$  law exactly (ratios 1.00–1.10 across the tested grid); the perturbation amplitude scales as  $A^* \sim \sqrt{\tau}/\delta^2$  under the frozen convention  $\Delta\lambda = O(A^2\delta^4)$ ; and the kill localizes through the port block. *The decision:* the reduction chain of the wall is an *exact reformulation at RH strength* — there is no unconditional slack anywhere in it; the identities are exact, and the one-sidedness is open at exactly that strength. Recorded as a hardness decision, not a contract demand. No marker move; no RH statement.

**The port limit contract (`PRIME.PORT.LIMIT.01` [O]).** *Registered* 2026-08-09, OPEN (sandbox probes `port_limit_operator_probe.py`, `port_cocycle_window_probe.py`, `weyl_m_criticality_probe.py`). *Measured:* the compressed operator-norm convergence has slope  $-2.80$  carrying the wall exactly;  $\lambda_{\max}$  increases to 1 from below; the atom-at-edge Weyl- $m$  class has  $\sigma \rightarrow -1$  and  $w_1 \rightarrow 0.72$ . *The demand:* the fixed-space limit theorem — the existence and identification of the limiting port operator on a fixed space, with the wall carried through the limit. *Kill:* a measured break of the convergence slope or a rung family on which the compressed operator fails to carry the wall. No marker move; no RH statement.

**The scalar port inequality (`PRIME.PORT.SCALAR.01` [O]).** *Registered* 2026-08-09, OPEN (sandbox probe `port_scalar_schur_probe.py`). *Measured:* the wall is *one* scalar inequality,

$$\sigma_h = [1 - \check{\nu}_{m^*} K_h(y^*, y^*)] - b^\top (I - E_{\text{rest}})^{-1} b \geq 0,$$

exactness  $6.2 \times 10^{-11}$  on 42 rungs, with  $\sigma_h$  tracking  $\tau_h$ . *The demand:* one-sidedness — a proof that  $\sigma_h \geq 0$  along a cofinal window family. *Kill:* a rung on which the scalar reduction fails at certificate grade. No marker move; no RH statement.

**The port cocycle window (`PRIME.PORT.COCYCLE.01` [O]).** *Registered* 2026-08-09, OPEN (sandbox probe `port_cocycle_window_probe.py`). *Measured:* a fixed  $12 \times 12$  window carries the wall exactly, with compressed convergence slope  $-2.80$ . *The demand:* the closed-form limit of the window cocycle. *Kill:* a depth at which the fixed window stops carrying the wall. No marker move; no RH statement.

**The Case sum-rule contract (`PRIME.CASE.SUMRULE.01` [O]).** *Registered* 2026-08-09, OPEN (sandbox probes `case_sum_rule_probe.py`, `testing_mnt_law_probe.py`). *Measured:* the deployed Jacobi chains are measured Killip–Simon class. *The demand:* prove a Case C0/C1 sum rule implying the unconditional testing bound  $T_h \leq 1$ ; the target is exactly the closed criticality budget of round 39 (PNT growth compensated by window geometry plus Christoffel). *Kill:* a measured departure of the deployed chains from the Killip–Simon class. *Route update (2026-08-09, round CVIII) — retyped conditional:* the free-tail bridge is numerically delivered (sandbox `freetail_case_bridge_probe.py`, 8/8: the finite Killip–Simon P2 identity exact at  $5 \times 10^{-11}$ ; the P2 energy  $h$ -uniformly bounded 0.63–0.75 over all 42 rungs), *but* the route is retyped conditional by `[verification/v889_route_decisions.py]`: PNT-INSUFFICIENT (the smooth reference violates the

port margin on every rung) and HOMOTOPY-INDEFINITE-ARITHMETIC (the required input is the signed pair-correlation functional). The demand becomes: state and prove the weighted prime-pair correlation inequality, or accept conditionality. No marker move; no RH statement.

**The port Riemann–Hilbert contract (PRIME.PORT.RIEMANNHILBERT.01 [O]).** *Registered* 2026-08-09, *OPEN* (sandbox probes `port_riemann_hilbert_setup_probe.py`, `mellin_model_operator_probe.py`). *Measured*: the discrete  $2 \times 2$  IKS Riemann–Hilbert problem of the dressed port operator is *well-posed* — the two generators are explicit and the resolvent closure holds at  $1.2 \times 10^{-12}$ . *The demand*: nonlinear steepest descent for  $h \rightarrow \infty$ . *The calibration constraint (load-bearing)*: the smooth PNT continuum model *violates* the floor at finite depth ( $\tau_{\text{model}} -1.26 \rightarrow -0.11$ ) — the prime discreteness is load-bearing, and the unperturbed problem must keep the lattice. *Execution note (2026-08-24, rounds 244–255)*: the finite formulation this demand must consume is now certified — the bordered end-object is a canonical  $2 \times 2$  RHP  $\times$  a one-dimensional budget cocycle with exact evaluation infrastructure (`[verification/v958_bordered_tau_readout_dictionary.py]`); the irreducible- $3 \times 3$  steepest-descent variant is dropped without replacement, and the admissible target form is an *integrated* statement over the compressed CD readout (no uniform  $Y_n$  parametrix) — the campaign continues as PRIME.PORT.RHP.FULLSOURCE.BASEFIBER.01 (the block below). This umbrella stays [O]. *Kill*: a failure of the IKS well-posedness on a deployed rung. No marker move; no RH statement.

**The rounds-40/41 flow corridor (the determinant chain, the parametrix decision, the certified head, and the complete-comb anatomy).** *Recorded* 2026-08-09 (a campaign note; nine frozen sandbox probes, every spec amendment fail-first on record; the promoted verification modules `[verification/v883_tau_chain_parametrix.py]`, `[verification/v884_certified_head_positivity.py]`, `[verification/v885_ladder_flow_anatomy.py]` and `[verification/v886_register_gapcode_null.py]` embed their discovery probes byte-exact). The campaign gives the wall’s sign an exact finite determinant chain ending in a candidate tau function, the decided parametrix constraint (the Euler-product comb is globally rigid — every smoothing-based parametrix is dead), the first rigorous positivity statement of the program (exact integer Sylvester certificates on the three head rungs), the complete flow/induction anatomy of the ladder (degree-2 Lax closure; no monotone one-prime induction in any insertion order — positivity is a property of the complete comb), and the register-gapcode disenchantment. The registrations and updates follow as the blocks below. Full narrative in the prime-front paper; audit cards in Paper 3. No marker moves; no RH statement.

**The tau contract (PRIME.PORT.TAU.01 [O]).** *Registered* 2026-08-09, *OPEN* — the finite- $h$  determinant chain delivered (`[verification/v883_tau_chain_parametrix.py]`, 5/5, TAU-CHAIN-EXACT). *Exact*: the Schur determinant identities  $\sigma_h = \det(I - E)/\det(B)$  and  $\det(I - E) = \det(I - R)\det(I - D_P)$  are machine-exact ( $4.8 \times 10^{-13}$  in log space) with positive bulk factors, hence  $\text{sign } \tau_h = \text{sign } \sigma_h$  is the sign of the wall with  $\tau_h := \det(I - D_P)$ ; the Jacobi variational identity  $\frac{d}{ds} \log \det(I - sD_P) = -\text{tr}[(I - sD_P)^{-1}D_P]$  holds on the true family ( $10^{-6}$ ); the determinant wall indicator is false on Epstein and scramble. *Measured*: the  $\kappa/\rho$  readout —  $\text{slope}(\log \sigma) = -2.93$  vs  $2 \times$  the Mellin fluctuation slope  $-2.53$  ( $\sigma \sim \text{fluctuation}^2$  within 16%). *The demand (the symbolic half)*:  $\tau_h$  is the tau function of the discrete  $2 \times 2$  IKS Riemann–Hilbert problem of the v881 generators, for arbitrary  $h$  without numerics. *Execution note (2026-08-24, rounds 224–226)*: the demand is *executed at the finite-identity level* — Sylvester as a full characteristic identity, the  $s$ -dressed Schur family, the confluent diagonal, the IKS dressing identity with source-explicit dressed generators, the Fredholm/JMU variational identity in two independent times and the relative tau transfer, all machine-exact on the actual builders (TAU-IKS-EXACT, 18/18; `[verification/v955_tau_iiks_toda_dictionary.py]`); per the note-DLI claim splitting the finite theorem set is carved out as PRIME.PORT.TAU.FINITE.IKS.01 [E] (the block below) and the sign question stands alone as PRIME.PORT.TAU.NOPOLE.COFINAL.01 [O]; this umbrella stays [O] for the fully symbolic arbitrary- $h$  statement. *Kill*: a heavy rung on which the determinant chain fails at certificate grade. No marker move; no RH statement.

**The finite IKS/Toda theorem set (PRIME.PORT.TAU.FINITE.IKS.01 [E]).** *Registered and delivered* 2026-08-24 (rounds 224–226, notes DL/DLI/DLII;

[[verification/v955\\_tau\\_iiks\\_toda\\_dictionary.py](#)], probes embedded byte-exact, 18/18+15/15+17/17 first pass). *Exact*: the wall determinant is the exact IKS Fredholm tau function at the finite-identity level (Sylvester characteristic identity on the whole  $s$ -grid; the  $s$ -dressed Schur family  $\det(I - sE) = \det(I - sR) \det(I - sD_P(s))$  with the fixed-kernel alias named and blocked; the confluent diagonal  $\leq 3.2 \times 10^{-13}$ ; the IKS dressing identity  $\leq 7.9 \times 10^{-13}$  with the dressed port generators source-explicit in closed form; the Fredholm/JMU variational identity in two independent times; the relative tau transfer exact over  $-212.84 / -195.50$  log-unit collapses); the  $h$ -direction carries an exact *predictive* degree-1 Lax dynamics (rank-1  $h$ -step at  $2.9 \times 10^{-16}$ ; blind next-generator and tau-step predictions at  $3.4 \times 10^{-16} / 3.2 \times 10^{-11}$ ; zero curvature identically, world-blind; mpmath wards at dps 80/120); and the Toda dictionary holds:  $\tau_{w,n} = D_n(\mu - \nu)/D_n(\mu)$ ,  $r_n = h_n(\mu - \nu)/h_n(\mu)$ , the Hirota coefficient source-pure from a scaled signed Stieltjes recursion. *Typed honestly*: the sign is WALL-EQUIVALENT ( $\hat{\gamma}_n > 0$  for all  $n \iff$  quasi-definiteness of the signed defect measure  $\mu - \nu \iff$  the wall; MAIN carries all 184  $r_n > 0$ , the controls flip at  $n = 25/21/27$ ); the  $s$ -time has no fixed degree (CONDITIONING-ONLY) and across rungs there is no common carrier; the general- $h$  statements are classical theorems (Sylvester, Schur, CD confluence, IKS dressing) instantiated on the actual builders. *Kill*: any of the finite identities failing on a deployed rung. No RH statement.

**The tau sign contract (PRIME.PORT.TAU.NOPOLE.COFINAL.01 [O]).** *Registered* 2026-08-24, OPEN (the honest open half of the note-DLI claim splitting). *The demand*:  $\tau_h(1) > 0$  cofinally — equivalently, no Malgrange divisor of the  $s$ -family up to  $s = 1$ . *Equivalent coordinates* (all machine-exact dictionaries, [[verification/v955\\_tau\\_iiks\\_toda\\_dictionary.py](#)] / [[verification/v956\\_signedmoment\\_halffilling\\_duality.py](#)]): quasi-definiteness of the signed defect measure  $\mu - \nu$  through the full free moment window; all Christoffel scalars  $r_n > 0$ ; all Jacobi couplings  $\beta_n > 0$  in every one of the four source-pure coordinate paths; the positive orientation of the signed Stokes multiplier of the two-sided midpoint problem. *Measured seeds*: MAIN survives the entire free moment window plus  $O(1)$  forced-tail degrees (offsets 0/2/2/3/1) while the controls die at 11–15% of it;  $\lambda_{\max}(Q_{N_w}) = 0.999832..0.999999 < 1$ . *Kill* (of any mechanism candidate): failing the frozen micro falsifier — predict the forced tail 0/2/2/3/1 and the control flips 25/21/27 blind. No marker move; no RH statement.

**The parametrix decision (PRIME.PORT.LATTICE.PARAMETRIX.01, decided).** *Registered and decided* 2026-08-09 ([[verification/v883\\_tau\\_chain\\_parametrix.py](#)], 3/3+3/3, PARAMETRIX-MEASURED (FLUCTUATIONS-REQUIRED) + ZONES-MEASURED (BOTH-REQUIRED)). *Measured*: four worlds through the identical pencil pipeline on every heavy rung — true comb  $\tau = +1.7 \times 10^{-4}$  to  $+6.7 \times 10^{-7}$  (ward); lattice with fully smooth quadrature masses  $-78$  to  $-1706$ ; lattice with exact local  $\pm \frac{1}{2}$ -log-unit mass sums  $-78$  to  $-5537$ ; plain continuum  $-0.11$  to  $-1.26$  (ward) — and the zone refinement kills both ways (edge-only smoothing  $-3.7$  to  $-22.1$ ; interior-only  $-77$  to  $-7551$ ). *The decision* (a design constraint for every RH analysis of the wall): the exact  $\Lambda$ -mass-position pairing — the multiplicative Euler-product structure — is load-bearing in both zones; the main/unperturbed object must carry the full von Mangoldt comb; every smoothing-based parametrix is dead; the small parameters live exclusively in the flow/deformation direction (ladder order, isomonodromy/tau,  $s$ -deformation). No marker move; no RH statement.

**The certified head (PRIME.PORT.CERTIFIED.HEAD.01, delivered).** *Registered and delivered* 2026-08-09 ([[verification/v884\\_certified\\_head\\_positivity.py](#)], 5/5, CERTIFIED-HEAD; SPEC v2 amendments on record). *Proven*:  $\sigma_h > 0$  on kz 9 ( $h = 184$ ), kz 12 (151) and kz 13 (168) — the first rigorous positivity statement of the program: independent mpmath rebuild of the window lags (dps 60/100; conservative per-entry error  $\varepsilon_c = 10 \times$  dps-disagreement; cross-ward at  $1.1 \times 10^{-14}$ , which caught a sign bug in the rebuilt Arch near-branch), matrix shift by the total error budget ( $\sim 2 \times 10^{-18}$ ,  $\sim 14$  orders below the margins), then exact integer Bareiss/LDL on the  $Q = 10^{20}$  grid: all pivots positive (Sylvester)  $\Rightarrow K \succ 0 \Rightarrow \sigma_h > 0$  via the exact v866 equivalence; certified floors  $\lambda_{\min} \geq 3.633/3.327/3.842 \times 10^{-4}$ ; the Epstein control through the identical machinery refused at pivot index 10. *Typed honestly*: proven modulo the declared conservative transcendental-evaluation error model (validated-numeric standard); the ball-arithmetic closure and the 42-rung rollout

were the named follow-ups — both since delivered ([`verification/v887_certified_ladder_complete.py`], [`verification/v897_certified_interval_ladder.py`]: the informal error model is retired from the ladder). *Kill*: a certified pivot failure on a head rung under the frozen error model. No marker move; no RH statement.

**The leading-sign contract (PRIME.PORT.LEADING.SIGN.01 [O]).** *Registered* 2026-08-09, OPEN (the moonshot of the round-40 review). *The demand*:  $\sigma_h = \kappa\rho_h + o(\rho_h)$  with a closed-form  $\kappa > 0$  via a relative Case/Killip–Simon identity — the one-sidedness of the scalar wall as an asymptotic law. *Measured seeds*: the  $\rho$  candidate is the squared lattice fluctuation scale —  $\text{slope}(\log \sigma) = -2.93$  vs  $2\times$  the Mellin fluctuation slope  $-2.53$  (within 16%; [`verification/v883_tau_chain_parametrix.py`]); the parametrix decision forces the  $h$ -ladder direction —  $D_P(h)$  as sections of one arithmetic limit object,  $\kappa$  from the cocycle increment structure, not from comb deformations. *Kill*: a measured break of the  $\sigma \sim \text{fluctuation}^2$  law on the deep ladder. No marker move; no RH statement.

**The tail contract (PRIME.PORT.TAIL.01 [O]).** *Registered* 2026-08-09, OPEN (the final composition contract of the port route). *The demand*:  $|\sigma_h - \kappa\rho_h| \leq (\kappa/2)\rho_h$  for all  $h \geq H$  explicit (the tail half), plus the finite head by certified numerics, plus the Lean composition through `krein_cofinal_weil`. *Head status*: delivered for kz 9/12/13 (PRIME.PORT.CERTIFIED.HEAD.01, [`verification/v884_certified_head_positivity.py`]); the rollout to all 42 head rungs and the ball-arithmetic closure are the named next steps; the tail half waits on PRIME.PORT.LEADING.SIGN.01. *Promotion note (2026-08-09, round CVIII)*: the rollout is *complete* —  $\sigma_h > 0$  proven on all 42 reachable rungs, promoted in [`verification/v887_certified_ladder_complete.py`] (PRIME.PORT.CERTIFIED.LADDER.01; the 40-rung parent rollout stays sandbox, runtime-excluded); the remaining named steps are the ball-arithmetic closure and the Lean composition. *Interval note (2026-08-09, round CXV)*: the ball-arithmetic closure is *delivered* — [`verification/v897_certified_interval_ladder.py`] (PRIME.PORT.BALLLADDER.01) re-proves all 42 reachable rungs with rigorous interval-arithmetic shifts (15 exact-rational + 27 validated-precision) and retires the informal lag-evaluation error model; the remaining named steps are the Lean composition and the asymptotic tail. *Kill*: a certified negative  $\sigma_h$  on any head rung. No marker move; no RH statement.

**The isomonodromy-flow measurement (PRIME.PORT.ISOFLOW.01).** *Registered* 2026-08-09, measured ([`verification/v885_ladder_flow_anatomy.py`], 8/8, ISOFLOW-MEASURED (FLOW-OPEN)). *Measured*: the gauge-fixed IKS generator flow over 41 consecutive ladder pairs closes in degree 2 (residual medians 0.244/0.248 with  $\{Y^2f, Y^2g\}$ ; degree 1 fails at 0.498/0.457) — isomonodromy-compatible, gauge stable (consecutive cos median 0.943),  $\log \tau$  monotone  $-58.6 \rightarrow -396.6$ . *Typed honestly*: the controls close similarly — degree-2 closure is a structural precondition, not a discriminator; the arithmetic stays in the value. *The demand*: the closed-form degree-2 deformation polynomial in a well-conditioned coordinate (the LAX2 block below). No marker move; no RH statement.

**The window-ladder transfer (PRIME.PORT.LADDER.TRANSFER.01).** *Registered* 2026-08-09, measured ([`verification/v885_ladder_flow_anatomy.py`], 5/5, LADDER-MEASURED (PERTURBATIVE + NON-MONOTONE)). *Measured*: over 40 consecutive steps the first-order transfer is marginally accurate (median 0.458 at bar 0.5) — the step is first-order graspable — but raw  $h$  order is not a natural flow (only 52% of steps increase  $\lambda$ ;  $\|\Delta\|/\text{gap} = 594$ –2026: the accuracy comes from port concentration, not step smallness). *The typed lesson*: the induction needs a better order than  $h$ ; the nested one-prime ladder was the named candidate and is decided in the block below. No marker move; no RH statement.

**The one-prime ladder decision (PRIME.PORT.PRIMELADDER.01, decided).** *Registered and decided* 2026-08-09 ([`verification/v885_ladder_flow_anatomy.py`], 13/13, PRIMELADDER-MEASURED). *Measured (exact per-atom lag rows; linearity warded at 0; endpoint warded vs the true floor at  $2.5 \times 10^{-9}$ )*: the pure-Arch start is catastrophically subcritical ( $\tau(0) = -9.9 \times 10^{13}$  at kz 9;  $G_+$  not even PD at kz 13/26/40); insertion crosses into positivity only at  $u^*/U = 0.991$ –0.999 in every insertion order; the edge-first ordering reaches stable positivity only at 100% of the atoms

(the  $\sim 39\%$  edge mass alone does not carry the wall); per-atom increments 54–64% positive; the deterministic mass scramble drops the endpoint to  $-474.7$  (control fires). *The decision*: positivity is a property of the complete comb —  $v877$ 's truncation reversal sharpened to atom resolution; no monotone one-prime induction exists in any order; induction must run in a flow coordinate, not in comb content. No marker move; no RH statement.

**The degree-2 Lax coefficient flow (PRIME.PORT.LAX2.01).** *Registered* 2026-08-09, measured ([verification/v885\_ladder\_flow\_anatomy.py], 11/11, LAX2-MEASURED (FLOW-NOISY / LAX2-PARTIAL)). *Measured*: the twelve deformation-polynomial coefficients are extracted on all 41 ladder steps, but normalized by the  $\alpha$  step the coefficient series behave as noise (median roughness 1.365 vs lawful bar 0.7); the operator-level IKS prediction misses the bar narrowly (0.350 vs 0.3); the *integrated* tau-transfer identity  $\log \det(I + (I - D_k)^{-1} \Delta D)$  holds exactly (warded at  $5.6 \times 10^{-11}$ ). *The named lead*: the dictionary columns  $f, Yf, Y^2f$  are nearly collinear — the raw coefficient basis is ill-conditioned; a well-conditioned reparametrization may expose a lawful flow the raw basis masks. *The demand*: the flow in a conditioned coordinate; then couple to the PRIME.PORT.TAU.01 variational identity. *Adjudication note* (2026-08-24, round 225): the demand is adjudicated — **the word “2” was a hypothesis**: with the canonical basis-invariant conditioning the true minimal degree in the closing direction is *one*, and the closing direction is  $h$ , not  $s$  (LAX1-H-EXACT: rank-1  $h$ -step, blind-predictive degree-1 connection, zero curvature world-blind; the  $s$ -time has no fixed degree, CONDITIONING-ONLY, and across rungs there is no common carrier — [verification/v955\_tau\_iiks\_toda\_dictionary.py], PRIME.PORT.TAU.FINITE.IKS.01 [E]). Per the contract's no-go rule, no further Lax cosmetics in the  $s$ /across directions. *Kill*: a conditioned basis in which the flow remains noise at the frozen bars (executed: the  $h$ -basis closes exactly). No marker move; no RH statement.

**The half-filling / free-moment theorem set (PRIME.FREEMOMENT.JFRACTION.01 [E] + PRIME.JACOBI.DUAL.REVERSAL.01 [E]) and the positive-prefix contract (PRIME.FREEMOMENT.POSITIVEPREFIX.01 [O]).** *Registered and delivered / registered* 2026-08-24 (rounds 228–231, notes DLIV–DLVII; [verification/v956\_signedmoment\_halffilling\_duality.py], probes embedded byte-exact, 14/14 + 17/17 + 17/17 + 13/13 first pass). *Exact (JFRACTION)*: the half-filling law  $N_w = \lceil \# \text{supp}(\tilde{\mu})/2 \rceil$  on all five windows with the wall dying immediately past the cap (first  $r$ -flip at  $N_w + 0/2/2/3/1$  — the wall is *maximal*); the complement identity and  $h$ -duality at dps 80; the free moment window law (the largest Hankel window inside the  $S$  free moments is  $n = (S + 1)/2 = N_w$ ; MAIN exhausts the entire free window, the controls die at 11–15%); the inertia theorem; the four-path dictionary (moment prefix = Jacobi chain = Euclidean chain = value chain, exact in rationals; Epstein's first wrong Euclidean pivot exactly at its flip  $n = 25$ ). *Exact (DUAL.REVERSAL)*: the full chain reversal ( $\alpha_m^\# = \alpha_{S-1-m}$ ,  $\beta_m^\# = \beta_{S-m}$ ; NO-SOURCE-REVERSAL — not a builder transformation) and the  $L$ -gauge midpoint connection theorem ( $Y_{N-1+l}^\# = JY_{N-1-l+1}W$ ,  $W = \text{antidiag}(1/L, L)$ , built from the node polynomial alone; exact in rationals, real w9 at  $9.1 \times 10^{-94}$ , dps 100); the gauge is  $h$ -free — the wall's orientation sits exclusively in the  $h$ -chain it does not see (SIGNED-STOKES-WALL-EQUIVALENT). *The open contract (POSITIVEPREFIX)*:  $\beta_n > 0$  through the entire free moment window as a theorem with an identified mechanism — the wall itself in free-moment coordinates; half-filling is the *location*, its positive reachability is the result to be proven; named sub-deliverables: the source-pure scalar normalization beyond the node-log pole remover, and a source-side critical filling  $t_c$  as output. *Kill (of any mechanism candidate)*: the frozen micro falsifier — a parametrix must predict the forced tail 0/2/2/3/1 and the control flips 25/21/27 blind; a model that only paints the positive side is a fit, not a mechanism. Design rules unchanged: the full von Mangoldt comb stays in the main problem, no PNT smoothing. No marker move beyond the two [E] registrations; no RH statement.

**The bordered-RHP / tau-readout theorem set (PRIME.PORT.RHP.BORDERED.READOUT.01 [E]) and the full-source campaign contract (PRIME.PORT.RHP.FULLSOURCE.BASEFIBER.01 [O]).** *Registered and delivered / registered* 2026-08-24 (rounds 244–255, notes DLXXIII–DLXXXV; [verification/v958\_bordered\_tau\_readout\_dictionary.py], probes embedded byte-exact, 23/23+29/29+24/24 + 20/20 + 22/22 + 26/26 first pass, plus the module-own exact-rational PSD base theorem).

*Exact (BORDERED.READOUT):* the bordered-RHP dictionary (the Uvarov border column carries  $F_n$  as its  $z^{-1}$  coefficient; the budget is a CD-kernel functional of the terminal  $Y_n$  data alone,  $S_{n-1} = \iint K_n d\sigma d\sigma$ ; the Uvarov tau step and the  $T$ -sourced bordered LAX1 flow, mp-exact through all 184 w9 degrees); the Schlesinger rank-1 insertion ( $\det Y_3 = -h_{n-2}F_{n-1}$ ; IKS generator rank exactly 2; canonical  $2 \times 2$  RHP  $\times$  one-dimensional budget cocycle carrying the full pivot chain — the irreducible- $3 \times 3$  campaign dropped without replacement); the centering congruence ( $B^\circ = B - \rho_0$ , exact in rationals), the  $\sigma^\circ$  dictionary, the dual-norm identity  $Q_m = t^T H^{-1} t$  and the tail-kernel compression  $T = \iint [K_N - K_8] d\sigma^\circ d\sigma^\circ$  (dps 160 through  $N = 184$ , rel  $3.3 \times 10^{-160}$ ); the three exact error formulas ( $\delta h_n = (R_1)_{12}$ ; the central rest identity; the error-kernel sandwich at  $1.1 \times 10^{-81}$ ); the contour formula  $R_1 = \frac{1}{2\pi i} \oint (R(z) - I) dz$  ( $3.2 \times 10^{-14}$ ); the exact augmented tau quotient  $D = B^\circ - Q_7 - T = \tau^{\text{aug}}/\tau$  at full depth (three routes,  $10^{-11}$  / mp  $10^{-16}$ ) with the pole classification (rank-1 border poles removable by residue closure; the  $D$ -layer essentially singular,  $\exp(c/(z-s))$  type); and the PSD base theorem (positive basis  $\Rightarrow$  Gram PSD  $\Rightarrow$  border readout  $\geq 0$ ; bordered PSD = wall + budget by exact Schur complement; the frozen signed witness fires — positive-basis structure alone proves the wrong statement). *Typed honestly (measurements, not identities):* the corner census (CORNER-IMPORTED-ONLY, PSD 0/42 for all three sealed candidates) and the head-heavy budget profile (IRREGULAR-BULK, Gini 0.91); amplification extensive ( $\mathfrak{A} \sim N^{1.01}$ ), annihilation neutral, the fiber fully in the error; the gauge thesis refuted (NO-GLOBAL-GAUGE, MODEL-ORIENTATION-BLIND — BASE-BLOCKED-BY-DRIFTING-LAYER on record); the diagonal pairing thesis refuted (diagonal 19%, wrong sign on scramble; REL-STRUCTURELESS, PAIRING-EXTENSIVE). *The open contract (FULLSOURCE.BASEFIBER, the campaign; the coupled-tau dictionary and telescope certified in [verification/v959\_coupledttau\_terminal\_dictionary.py]; carried at the wave-4 state since 2026-08-25 — surface and orientation certified in [verification/v960\_terminal\_surface\_closure.py] and [verification/v961\_midpoint\_orientation\_dictionary.py], block below):* per the round-256 causal decision the separate treatment of base and fiber is *artificial* — the campaign object is the indivisible full-source pair  $(\tau, \tau^{\text{aug}})$  with  $D = \tau^{\text{aug}}/\tau > 0$  plus the no-pole statement, on the full von Mangoldt comb; admissible target form stays an *integrated* statement over the compressed CD readout (no uniform  $Y_n$  parametrix; the v959 coupled recursion is the one-dimensional budget cocycle of the v958 formulation in closed form). *The precise localization after rounds 256–259: the last fiber edge is a positive tau cross-ratio* — the entire fiber positivity is the one terminal inequality  $h_{N-1}/F_{N-1}^2 > 7/5$  (equivalently  $q_N < 1$ ; the 5/7 floor typed FLOOR-IMPORTED, derivation open; the measured driver is the rank-perfect  $F$ - $h$  covariance +1.000: the ratio is the object) — and *the last base edge is a globally oriented prefix resummation* — the determinant sign is carried by coherent cancellation over the configuration ensemble (a resummation gap, not a selection gap). The two edges are registered as the successor contracts PRIME.PORT.COUPLEDTAU.TERMINAL\_CROSSRATIO.01 [O] and PRIME.PORT.FULLSOURCE.PREFIX\_RESUMMATION.01 [O] (blocks below; wave-4 state: the 42-rung surface of the fiber edge is *certified* in v960, the orientation edge has its blind Maslov GO and its named successor in v961). *The formalized target form (wave 4):* both edges are corollaries of the one master theorem `augmented_prefix_positive (rh/lean/RH/Window.lean)` —  $A_{w,n} = \begin{bmatrix} H_n & u_n \\ u_n^T & B \end{bmatrix} \succ 0$  for every admissible MAIN window and prefix degree, with `MainWindow` an opaque predicate (the three reviewer counterexamples are permanent Lean guards); modulo this one **sorry** the two ex-edges are Lean-proved, and per the wave-4 measurements the predicate must encode *metric* exactness (sub-gap-exact comb placement), not combinatorial Euler properties. *The north star reformulated (wave 5, [verification/v962\_halffilling\_pinning\_theory.py]):* by the moment counting theorem T1 (the free pivots are exactly  $h_0..h_{N_w-1}$ ), the crossing budget theorem T2 (the crossing count is  $S_-$ , world-blind) and the main window reduction T4, the campaign's open center is, in reinstform, **why is the signed prime moment form quasi-definite up to the maximally free order?** —  $\forall n < N_w : h_n > 0$ , exactly equivalent to  $\min C \geq N_w$ , and via T4 the base half of the master **sorry** (Lean: `free_window_positivity`, `rh/lean/RH/Window.lean`; carried on the successor contract PRIME.PORT.RHP.FULLSOURCE.QUASIDEFINITENESS.01 [O], block below). *The canonical form (wave 6, [verification/v963\_lstar\_reduction\_dictionary.py]):* by the r283 A2 chain the north star is exactly **the missing lemma L\*** — for every real polynomial  $p$  with  $\deg p < N_w$ :

$\int p^2 d\nu < \int p^2 d\mu$ , equivalently  $\lambda_{\max}(E_{N_w}) < 1$  (the  $\nu$ -dressed  $\mu$ -CD kernel is a strict contraction at half-filling depth) — a two-measure moment problem instead of a determinant cascade; registered as `PRIME.LSTAR.SUBORDINATION.01` [O] (block below) with the standalone problem document `rh/problem/lstar_problem.tex` and the Lean statement `lstar_subordination` (the direction  $L^* \Rightarrow$  free-window positivity is proved). *The measured state after rounds 250–259 (rounds 250/254/255 experiments-side, SPEC f9bf299d / 1083f7a0 / 45875904):* OUTER-MODEL-FAILS (rates carried — plateau  $((b-a)/4)^2 = 0.2500$  exact within 0.040,  $h$ -rate slopes  $\leq 0.007$  dec/degree — pointwise and fiber failed; the naked electrostatic model beats the dressed one 2.2–3.8 $\times$ ); OFFDIAG-EXTENSIVE ( $D_{\text{part}} \approx 0.195N$ ; the atom diagonal carries 78.6% of the world gap; the round-254 compactness-inversion headline and the Epstein fiber anomaly are *downgraded on record* by the round-256 firewall — both are base contamination of the indefinite continuation and may not be cited as border/fiber discriminators); COCYCLE-LAW (round 257: the drain is monotone on MAIN through the entire free window and its increment law breaks exactly at the control flips 21/25/27 — base and fiber are coupled in one scalar law; both augmented border variants cross the wall: no bordered spectral no-pole shortcut); and the round-255 one-bit reading is *revised* (round 259: the same short-range swap structure sits on the mains with fully overlapping gap distributions — the base flip stays a threshold event of an extensively accumulated quantity, but the missing information is not one-dimensional). *Kill (of any mechanism candidate; unchanged plus the v959 no-gos):* a base layer must produce the control flips on Epstein/smooth (orientation, not just rate); a fiber main term must separate scramble by  $\geq 1$  decade before readout and survive the  $\times 2$  share window; mass majorants, level crossing, one-swap statistics and  $k$ -swap truncations are stop-listed (measured dead in rounds 258/259). No marker move beyond the [E] registrations; no RH statement.

**The coupled-tau / terminal-budget theorem set** (`PRIME.PORT.RHP.COUPLEDTAU.TERMINAL.01` [E]) and the two terminal-edge contracts (`PRIME.PORT.COUPLEDTAU.TERMINAL_CROSSRATIO.01` [O], `PRIME.PORT.FULLSOURCE.PREFIX_RESUMMATION.01` [O]). *Registered and delivered / registered* 2026-08-25 (rounds 256–259, notes DLXXXVII-DXC; [verification/v959\_coupledtai\_terminal\_dictionary.py], probes embedded byte-exact, 21/21 + 26/26 + 27/27 + 19/19 first pass, plus the module-own exact section S0). *Exact (COUPLED-TAU.TERMINAL):* the positive-prefix firewall (MAIN  $p_{\max} = 184 = N$  with zero negative modes, the only positive-prefix world; controls INDEFINITE-CONTINUATION at 21/25/27; every prefix kernel numerically PSD at  $5.7 \times 10^{-16}$ , re-proven as exact algebra in-module:  $H_p$  PD iff the whole  $h$ -prefix is positive,  $\det H_p = \prod h_i$ , Schur-complement must-fail firing) with the exact  $2 \times 2$  causal census and its two contamination downgrades on record (the r254 scramble compactness is base contamination,  $f_{\text{neg}} = 0.9311$ ; the Epstein fiber anomaly lives entirely in the indefinite continuation  $k \geq 25$ ); the coupled two-term recursion  $\tau_{n+1} = a_n \tau_n$ ,  $\tau_{n+1}^{\text{aug}} = a_n \tau_n^{\text{aug}} + b_n \tau_n$  with  $a_n = c_n^2 h_n$ ,  $b_n = -(c_n F_n)^2$ , equivalently  $D_{n+1} = D_n - F_n^2/h_n$  (anchor rel  $4.8 \times 10^{-11}$ , mp dps 220 through the scramble flip), with the bilinear tau form  $\tau_n^{\text{aug}} \tau_{n+1} - \tau_{n+1}^{\text{aug}} \tau_n = (c_n F_n)^2 \tau_n^2$  derived and gated in-module — the base alone carries the sign, the border only nonnegative magnitude, and the cocycle breaks exactly at the control flips; the blind coefficient-field pass of the frozen micro-falsifier (control flips 3/3 and forced tail 0/2/2/3/1 on all five windows; the parametrix half open); and the Schur budget telescope  $D_N = B - S_{N-1}$  (mp dps 220 at  $2.4 \times 10^{-16}$ ;  $B = S_{N-2} + 5/7$  by construction) with the consequence gate margin  $> 0 \iff q_N < 1$  and the corner-ray  $t^*$  normal form (slope exactly 5/7, root exactly  $q_N$ ). *Typed NUMERICAL (measurements, not laws):* the  $q$  census (42/42, flat, min margin 0.0139); the  $F$ - $h$  covariance +1.000 (a measured structural fact); the 5/7 surface (FLOOR-IMPORTED, derivation open); the r259 ensemble-coherence findings and the parametrix orientation blindness. *Honest negatives (load-bearing):* the level-crossing no-go (branches cross 5–16 degrees too early; MAIN crosses 45 times with zero true flips), the one-swap no-go (r255's ORIENTATION-LOWDIM(1) revised on record), the mass-majorant no-go (all three sealed majorants 0/42, median  $\times 490$ , break at  $0.97N$ ). *The open contract (TERMINAL\_CROSSRATIO, the fiber edge):* prove  $h_{N-1}/F_{N-1}^2 > 7/5$  on the cofinal ladder — equivalently derive the true terminal floor source-pure; the bound must bind the *ratio* directly (no separate  $F$ -majorant against an  $h$ -minorant), hold against the flat measured margin profile (min 0.0139, no  $N$ -decay to spend) and survive the near-terminal zone at  $0.97N$ ;

*kill list (stop-listed)*: the mass-majorant class, compression, constants, band locality, any one-sided bound on the covariant pair. *The open contract (PREFIX\_RESUMMATION, the base edge)*: the closed analytic form of the coherent ensemble resummation — the determinant as a globally oriented signed partition sum, completing the coefficient-field pass to a parametrix-level pass (predict the control flips and the forced tail blind from a closed model, no minors/slogdets, AST-scanned); *kill list (stop-listed)*: two-branch level crossing, one-swap/ $k$ -swap statistics, saddle-point dominance, the r255 parity-rule and hysteresis hybrids (0/3+0/5 blind). *Wave-4 progress (2026-08-25, block below)*: on TERMINAL\_CROSSRATIO the 42-rung surface is now *certified* (41 mechanism target-blind + kz15 exact-finite; the cofinal front is the single hypothesis H5 / lemma L2). *Wave-7 progress (2026-08-25, [verification/v964\_lstar\_coherence\_census.py])*: the H5/L2 front has its *first named classical theorem candidate* — the exact van der Corput inequality at  $H = \lceil \sqrt{m} \rceil$  (constant-free, arithmetic-free) delivers  $\delta' = +0.309 > 0.21$  world-blind and certifies 6/7 exceptions +38/42 rungs; only the kz15 razor blocks 7/7 (exact-finite per r270); registered as the successor contract PRIME.PORT.L2.VDC\_LEMMA.01 [O] (block below) — the remaining input is the chain origin of the  $P$ -variance scaling. On PREFIX\_RESUMMATION round 261 closed the finite-combinatorial half honestly negative and the named path is the oriented midpoint theorem (Maslov census GO; successor contract PRIME.PORT.RHP.MIDPOINT.ORIENTED\_THEOREM.01 [O], executed as rounds 279–281 and consumed in wave 5, block below). Mincut base 4 / refined 5 unchanged; no RH statement.

**The terminal-surface closure and the midpoint-orientation theorem** sets (PRIME.PORT.COUPLEDTAU.SURFACE\_CLOSURE.01 [E], PRIME.PORT.RHP.MIDPOINT.ORIENTATION\_DICTIONARY.01 [E]) and the two wave-4 contracts (PRIME.PORT.RHP.MIDPOINT.ORIENTED\_THEOREM.01 [O], PRIME.PORT.WALL.METRIC\_FIREWALL.01 [O]). *Registered and delivered / registered* 2026-08-25 (rounds 260–278, notes DXCII–DCXI; [verification/v960\_terminal\_surface\_closure.py] with 31/31 + 26/26 + 26/26 + 28/28 + 29/29 + 29/29 + 28/28 + 28/28 + 24/24 + 24/24 first-pass gates and [verification/v961\_midpoint\_orientation\_dictionary.py] with 32/32 + 25/25 + 29/29 + 31/31; probes embedded byte-exact and executed in their sealed -smoke stage, SPEC SHAs pinned, full records re-verified by rh/verification/run\_rh.py — since this wave a gate of bash build.sh audit). **Exact (SURFACE\_CLOSURE): the 42-rung surface census  $q_N < 1$  is certified on all 42 rungs as a finite fact** — the two-branch theorem ( $U$  = the sealed C1b triangle form,  $M = \sqrt{5/7}$ ,  $Z$ -identity at  $1.8 \times 10^{-14}$ ; cheap branch  $g_w \geq 0$  exactly 35/42 including both mains *without cancellation*; exception set exactly  $\{kz15, kz20, kz22, kz36, kz38, kz39, kz52\}$ ), the sealed phase bounds (edge truncation: kz22, the first source-pure exception certification ever; block alternation: +0.26 to +0.41 decades on all 42, six exceptions certified, no detector fires), the kz15 interval certificate (outward-rounded intervals, dps 640, end width  $1.5 \times 10^{-92}$ , margin +0.02680, dps-halving ward — typed an exact finite enclosure, no mechanism), and the universal pair theorem (H1–H4 structural, Lean-proved in rh/lean/RH/PairBound.lean; H5 = the margin  $|Z_{\text{local}}| + \varepsilon < \sqrt{5/7}$  the single window-dependent hypothesis). **Exact (ORIENTATION\_DICTIONARY)**: the base Casoratian is the pivot chain itself ( $\hat{\pi}_{n+1}q_n - \hat{\pi}_nq_{n+1} = h_n$ ,  $c' = 1$ ; right solution triple-constructed, the node polynomial is the right Dirichlet boundary), the  $h$ -free midpoint form *is* the node polynomial (the r231 sign-blindness as an identity), the augmented telescope  $D_{n+1} = B - W_n^{\text{aug}}/W_n^{\text{base}}$  exact against the independent r263 route, and the Maslov census GO (rule R2 = Jacobi interlacing/reality: blind 42/42, mains safe at full depth, controls fire exactly at flip+1, one-way, not  $h$ -equivalent). *Typed NUMERICAL (measurements, not laws)*: the L2 scaling anatomy (BOUND-COARSENESS: the truth margin rises with  $N$ ; flip condition  $\delta' > 0.21$  of available 0.45, address c3), the Euler surgery ladder (PERTURBATION-INSENSITIVE + the firewall map: the wall, not the cancellation rate, is what is arithmetically special), the metric-firewall continuum ( $D \sim \theta^b$ , support exactness first), the exact position gradients (Hellmann–Feynman class, identity fine type) with the predictive small-primes-loaded u-profile, and the perturbative-only stability law. *Honest negatives (load-bearing)*: the eventual-triangle no-go (kz52 itself an exception), three identity routes dead with break loci, the retracted w13 Gram provenance, the KYP memory no-go from structure (exact double obstruction o1/o2 on

46/46 worlds; the admissible Riccati memory target-inverse), and the raw atom-Sturm census refuted as winding quantity (STURM-CHAIN-VERIFIED honestly not awarded). *The open contract (ORIENTED\_THEOREM, the orientation edge)*: the contraposition form — an independent index obstruction that forbids the R2 interlacing break before half-filling on MAIN; the detection direction (break at flip+1, after the flip) forces exactly this form, and the MainWindow arithmetic must enter here as *metric exactness*; *wave-5 state*: rounds 279–281 are executed and consumed ([verification/v962\_halffilling\_pinning\_theory.py], block below) — the theorem share is promoted, the candidate generic two-sided obstruction is refuted exactly, and the center is precisified to the budget localization (free-window survival). *The open contract (METRIC\_FIREWALL, the firewall lane)*: the global form of the wall’s metric stability — a wall lemma quantitatively stable under family/assignment permutations at small dose while consuming the metric data exactly; the certified state is graded-continuous with exact gradients but perturbative-only ( $\theta^*$  25–170× below the smallest measured dose,  $L_D$  window-specific, no uniform lemma candidate); the MAIN positivity itself (H5 / the master theorem) is not part of this demand and stays the open center. *The surface fence (binding)*: the surface closure is a census-plus-certificate result, *not* a cofinal theorem — the window quantifier stays. Mincut base 4 / refined 5 unchanged; no RH statement.

**The half-filling pinning theory (PRIME.WALL.HALFFILLING\_PINNING\_THEORY.01 [E]) and the two wave-5 contracts (PRIME.PORT.REPRESENTATION.CONTEST.01 [O], PRIME.PORT.RHP.FULLSOURCE.QUASIDEFINITENESS.01 [O]).** *Registered and delivered / registered* 2026-08-25 (rounds 279–281, notes DCXII/DCXIV/DCXV; [verification/v962\_halffilling\_pinning\_theory.py] with  $32/32 + 29/29 + 28/28$  first-pass gates; the three probes embedded byte-exact in their sealed-smoke stage, SPEC SHAs pinned, full records re-verified by rh/verification/run\_rh.py; plus the module-own exact section S0 in pure Fractions). *Exact (the four theorems as named gates)*: (T1) the **moment counting theorem** — the pivot  $h_n$  consumes  $m_0..m_{2n}$ , the  $S$  free moments are  $m_0..m_{S-1}$ , hence the free pivots are exactly  $h_0..h_{N_w-1}$  and  $h_{N_w}$  is the first forced one: half-filling is the end of the free moment space, arithmetic-exact for all  $S = 2..2000$  with the exact-rational freedom demonstration — “why half-filling” is answered by counting; (T2) the **crossing budget theorem** —  $\#\{n < S : h_n < 0\} = S_-$  exactly (Jacobi/Sylvester), world-blind; (T3) the **two-sided parity theorem** — node-sign pattern, gap parity and the census bilanz,  $h$ -blind at every degree, with the 87376-case exhaustion; (T4) the **main window reduction** — given T1 and T2 the entire open statement of the wall is  $\min C \geq N_w$ , exactly equivalent to  $\forall n < N_w : h_n > 0$ : one placement question, the north star in reinstform. *The refutations (named negative gates, load-bearing)*: NO-UNIVERSAL-O1-PINNING (exact; the one-negative toy has offset  $N_w-2$ , unbounded in  $S$ ; only provable ceiling: the pigeonhole  $\min C \leq S - S_-$ ;  $C \leq 5$  is a 42-rung measurement), NO-EXTREMALITY (measured + exact witness: three mp-confirmed directions push the w9 crossing  $184 \rightarrow 185$ ; a 4-atom one-negative measure has  $\min C = N_w + 1$  exactly), NO-GENERIC-MASLOV-OBSTRUCTION (exact; three rational counterexamples — the two-sided machinery carries no arithmetic), NO-SIMPLE-OFFSET-LAW (measured; six sealed source-pure candidates,  $\max |\text{sp}| = 0.273 \ll 0.75$ , all world-blind). *Lean state*: T1 and T4 *proved* (moment\_counting\_free\_pivots, main\_window\_reduction); T2 *stated* (crossing\_budget, sorry with v962 reference); the fog-free central hole is free\_window\_positivity (rh/lean/RH/Window.lean, via T4 the base half of the master sorry); the exact one-negative instance is the permanent guard upper\_pinning\_not\_universal. *The open contract (REPRESENTATION.CONTEST)*: the sealed contest of representations for the localization statement — which coordinate system (dual chain, moment/Weyl, gradient/criticality, Jacobi interlacing, tau/cocycle) carries a MAIN-separating, non-restatement handle on free-window survival; *forbidden (binding)*: no derived 5/7, no bound mechanism claim, no asymptotic law, no localization claim from any census, no undeclared  $h$ -reading candidates, no world-blind carrier; the v962 refutation set binds the candidate space. *The open contract (FULL-SOURCE.QUASIDEFINITENESS, the open center)*:  $\forall n < N_w : h_n > 0$  on every admissible MAIN window — why is the signed prime moment form quasi-definite up to the maximally free order? *Forbidden (binding, the measured kill lists)*: no bookkeeping-only proof (the Lean counterexample guards are permanent), no world-blind mechanism, no Weyl  $\mu/\nu$  perturbation

route (range 1.1%), no dual restatement, no variational local-maximum route, no simple off-set law, no universal  $O(1)$  upper pinning. Both rounds (`representation_contest_probe.py`, `fullsource_quasideterminateness_probe.py`) were in flight at the wave-5 cut and are *executed and consumed in wave 6* (`[verification/v963_lstar_reduction_dictionary.py]`, block below): REPRESENTATION.CONTEST adjudicated CONTEST-ALL-DEAD (the four classical languages closed with exact reasons), FULLSOURCE.QUASIDEFINITENESS adjudicated MECHANISM-PARTIAL(A2) —  $L^*$  established as the canonical reduction, the open center carried forward on PRIME.LSTAR.SUBORDINATION.01 [O]. Mincut base 4 / refined 5 unchanged; no RH statement.

**The  $L^*$  reduction dictionary (PRIME.LSTAR.REDUCTION\_DICTIONARY.01 [E]) and the  $L^*$  subordination contract (PRIME.LSTAR.SUBORDINATION.01 [O]).** *Registered and delivered / registered* 2026-08-25 (rounds 282–285, notes DCXVI/DCXVII/DCXIX–DCXXI, plus the standalone problem document of note DCXX; `[verification/v963_lstar_reduction_dictionary.py]` with 30/30 + 26/26 + 30/30 + 33/33 first-pass gates; the four probes embedded byte-exact in their sealed -smoke stage, SPEC SHAs pinned, full records re-verified by `rh/verification/run_rh.py`; plus the module-own exact section S0 in pure Fractions). *Exact (the reduction as named gates):* the  $\mu$ -frame congruence  $\text{minor}_k(D_\mu - G) = D_k(\tilde{\mu})$  at every  $k \leq S_+$  (truncation must-fail loud); the frame contraction  $h_0..h_{m-1} > 0 \iff \lambda_{\max}(E_m) < 1$  (both truth values); the crossing dictionary  $\lambda_{\max}$  crosses 1 exactly at  $\min C + 1$  (CAPACITY-CEILING typed where  $\min C = S_+$ ); the pigeonhole ceiling  $\min C \leq S_+$  re-proved in the frame; capacity-as-counting refuted by the rank-one pair (metric values  $1.25 \times 10^{-13}$  vs  $1.25 \times 10^2$  decide); the r285 bookkeeping  $\lambda_{\max} = \text{maxdiag} \times (1 + \text{assist})$  with the sign-exact budget equivalence; and the r282 four-language elimination as named negative gates (CONTEST-ALL-DEAD — SOS iff empty negative register, Kasteleyn orientation iff  $S_- = 0$  by full  $2^S$  exhaustion, Hamiltonian-PSD =  $h > 0$  bookkeeping, dual-pair sync by theorem; the common deep reason: every classical positivity language forces positivity exactly for the positive measure class, and MAIN is signed). *Typed NUMERICAL (measurements, not laws):* the one-atom-vs-collective dichotomy ( $n_{\text{DIAG}} = 187$  vs crossing 185; controls collective), the sub-classical edge growth ( $p = 0.38$ , classical coverage 0/42 — the (D) route must be a discrete bound), the ensemble position (MAIN’s wall assist the extreme LOW-OUTLIER, pct 0.00), and the first two positively MAIN-separating detectors of the program (assist-at-crossing 0.0195 vs dead 1.69..2.99; random-sign  $z = -3.15$  destructive vs dead  $+4.95.. + 12.44$ ). *The open contract (SUBORDINATION, the canonical open center):* for every admissible anchor  $z$  and every real polynomial  $p \neq 0$  with  $\deg p < N_w = (S + 1)/2$ :  $\int p^2 d\nu_z < \int p^2 d\mu_z$  — equivalently  $\lambda_{\max}(E_{N_w}) < 1$ ; *state (wave 7, [verification/v964\_lstar\_coherence\_census.py]):* 57/57 measured (the 42 plus the 15 extension anchors,  $N_w$  up to 1218), **the margin decay resolved harmless** (r286: no counterexample, the  $O(1)$  census offset survives beyond the cap, flattening power law  $\alpha \approx 3.05$  with driver  $c_w \rightarrow 1$ , and the margin falls *slower* than the local loading speed — no acute falsification threat; the unbounded family stays open); the coherence half has a *named carrier class* (r288: antiphase next-nearest ARCH–ARCH pairs,  $z_v = -3.149$ , total control reversal, finest alignments below phase resolution); **the diophantine route is excluded** (r289 METRIC-ONLY: the rational twin keeps the full signature identical while every exact log-relation is destroyed; metric coherence threshold  $10^{-3}..3 \times 10^{-3}$  of the local gap; Baker  $\sim 7900x$  too weak *and unnecessary*); *state (wave 8, [verification/v965\_lstar\_curvature\_arc.py]):* the profile-functional question is **functionally negative-sealed** — after five sealed class families (local scalars, non-local/low-rank, the curvature two-form in the mixed metric, the metric-reconciled two-form, the hardened partial-free  $|\text{sp}|$  form) no closed predictive profile functional of the working set exists over the honest bar (r294 F10-FRAGILE + r295 F10-SP-MAJORITY: F10 unpromoted), while the working-set geometry is now *measured* (anisotropic tube, one-degree budget ridge, rank-1 DENS curvature valley, simple  $h_{184}$  flip zeros); the successor question is PRIME.LSTAR.CLOSED\_FUNCTIONAL.01 [O] (block below); no round is in flight at the wave-8 cut; the standalone document `rh/problem/` is the external reference; *kill lists (inherited, binding):* the v962 refutation set, the r282 four-language closure, the r283 forbidden list (no triangle/Gershgorin bound — dead at  $n = 21 \ll 185$  — no target-inverse quantity, no truncation, no posthoc window), the r284 shielding/Nyquist refutations, the r285 no-fixed- $\delta/\delta'$

trend, plus the wave-7 kills (no diophantine/linear-forms route, no plain phase-dispersion separator, no sinc/CD-only sign mechanism), no derived 5/7, no bound mechanism claim, no asymptotic law. Lean: `lstar_subordination` stated (`rh/lean/RH/Window.lean`, documented `sorry`); the direction  $L^* \Rightarrow$  free-window positivity is *proved* (`lstar_implies_free_window`). *State (wave 11, [verification/v968\_architecture\_adjudication.py])*: the formal frame is **restructured** — the master theorem is now *proved from*  $L^* +$  terminal (`lstar_terminal_implies_master`, r305, with `augmented_prefix_positive` and `free_window_positivity` as corollaries; sorries  $9 \rightarrow 5$ , the two TRUE holes standing alone: `lstar_subordination` and `terminal_positive_main`), the window source is a real Lean construction (r310/r310b: nodes =  $\log p^k$ , weights =  $\Lambda(p^k)$  derived; source exactness, mass conservation of any folding, the strong stabilization form), and the surrounding architecture is fully adjudicated (Lane A closed as cone language, B not overtaking with the rank-one identities banked, C parked); additional kill entries: no fixed-head route (r307), no constructive PSD-G rule in the planned cone language (r312), no transfer-operator contraction at chain level (r313), no two-regime pointwise majorant under a mid-ladder freeze (r316). *State (wave 12, [verification/v969\_forks\_and\_redteam.py])*: the indefinite fork is adjudicated (r318 P2-MAIN-SPECIFIC: P1 banked *as language* — the index statement is a total restatement of  $L^*$ , no lever; classical sign regularity dead on MAIN) and the antiphase fingerprint is closed as ALGORITHM-ARTIFACT (r322: a least-norm-proximal Dykstra-basin property, not a psd-solution-set property; the selection geometry is the named residual); the Lean statement class is red-team hardened (the r319 U-findings repaired by the r320 retypes + guards, census 5); **per the binding reviewer adjudication the lane is internally paused** (allocation 0%; the selection geometry is not pursued further; reopening only under the six documented conditions of the r322 round notes; the external memo candidate: the two-weights sampling formulation  $\|J_N\| < 1$ ), and the hole itself is typed **window-local** (Level B of the wave-12 three-level proof graph — see the red-team block below). Additional kill entries (wave 12): no Pontryagin/Krein index lever, no sign-regularity/variation-diminishing route, no antiphase sign-law theorem without a canonical selection rule. *State (wave 13, [verification/v970\_extraction\_and\_composition.py])*: the pause held — the planned r323 fork was cleanly aborted at the end of its design phase, before any write access (nothing counts as a round); the Lean census is now  $5 \rightarrow 8$  (three typed Level-C extraction statements, r326 — none touches this hole; `lstar_subordination` byte-identical). Mincut base 4 / refined 5 unchanged; no RH statement.

**The  $L^*$  coherence census and the vdC L2 lemma contract (PRIME.LSTAR.COHERENCE\_CENSUS.01 [E], PRIME.PORT.L2.VDC\_LEMMA.01 [O]).** *Registered and delivered / registered* 2026-08-25 (rounds 286–289, notes DCXXIV/DCXXII/DCXXV/DCXXVI; [verification/v964\_lstar\_coherence\_census.py] with  $29/29 + 25/25 + 31/31 + 42/42$  record gates; the four probes embedded byte-exact in their sealed `-smoke` stage, SPEC SHAs `0a44ac4e/761d88fa/da46f7ee/91cdc2b1` pinned; plus the module-own exact section — the Fejér window-sum identity, the van der Corput inequality with its Cauchy–Schwarz proof route, the Abel/Erdős–Turán bound, the mutant must-fails, the 2-power resonance count  $\binom{8}{2} = 28$ , in pure Fractions). *Exact (the vdC theorem)*:  $|\sum_j P_j|^2 \leq \frac{m+H-1}{H} \sum_{|h|<H} (1 - \frac{|h|}{H}) A(h)$  — constant-free, arithmetic-free, proved module-own; at the frozen window  $H = \lceil \sqrt{m} \rceil$  it delivers  $\delta' = +0.309 > 0.21$  world-blind (admissible: L2 is the generic half), certifies 6/7 exceptions +38/42 rungs, and only the kz15 razor blocks 7/7 (exact-finite per r270, reserve 0.0268); F1 discrepancy/Erdős–Turán dies loud ( $\delta' = -0.21, 0/7$ ); the autocorrelation map relocates the c3 address: on 39/44 worlds the block sequence *reinforces* — the cancellation is already absorbed into the  $P_j$  magnitudes, the frames win through the  $\sqrt{m}$  count economy. *Typed NUMERICAL (measurements, not laws)*: the r286 extension census (15 new anchors  $N_w = 942..1218$ , all margins mp-sign-safe positive, min  $+1.806 \times 10^{-8}$ ; off census  $O(1)$ , max +4; flattening power law  $\alpha \approx 3.05$ ; HARMLESS quantified;  $q_N$  reconciled; EXTRAP-CALIBRATED 15/15), the r288 carrier map (antiphase next-nearest ARCH–ARCH pairs,  $z_v = -3.149$ , total control reversal, finest alignments below phase resolution), the r289 twin table and jitter ladder (metric coherence threshold  $10^{-3}..3 \times 10^{-3}$  of the local gap). *Honest negatives (load-bearing)*: SAMPLING-BLIND (the sinc/CD predictor partial,  $0.728/0.878 < 0.90$ ), SOURCE-SEPARATOR-NOT-

FOUND (plain phase dispersion world-blind, ensemble-generic), DIOPHANTINE-EXCLUDED(METRIC-ONLY) — the rational twin (every tent center replaced by a rational  $p/q \times \Delta$  at position cost  $\leq 2.1 \times 10^{-9}$ , every exact log-linear-form relation destroyed) keeps the full signature *identical*: the effect of the diophantine structure is null, Baker territory cleanly excluded ( $\sim 7900\times$  too weak in the exponent *and unnecessary*). *The open contract (VDC\_LEMMA, the L2 chain statement)*: prove the vdC form  $|R| \leq \sqrt{\frac{m+H-1}{H}} \sum_{|h|<H} (1 - \frac{|h|}{H}) A(h)$  at  $H = \lceil \sqrt{m} \rceil$  as a *chain* statement — derive the  $P$ -variance scaling (measured small and large-lag structureless) from the chain/source, source-pure and window-independent, so the bound closes the L2 lemma (H5, address c3,  $\delta' > 0.21$  of available 0.45) as a theorem; kz15 needs no new mechanism (exact-finite per r270); *kill list (inherited)*: l1 pair hierarchies, state-space/KYP, s-flows, mass splits, H-peek (selection-by-answer, sealed away), fit primitives; world-blind validity is *admissible* here (the declared difference to the  $L^*$  lane); *state (wave 9, [verification/v966\_l2\_reduction\_chain.py], block below)*: the remaining input stands **reduced to one inequality** — the four-round chain  $r297 \rightarrow r300$  froze the target level  $\sigma^* = -0.516$ , proved the exact sign-preserving transfer to the difference measure, settled the ratio half structurally and left NEFF-TARGET (slope( $n_{\text{eff}}$ )  $\geq +0.908$ , measured  $+0.963$ , margin 0.055; PRIME.L2.NEFF\_TARGET.01 [O]); *state (wave 10, [verification/v967\_l2\_cascade\_closure.py], block below)*: the reduction route is **documented stopped** — the r303 regress audit retyped the cascade  $r297 \rightarrow r302$  as an exact dictionary around one measured core  $S = \sigma^* - \text{sl}_D = +0.0547$  (REGRESS-CONFIRMED) and the r304 stop case fired (LAW-LONGRANGE: the global lag profile is a stable period-4 comb, no  $k_0$ ); the honest stop state is L2 generic  $\Leftrightarrow$  anti-concentration of the explicit  $dc$  block field; return later with new tools. *The successor front on SUBORDINATION*: the firewall reads the tent-split fraction *profile* metrically, not its arithmetic — the open question was the profile functional (rounds 290–295 executed and consumed in wave 8, [verification/v965\_lstar\_curvature\_arc.py], block below: the question is functionally negative-sealed over five class families and carried forward as PRIME.LSTAR.CLOSED\_FUNCTIONAL.01 [O]). Mincut base 4 / refined 5 unchanged; no RH statement.

**The  $L^*$  curvature arc and the closed-functional contract (PRIME.LSTAR.CURVATURE\_ARC.01 [E], PRIME.LSTAR.CLOSED\_FUNCTIONAL.01 [O]).** *Registered and delivered / registered* 2026-08-26 (rounds 290–295, notes DCXXVII/DCXXIX–DCXXXIII; [verification/v965\_lstar\_curvature\_arc.py] with  $32/32 + 30/30 + 36/36 + 43/43 + 42/42 + 44/44$  full record gates; the six probes embedded byte-exact in their sealed **-smoke** stage, SPEC SHAs `f953dd71/bb512c17/050821ff/33c44cc6/88c6fd1e/4d7d8095` pinned; plus the module-own exact section — the r292 polarization identity  $H(u,v) = [d^2(u+v) - d^2(u-v)]/4$  with its wrong-prefactor must-fail, the r291 budget telescope with the dropped-atom must-fail, the r293  $\alpha = 1$  simple-zero doubling law with the double-zero mutant, the sealed r294/r295 decision bars as exact logic with tipping must-fails, in pure Fractions). *Measured (the geometry of the working set)*: a soft-shouldered, strongly anisotropic TUBE in the exact  $\theta_{\text{eq}}$  coordinate (killfrac 0.38/0.62/1.00 at  $5 \times 10^{-4}/10^{-3}/2 \times 10^{-3}$ ; world axes kill 5–50x earlier; the SMOOTH axis a privileged yet gradient-orthogonal killer,  $\cos = -3 \times 10^{-5}$ ); a REAL one-degree ridge (min  $C = 185$  at factors 1.8) whose lift is a first-order budget phenomenon (one threshold in  $(1.280, 1.291]$  over all 18 matched-dose cases,  $k_{\min} = 9$ , one TOP6 retraction; no fixed point, plateau 185; MAIN-specific — EPSTEIN first-order flat); a RANK-1 DENS curvature valley (92.5%,  $\lambda_{\text{top}} = -0.418$ ; not SMOOTH, not the ridge; EPSTEIN second-order structureless); and SIMPLE  $h_{184}$  flip zeros ( $\alpha = 1.000 \pm 0.003$ ; the TOP6 retraction a second simple zero at  $f_{\text{ret}} = 7.107$ ; MSTAR-NO-LAW — no fixed safety distance in any unit). *Honest negatives (load-bearing, the frozen content)*: ALL-FUNCTIONALS-BLIND over FIVE sealed class families (local scalars best  $|\text{sp}|$  0.263; non-local/low-rank best  $0.471 < 0.881$ ; curvature two-form mixed  $0.884 < 0.907$ ; the metric-reconciled F10 — home win  $+0.024$  and partial  $+0.423$ , but r294 F10-FRAGILE: win 5/5, partial 2/5; the hardened partial-free form — r295 F10-SP-MAJORITY 14/20, margin median  $+0.028$ , min  $-0.079$ : the sealed three-clause HARDENED bar fails on two clauses); R293-LUCK (the r293 partial  $+0.423$  and the r294 5/5 were both top-of-distribution; the composition-matched subsample gains  $+0.067$  and misses the level clause by 0.004); L2-NOT-DENS + L2-VIA-CONSERVATION (near-tautological,  $\eta_0$  share  $2.9 \times 10^{-30}$ ) — F10 is NOT promoted, the sealed rejections stand

verbatim. *The open contract (CLOSED\_FUNCTIONAL, the successor question)*: exhibit a closed, source-pure functional of the tent-split fraction profile that predicts the survival depth of the working set at HARDENED grade (the sealed r295 three-clause bar:  $\geq 18/20$  wins AND margin median  $\geq +0.02$  AND no loss beyond 0.02) with a stable partial channel — or prove non-existence in a named class; *kill list (sealed, binding)*: the five class families above, MSTAR-NO-LAW, no diophantine/linear-forms route (wave-7 METRIC-ONLY inherited), no plain phase-dispersion separator, no sinc/CD-only sign mechanism, no re-training of the hash-sealed  $H_{tr}$  (3447ed198a56), no posthoc bar moves; *open fronts (named)*: the loss-corpus forensics (K05/K07/K19, margins  $-0.041/-0.057/-0.079$ ), the rank-2 DENS core (top DENS eigenaxis alone  $+0.855$ , DENS share 0.989, mechanism-constant across w7/w9/w11:  $0.795/0.989/0.834$ ), the substantive “why L2” question, and larger window corpora for the w7/w11 knife edge; *measurement standard (sealed)*: PARTIAL-STD = the partial on the full sealed 147-point fresh protocol (r295 pins: median  $+0.279$ , IQR  $[+0.182, +0.401]$ ). Lean: no new statement — the wave carries no new formal target form, the hole stays `lstar_subordination`. *State (wave 9, [verification/v966\_l2\_reduction\_chain.py])*: the DENS-identity fork is **closed** — r296 adjudicated DENS-WORLD-BLIND under the sealed rules (the coupling number  $\cos(e_{top}, \nabla \lambda_{max}) = +0.394$  below the sealed 0.40 bar, noise max 0.331, miss 0.006 honest; the  $\lambda_3$  control couples *harder*,  $-0.574$ : a near-crossing eigenvalue-band property, not top-specific; every arithmetic candidate  $\leq 0.38$ ; the moment subspace (0.825) construction-adjacent, holding on both dead controls  $0.571/0.603$ ): front (2) — the rank-2 DENS core — is answered negatively at the identity level (a mechanism hint, not an arithmetic identity, not the  $L^*$ -gradient direction), the lane routed to L2 per the pre-adjudicated reviewer fork; the remaining open fronts are (1), (3) and (4). *State (wave 11, [verification/v968\_architecture\_adjudication.py])*: **Lane A (the constructive PSD-G rule fixed by the r308 feasibility census) is closed as a cone language** — r311 STRICT-SOURCE-CONE (exact rational in-span Farkas certificates on  $MINI\#0 + SM1\#0$ ; the r308 world separation demarcated as the budget sign / a chordal restatement) and r312 COEFFICIENT-SIGN-WALL with the mechanism *named* (block-psd membership without rank-one-SDD membership, obstructed at the budget/border sign; Farkas mass on the border/budget row,  $t$ -share  $0.583/0.725$ ) — no GO per the reviewer tree, the resources moved to the fiber; the PSD-G rule stays open as a *question*, closed as a *lane*; the closed-functional demand itself is untouched (the sealed r295 bar stands). Mincut base 4 / refined 5 unchanged; no RH statement.

**The L2 reduction chain and the NEFF target contract (PRIME.L2.REDUCTION\_CHAIN.01 [E], PRIME.L2.NEFF\_TARGET.01 [O]).** *Registered and delivered / registered* 2026-08-26 (rounds 296–300, notes DCXXXV–DCXXXIX; [verification/v966\_l2\_reduction\_chain.py] with  $39/39 + 27/27 + 28/28 + 32/32 + 31/31$  full record gates; the five probes embedded byte-exact in their sealed `-smoke` stage, SPEC SHAs `ffb413c8/e42a76eb/05e831be/f432e944/55218b5d` pinned; plus the module-own exact section — the  $\sigma^*$ -composition arithmetic ( $\delta'(\sigma^*(pad)) = pad$  exactly, the pad-dropped mutant composes to 0), the decomposition identity  $S_F = B(\omega, \omega) + B(\Delta, \omega + \beta)$  on the exact positional Fejér block form (toy  $3 = 1 + 2$ ; the  $|\Delta|$  and wrong-weight mutants break by exactly 4 and  $2/3$ ), the participation identity  $D \times n_{eff} = L_1^2$  with ratio multiplicativity and the machine-checked  $DIAG \leftrightarrow NEFF$  equivalence (margin 0.055 exact in both coordinates; signed-sum mutant  $8/3$ ), the kernel envelope  $F_H(s) \leq \min(H, 1/(Hs))$  exact-rational with equality witnesses on both branches, and the five sealed verdict bars as exact decision logic with tipping mutants, in pure Fractions). *Delivered (the chain)*: the r296 DENS fork closed honestly (DENS-WORLD-BLIND, above); the r297 target inequality frozen ( $\sigma \leq \sigma^* = -0.516$ , measured  $-0.714$ , margin 0.198 — provenance missing, not room; the chain exactly orthogonal on the WINDOW measure); the r298 exact sign-preserving transfer (TRANSFER-DOMINANT: the window main term empty,  $S_F \approx B(P\Delta, P\Delta)$  — the vdC input IS the Fejér energy of the difference measure); the r299 decay split (LOWPASS 0.93; full-support overlap  $42/42$  —  $\Delta$  a pure c-value difference; ET/Abel dead  $+1.948$ ; the live edge  $sl_D = -0.571 \leq \sigma^*$  with falling ratio  $-0.168$ ); the r300 diagonal anatomy ( $sl_D = 2sl_{L1} - sl_{neff}$  exact; both magnitude routes closed honestly — the  $|dw|$  sum-rule attach BREAKS, the fill decay  $-0.225$  is invisible to  $\max \times \text{mass}$  bounds; RATIO-BOUNDED-STRUCTURAL:  $R_{env}$  med 1.61 falls at  $-0.122$ ). *Honest*

*negatives (load-bearing, the magnitude catalog):* r297 B1 pair-gap ( $-0.535$  vs needed  $\leq -1.178$ ) and B2 mass imbalance ( $+0.244$  growth); the r299 ET/Abel wall ( $+1.948$ , MASS-TARGET missed by  $2.46$ ); the r300 chain-norm and  $\max \times L_1$  closures; no pointwise c-value convergence (cconv  $0.86$ ,  $+0.045$ ); the lobe-width heuristic refuted ( $\text{sp} = -0.47$ ). *The open contract (NEFF\_TARGET, the one remaining inequality):* prove  $\text{slope}(n_{\text{eff}}) \geq 2\text{sl}_{L1} - \sigma^* = +0.908$  for the participation number  $n_{\text{eff}} = L_1^2/D$  of the aggregated difference profile (measured  $+0.963$ , margin  $0.055$ , thin) — exactly equivalent to DIAG-TARGET  $\text{sl}_D \leq \sigma^*$ , and through the settled chain it closes  $\delta' > 0.21$  on the generic half:  $[\text{NEFF-TARGET}] \Rightarrow \text{sl}_D \leq \sigma^* \Rightarrow \sigma \leq \sigma^* \Rightarrow \text{r297 target} \Rightarrow \text{v964-S0 vdC} \Rightarrow \delta' > 0.21$  (exceptions: 6 via r287 F2, kz15 exact-finite via r270); *the named bridge (correlational, not causal):* the carrier count grows where the boundary discrepancy falls ( $D_{\text{rank}}(\Delta)$  slope  $-0.117$ ,  $\text{sp}(D_{\text{rank}}, n_{\text{eff}}) = -0.81$ ) — a proof must turn the correlation into a mechanism; *kill list (measured dead, binding):* every max/mean magnitude majorization (r297 B1/B2, r300 chain-norm CS on  $|dw|$ , r300  $\max \times L_1$ ), the ET/Abel position-bound composition, pointwise c-value convergence, the lobe-width heuristic, H-peek, fit primitives; world-blind validity admissible (L2 is the generic half);  $D_{\text{rank}}$  statistics typed MEASUREMENT-ONLY (forbidden as proof premises). *State:* round 301 (NEFF-EQUIDISTRIBUTION) was in flight on exactly this question at the wave-9 cut; *wave-10 resolution (rounds 301–304 executed and consumed, [verification/v967\_l2\_cascade\_closure.py], block below):* the demand is **retyped to the documented stop state** — the r303 regress audit shows NEFF/UNIF/ATOM are *one* algebraic slack  $S = \sigma^* - \text{sl}_D = +0.0547$  (invariance  $\leq 9 \times 10^{-16}$ ; the  $\frac{1}{2}$ -conversion refuted), so the rest is no longer a ladder of targets but the stop state: L2 generic  $\Leftrightarrow$  anti-concentration of the explicit *dc* block field (core slack  $S = +0.0547$  measured); the global-profile mixing route is closed (period-4 comb, LAW-LONGRANGE); the mechanism is two-scale; return later with new tools — *the hard reviewer rule binds: a round counts as progress only if it adds new information, not if it re-expresses the same slack in a new coordinate.* Lean: no new statement — NEFF-TARGET is a measured ladder-slope aggregate, not an exact finite identity, and a universal ladder form would be refutable (the r273 guard pattern); the hole stays **lstar\_subordination**. Mincut base 4 / refined 5 unchanged; no RH statement.

**The cascade closure and the documented lane stop (PRIME.L2.CASCADE\_CLOSURE.01 [E]).** *Registered and delivered* 2026-08-26 (rounds 301–304, notes DCXL/DCXLII–DCXLIV; [verification/v967\_l2\_cascade\_closure.py] with  $32/32 + 30/30 + 26/26 + 37/37$  full record gates; the four probes embedded byte-exact in their sealed **-smoke** stage, SPEC SHAs **6f8cc404/36df9424/375e9f2b/2cc5d23f** pinned; plus the module-own exact section — the r303 margin-invariance algebra (all four target margins  $m_D = m_{\text{NEFF}} = m_{\text{UNIF}} = m_{\text{ATOM}} = 3/50$  exact on the rational slope set; the halved-need mutant breaks by exactly  $\text{sl}_{L1} = 1/5$ ; the  $\sigma$  decomposition  $0.198 = 0.055 + 0.143$  exact; the  $\frac{1}{2}$ -conjecture refutation  $0.099 \neq 0.055$ ), the r302 coherence identity  $(1 + CV^2) \text{surv}^2 n_{\text{eff}}^{\text{atom}} = n_{\text{act}} \chi$  on two rational toys (mutant breaks  $5/7$  and  $48/35$  exact), the r304  $\chi$ -lag decomposition  $\chi = 1 + 2 \sum_k T_k/Q = D/Q$  on both toys ( $5/7$  and  $5/19$  exact; factor-1 mutant  $1/7$ ), the r304 zero-sum tautology as a gate (mutants  $1/2$  and  $3/4$  exact;  $NC(1) = 0.556$  exact,  $NC(16) = 0.712 < 1$ ; the period-4 comb certificate with no  $k_0$ ), the  $-/-/+ /+$  sign-pattern certificate with the kz18/kz23 Fractions pins, the four sealed verdict bars as exact decision logic with tipping mutants, and the lane-stop composition gate, in pure Fractions). *Delivered (the rejections and retypings ARE the content):* r301 NEFF-SPLIT ( $n_{\text{eff}} = n_{\text{act}}/(1 + CV^2)$  exact with the perfect count link  $n_{\text{act}} = m$  on  $42/42$ ; the Rényi family at one exponent — echt anti-concentration; jackknife-stable; rest relocated onto UNIF-TARGET); r302 UNIF-DERIVED(B1+A2) — the *first* DERIVED of the lane, all hypotheses typed MEASURED (PROFILE-STATIONARY, KS  $0.043$ ; exact  $1/N$  transient onto  $m_2^\infty = 1.973$ ;  $\chi = 0.63$  destructive, falling; rest relocated onto ATOM-TARGET); r303 REGRESS-CONFIRMED + MIXING-INSUFFICIENT (the retyping: all four margins ONE number  $S = +0.0547$ , invariance  $\leq 9 \times 10^{-16}$ ; the cascade an exact dictionary, not six proof steps; the first causal coordinate — the  $\rho_1$  ladder is monotone and the flip kills the inequality at  $-0.044$  — but the  $\chi$  level misses by  $0.134$ ;  $n_{\text{eff}}^{\text{atom}}$  a pure marginal functional,  $10^{-15}$  on all builds); r304 LONGRANGE-STRUCTURE + LAW-LONGRANGE (the reviewer stop case: a stable, world-specific period-4 comb, no  $k_0 \leq 8$ ; the reviewer condition splits —  $NC(16) = 0.712 < 1$  holds, summability

fails at 1.563; the  $\chi$ -relevant structure is exactly short-range, the r303 gap at  $k \leq 3$ , lag-2-dominated; lag-8 matching hits the level, breaks the slopes: the mechanism is **two-scale**; the graduated ladder is monotone in  $\chi$ , not in the margin). *The consequence (sealed map (iii))*: **the global-profile mixing route of the L2 lane is documented closed** — L2 generic  $\Leftrightarrow$  anti-concentration of an explicit block field with long-range period-4 structure; return later with new tools. *What stands honestly*: the exact identity dictionary (count, coherence,  $\chi$ -lag), the two-scale split,  $NC < 1$ , the  $-/-/+ / +$  sign pattern (Fractions-exact kz18/kz23), the marginal-functional invariance ( $10^{-15}$  on all 2100 + 1008 builds). *The hard reviewer rule (binding)*: a round counts as progress only if it adds new information, not if it re-expresses the same slack in a new coordinate — no further coordinate change counts. Lean: no new statement — the stop carries no new exact target form (measured ladder aggregates; the r273 guard pattern); the hole stays `lstar_subordination`. Mincut base 4 / refined 5 unchanged; no RH statement.

**The architecture adjudication and the wave-11 claim split** (PRIME.LSTAR.ARCHITECTURE\_ADJUDICATION.01 [E], PRIME.SOURCE.RANKONE.UPDATE.IDENTITY.01 [E], PRIME.SOURCE.TENSOR.MECHANISM.01 [O], PRIME.L2.RENYI3.PROVENANCE.01 [O]). *Registered and delivered / registered* 2026-08-26 (rounds 305–316 incl. 310b, notes DCXLV–DCLVIII; [verification/v968\_architecture\_adjudication.py] with 27/27 + 28/28 + 30/30 + 33/33 + 34/34 + 32/32 + 32/32 + 29/29 + 29/29 + 35/35 full record gates; the TEN probes embedded byte-exact in their sealed -smoke stage, SPEC SHAs 3bb365e1/ec2bb008/d5147850/f8d99877/fac7a8df/6c32f749/6505dd10/841b3196/92d35a3a/5c28b12b pinned; plus the module-own exact section — the r306 Hill/Lagrange chain with both witnesses and the Rényi bridge with all constants, the r314 opening-flux lemma  $F_1 \equiv 0$  with the exact flux telescope, the collision counting identity  $3p_1p_2 - 2p_3 = 8(n + 3n(n-1))$  with the exact factor-8 mutant break, an r311-class exact rational Farkas mini-certificate (a PD target in the span with a forced negative generator weight, strictly outside the cone), the TEN sealed verdict bars as exact decision logic with tipping mutants, and the four-level composition gate, in pure Fractions). *Delivered (the reviewer's binding four-level structure)*: **Level 1** — real formal theorems (Lean, already green; no Lean change ships: `lstar_terminal_implies_master`, the four closed Inertia theorems, the real PrimeWindow construction, source exactness + mass conservation of any folding + the strong stabilization form, the r309 rank-one update identities, the Hill monotonicity + Rényi-3  $\Rightarrow N_2$  bridge; sorries  $9 \rightarrow 5$ , the two TRUE holes standing alone). **Level 2** — certified finite statements (Rényi-3 on 57 rungs at  $C = 1.069$ ,  $A = 2$ ; the fixed-head census 77/77; the block-Green identity + the strict cone with exact rational Farkas certificates; paired-cone soundness on 20714 steps; fold multiplicity exactly 2; 93% far-triple cancellation; the signed cubic identity with vanishing boundary, Fractions-exact on the real data; the 21 small- $m$  certificates). **Level 3** — negative architecture decisions verbatim (FIXED-HEAD-DEAD; PAIRED-CONE-NO-INDUCTION + WORLD-BLIND — B does not overtake A; the r308 discriminator = the budget sign / a chordal restatement; COEFFICIENT-SIGN-WALL — Lane A closed as cone language with the mechanism named: block-psd membership without rank-one SDD, obstructed at the budget/border sign; FLOQUET-EXPANDING; TRIPLE-TYPE-MAJORANTS-WRONG; PHI3-ALL-BLIND; TWO-REGIME-DEAD with the anatomy: the obstruction family cuts across the FCIX stratum, the first-5 protocol load-bearing, kz55/kz67 near-one-block). **Level 4** — open mechanisms typed (the constructive PSD-G rule as a question, the signed cubic flux theorem, the quantitative source functional, `crossing_budget`, the definitional source bridge,  $L^*$ , the terminal). *The claim split (binding — not one TRANSFER title)*: RANKONE.UPDATE.IDENTITY [E] carries *only* the exact update identities (determinant dictionary, Sherman–Morrison chain,  $\pm(\eta \pm e_t)$  border split, local Schur reserve, signed source updates); TENSOR.MECHANISM [O] carries the open question whether a source-pure cofinal control arises from them, registered *after* r313/r316 with the documented negatives. *The provenance contract (RENYI3.PROVENANCE)*: the pointwise GO stands as a certified finite statement ( $\sum q^3 \leq 1.069(\log m)^2/m^2$  on 57 rungs, growing reserve; fiber target  $n_{\text{eff}} \geq m/(1.034 \log m)$ ); the provenance is open — trilogy state: the identity exact (r314), the functional blind (r315), the two-regime form dead (r316); the R317 fork verbatim (reviewer adjudication pending): a near-critical family coordinate (TOP1/QMAX source-pure in

advance; kz53 needs a second coordinate) *or* a certified exception family (the r287-F2 pattern). *Framing*: the universality-class insight (base fine-metric — METRIC-ONLY, twin invariance; fiber recursive — EPSTEIN holds, SCRAMBLE breaks; shared algebra, not necessarily shared mechanism); the reviewer maturity raise  $4 \rightarrow 5$  (formal architecture 9.5) documented as a *reviewer assessment*, not a machine-checked fact. Mincut base 4 / refined 5 unchanged; no round in flight at this cut — R317 waits on the reviewer fork; no RH statement.

**The extraction-chain red team, the sliding bound, and the wave-12 adjudication (PRIME.REDTEAM.EXTRACTION\_AUDIT.01 [E], PRIME.L2.RENYI3.SLIDING\_BOUND.01 [O]).** *Registered and delivered / registered* 2026-08-27 (rounds 317–322, notes DCLXI–DCLXVI; [verification/v969\_forks\_and\_redteam.py] with  $38/38 + 25/25 + 39/39 + 25/25$  full record gates; the FOUR probes embedded byte-exact in their sealed -smoke stage, SPEC SHAs 04f5e5c0/f2d98683/e68883ad/761b51d4 pinned; plus the module-own exact section S0 — the two-sided concentration bracket  $q_{\max} \Phi_{H1} \leq \rho_2 \leq \Phi_{H1}$  re-proved from scratch with both mutant directions caught, the U1–U3 red-team witnesses in exact arithmetic (the  $B = -1$  window against the terminal type; the mesh-level-0 total node collision reduced to three rational facts about  $e$  on the certified bracket  $685/252 \leq e \leq 685/252 + 1/35280$ , kill  $0 < 0$  exact; the adversary  $(|F| + 1)^2 < 5/7$  false for every  $F$ ), the r320 separation/canonical-split certificates (witness mesh level 4 via  $2^7 < 3^5 < 2^8$ , reduction  $1/8 < 2/5$  exact; the level-0 failure), six sealed verdict bars with tipping mutants and the wave-12 composition gate, in pure Fractions; no Lean call — the r319/r320 Lean rounds are consumed as *reports*, their artifacts in `rh/lean/` re-verified by `lake build` and the RH suite). *Delivered (the audit and the repair)*: the r319 reconstruction found three kernel-checked *type inconsistencies* three levels above the boxes (U1: the r310b bridge bound neither  $u$  nor  $B$  — a  $B = -1$  window represents the empty spec against `terminal_positive_main`; U2: the mesh-level-0 tolerance admits the total node collision and  $p = X - 1$  forces  $0 < 0$  against the  $L^*$  type; U3: the pre-r320 `pair_margin_main` forces its source predicate empty) plus the mesh-vs-anchor cofinality seam — **the false cofinality direction is documented, not solved**; r320 repaired the interface completely (bridge retype with  $u/B$  fidelity + separation discipline + `budget_pos`; `SourceExact` for the additional arch/border finding; the canonical extraction with `canonical_split proved`; three permanent sorry-free guards; the sorry-free witness; census  $5 \rightarrow 5$  with two typed retypes). *The honest chain reading (binding)*: “two lemmata  $\Rightarrow$  RH” does not survive the literal reading — the two lemmata give *only window-local master positivity* for the opaque, currently witness-less `MainWindow`. *The three-level proof graph (reviewer adjudication, binding; the two holes are named **window-local** everywhere from wave 12 on)*: Level A — local window mathematics, *proved* ( $L^*(w) + T(w) \Rightarrow \text{MasterPositive}(w)$ ); Level B — the universal window family, the two arithmetic window-local holes; Level C — extraction to Weil/RH, *open* (source-exact realization + admissible transport + convergence/stabilization in the correct directed order + density/continuity). Two binding caveats: the sorry census does *not* measure the Level-C distance (the missing extraction architecture appears in no `sorry`), and `SourceExact` must not stand as a free assumption — the named open Lean target is `CanonicalPrimeWindow` + exactly one construction theorem `sourceExact_buildPrimeWindow` (load-bearing; the extraction-repair fork R325 in flight, not consumed). *The fiber fork (SLIDING\_BOUND)*: fork (b) sealed WHAC-A-MOLE(kz53, kz83) by contract (the spike family  $B = \{kz55, kz67\}$  is real,  $C_B = 1.0536$ ; the  $F_A$  top-3 are exactly the near-critical family, but as a *continuum*); fork (a) fired SLIDING-BOUND-GO( $G_{SQ}$ ):  $\sum_j q_j^3 \leq 1.3056 F_A(w)^2 (\log m)^2 / m^2$  for  $m \geq 73$ , 0/39 test violations, all four named violators inside at reserves 7.0–9.6 (honest:  $C_{\text{impl}} = 7.97$  is  $7.5\times$  looser — form, not sharpness; the demand is the two-part  $F_A/q_{\max}$  provenance split; the terminal sharpened to the  $Q_{\text{MAX}} \times M_2$  shape route — R324 in flight, not consumed). *The typed reviewer assessment (supersedes the wave-11  $4 \rightarrow 5$  scalar)*: finite dictionaries 9.5/10, formal window-local architecture 9/10, terminal mechanism 6/10,  $L^*$  mechanism 2/10, source/extraction architecture 2–3/10, complete RH proof path  $\sim 3/10$ , audit quality 9/10 — “stronger as a window-local research program; as a complete RH proof chain currently less far than previously assumed — the removal of a shortcut that did not exist.” Mincut base 4 / refined 5 unchanged; no RH statement. *State (wave 13, [verification/v970\_extraction\_and\_composition.py])*: both in-flight rounds are executed

and consumed — the extraction-repair fork R325 fired ELEMENTWISE-STABILIZATION-GO and the named open Lean target is *implemented and proved* in r326 (`sourceExact_buildPrimeWindow`; census  $5 \rightarrow 8$  with three named typed Level-C statements — the census caveat partially discharged; see the wave-13 block below); the terminal  $Q_{\max} \times M_2$  route composed MEASURED in R324 (PILEUP-GROWS-SUBCRITICAL +0.172 — subcritical also at the r328B chain-honest bar 0.188;  $m_0^*$  record  $10^{59.6}$ , chain-honest  $10^{238}$ ; PRIME.L2.RENYI3.MEASURED\_COMPOSITION.01 [O]).

**The elementwise extraction repair and the measured terminal composition — the wave-13 block** (PRIME.EXTRACTION.ELEMENTWISE.01 [E], PRIME.L2.RENYI3.MEASURED\_COMPOSITION.01 [O]). *Registered and delivered / registered* 2026-08-27 (rounds 323–327 incl. the r326 Lean round, notes DCLXVII–DCLXXII; [verification/v970\_extraction\_and\_composition.py] with 36/36 + 36/36 + 18/18 + 34/34 full record gates; the FOUR probes embedded byte-exact in their sealed `-smoke` stage, SPEC SHAs 9a6696f8/dc36cacb/31277f91/11e4fd40 pinned (r324 imports the embedded r324-pre, r327 the embedded r324-pre and r324); plus the module-own exact section S0 — the r324 interpolation  $M_3 \leq q_{\max} M_2$  with slack exactly  $1/36$  and double equality at the flat profile, the  $F_A$  identity  $F_A B \log m = m q_{\max}$  as exact constructional algebra, the r327 group ledger with the two-ancestor bound (both mutants break by exactly 1), the r325 onset formula  $\alpha^* = (n_g + 1)D_0/2$  as an exact toy stabilization, the  $m_0^*$  log-space solver verbatim (landing  $10^{59.6}$ ; the disclosed pre-freeze linear-space bug rebuilt and caught), the r328B chain-audit gate S0-T5d, six sealed verdict bars with tipping mutants and the wave-13 composition gate, in pure Fractions; no Lean call — the r326 Lean round is consumed as a *report*, the r323 abort as a *note*: the planned L\* fork was cleanly aborted before any write access, nothing counts as a round). *Delivered (ELEMENTWISE, the extraction repair)*: r325 adjudicated the three reviewer variants for the R319 cofinality seam — ELEMENTWISE-STABILIZATION-GO: on the native v749 class all three channels (comb/arch/pole) of the canonical tower windows stabilize *exactly* at the predicted finite anchor onset and are constant under dyadic mesh refinement — no mesh-cofinal ladder, no transport,  $H_{\text{cof}}$  **is not needed and is replaced as the target route** (variant B = the construction substrate, variant C = the quantified density step inside the rigorous interpolation envelope, no wall margin anywhere); r326 implemented the architecture (`rh/lean/RH/Elementwise.lean`, merged and green): `sourceExact_buildPrimeWindow` *proved* (SourceExact eliminated as a free assumption from the extraction route), the comb stabilization elementwise *proved* (onset predefined from the element's support alone), `weil_nonneg_of_windowlocal` *proved* (the extraction without the ladder: one finite instantiation per element); sorry census  $5 \rightarrow 8$  with three named typed Level-C statements (arch/pole transcription classical, source completion definitional). *The honest formulation of the new connection (binding, verbatim)*: “two window-local lemmata (over the canonical family) + the elementwise architecture + cited classics (density, v912 continuity, the Weil criterion)  $\Rightarrow$  RH; open links: the arch/pole transcription (classical), the compression bridge, the source completion (definitional).” *Registered (MEASURED\_COMPOSITION, the terminal end state)*: r324 — after the binding mid-round course correction, with the pre-work banked as r324-pre (FA-BOUNDED-DISTRIBUTIONAL + QMAX-SHARE-OPEN + NOT-COMPOSED;  $F_A$  de-black-boxed by the exact identity  $F_A B \log m = m q_{\max}$ ) — fired PILEUP-GROWS-SUBCRITICAL ( $e_{\text{tot}} = +0.172 < 0.224$ , decided): the  $O(\log m)$  scale count certifies pointwise ( $C_{\text{NSC}} = 2.0258$ , 0/39), the per-scale mass is refuted — the near-critical windows are *one heavy source scale*; the composition typed MEASURED:  $\sum_j q_j^3 \leq 8.941 (\log m) m^{+0.172}/m^2 \Rightarrow N_2 \geq m^{0.888}$  for all  $m \geq m_0^* = 10^{59.6}$ , the gap (274,  $10^{59.6}$ ) the disclosed extrapolation hypothesis — no cofinal claim (the independent chain audit r328B, carried as a chain annotation with the record sealed: the bar 0.224 used the atom-level need 0.888, the block-level need is  $908/1002 \approx 0.9062$  — the chain-honest bar is 0.188, the verdict survives  $0.172 < 0.188$  exactly, the chain-honest  $m_0^*$  is  $10^{238}$ , the going-forward bar is 0.188; exact v970 gate S0-T5d); r327 located the spike structurally (CAP-PARTIAL, anatomy SOURCE): on all 65 rungs the heaviest group of the argmax block is exactly *one*  $\beta/\omega$  fold pair (kz53 = one bulk/window coincidence at 88.8%); the  $\Lambda$ -pair cap refuted, **the group count certified** ( $C_{\text{NG}} = 2.6351$ , 0/39, named 4/4 — the third certified

$O(\log m)$  count); *the demand (the coincidence geometry)*: where can the single heavy bulk/window coincidence *sit* — a source-geometry question, not group count (certified) or multiplicity (proven); alternatively the MEASURED composition is the honest end state of the route. *Honest boundary*: the window-local positivity premise is proved nowhere; the two true window-local Lean holes are unchanged;  $L^*$  stays paused; nothing bounds anything beyond the measured rungs. Mincut base 4 / refined 5 unchanged; no RH statement.

**The wave-14 consolidation — terminal martingale, cover/growth/K2, the closed  $L^*$  margin chain, the Borodin duality and the second arithmetic** (PRIME.TERMINAL.DENSITY\_MARTINGALE.01 [E], PRIME.L2.COVER\_GROWTH\_K2.01 [E], PRIME.LSTAR.MARGIN\_CHAIN.01 [E], PRIME.LSTAR.DUAL\_HOLE.01 [E], PRIME.WORLD.DIRICHLET\_MATCHED\_FRAME.01 [E]). *Registered and delivered* 2026-08-27 (rounds 334–357, notes DCLXXXI–DCCIX; reviewer-authorized after the double substance condition r355 + r357; five modules [verification/v978\_terminal\_density\_martingale.py] / [verification/v979\_cover\_growth\_k2.py] / [verification/v980\_lstar\_margin\_chain.py] / [verification/v981\_lstar\_borodin\_duality.py] / [verification/v982\_dirichlet\_matched\_frame.py]; each re-derives its exact layer from scratch — sympy symbolic + pure Fractions, no probe imports — and re-runs the sealed round verdicts as exact decision logic on the frozen record aggregates with tipping mutants; the eighteen discovery probes stay sealed in `experiments/tfpt-discovery/`, pinned and re-verified by the RH suite). *Delivered (theorem-grade exact layers)*: the density martingale from mass conservation alone with the moment dictionary  $E[X_\infty^3] = m^2 M_3$ , the tilted tower, the exact hand-off and the pair ceiling  $\Gamma \leq 4$  (v978); the FAB identity, the data-free K1 formula  $R_{\text{ALG}} = 4^{1/3} \rightarrow 8/5$  and the mesh identity  $h - \nu u \in (0, 3/2]$  (v979); the one-line identity, the two-level dressed-scalar theorem, the pinning theorem ( $m'_2 = \text{margin}$  identically under orthonormal top-2 geometry), the rate-equality theorem and the  $\rho_K$  identities (v980); the Borodin complementation at half filling with rational conjugator, the equivalence  $L^* \Leftrightarrow R > \frac{1}{2}I$ , and the reciprocal dual weight  $u^\vee \propto c_j(1 - x_j)/|f|$  that makes the r354 anti-correlation *duality algebra by design* (v981); the derived  $\chi$ -archimedean side with bitwise trivial reduction, exact conductor shift and parity kernel (v982). *Registered (census, honest)*: the first certifying three-arm cover ( $m_0^* = 10^{22.6}$ ), the P02 predictor, the growth ceiling  $C_{\text{FAB}} = 14.93$  and every  $m_0^*$  of the lane are *Frame-A census* (binding r353: the second construction family breaks the ceiling and kills the reserve floor at  $m \sim 760$  — the sliding coverage is finite); K2 ( $\leq 11.87$ ) is the sole cross-family survivor, NU-free over three aspects and 0/126 on the second arithmetic; the  $L^*$  margin chain is complete and typed (weights dictionary / deficits candidate law / pinning theorem /  $\rho_r = \varphi$ -difference computable), the duality verdict is DUALITY-REPARAM-ONLY, and the lane is *frozen* as a specialist problem at `rh/problem/lstar_problem.tex` with the one-object question halved (the  $\sim 500\times$  smallness of by-design-coupled blocks, localized in local dual two-point data); SECOND-ARITHMETIC-LIVES: the r330 death was a frame artifact, the wall holds on  $n = 3$  arithmetics on 126 windows and the matched scramble breaks — the wall needs the arithmetic; the zeta-graded signatures (exact pinning, saturated reserve,  $\varphi$  suppression, counting constants) type the  $L^*$  margin legislation as a statement about how close zeta's world sits to the wall. *Honest boundary*: no cofinal claim anywhere (the r324 MEASURED composition stays the terminal end state); the equivalence is a reparametrization of the open condition; PRIME.LSTAR.SUBORDINATION.01 and PRIME.L2.RENYI3.SLIDING\_BOUND.01 stay [O] with wave-14 notes; everything census-side is finite, kein Satz. Mincut base 4 / refined 5 unchanged; no RH statement, no GRH statement.

**The kernel-Loewner positivity certificate at  $L = 0.3$  (PRIME.RDAGGER.KERNEL\_LOEWNER.01, Numerical/certified).** *Registered and delivered* 2026-09-01 (rounds 494 + 495; module [verification/v1017\_kernel\_loewner\_positivity.py]; re-derived, no probe imports). *Delivered*:  $Q_W(h) \geq 2.1 \times 10^{-3} \|h\|_2^2$  on  $\text{supp}(h) \subset [-0.3, 0.3]$  ( $2L = 0.6 < \log 2$ , prime term empty); float64 enclosed floor  $2.122 \times 10^{-3}$  after a  $3\times$  Hilbert–Schmidt tail charge — *not* interval arithmetic and *not* [E](repo convention: positivity [E] is Fractions or integer-Sylvester/mpmath.iv); G1 calibration vs defining digamma ( $\sigma_A(0) = -5.3721834192256654$ ); G2 Loewner after zero-extension; independent r495 translation identity 6/6 exact over  $\mathbb{Q}$  and doubled- $c_L$  false-world budget  $-2.188$ . *Honest*

*boundary:* r496 NO\_GO(compact-tail@ $L = 0.8$ ) is the named method boundary — not  $\lambda_*(L) \geq 0$  in general; no cofinal claim; no RH statement.

**The gapcode disenchantment (PRIME.DIVISOR210.GAPCODE.01).** *Registered* 2026-08-09, an honest negative/limit row ([verification/v886\_register\_gapcode\_null.py], 24/24, GAPCODE-MEASURED; five honest SPEC amendments on record, none loosening a criterion). *Measured:* the error-code reading of the 210-register is literally true on the prime side — the register bits  $\{2, 3, 5, 7\}$  are the first four parity checks of the sieve (rate 48/210; 16 syndrome labels; 0 bits of message capacity); the hard floor  $g_k \geq w(p_k \bmod 210)$  holds with zero violations in  $5.76 \times 10^6$  consecutive pairs;  $I(\text{state}; \text{gap}) = 1.14$  bits — and the code is the sieve: a zero-parameter wheel-Cramér null explains  $R^2 = 0.99839$  of the state profile, the arithmetic excess is 4.02% and shrinking ( $21.81 \rightarrow 4.02\%$  over  $X = 10^5$ – $10^8$ ), no knee at four checks, 210 on the foreign-wheel trend ( $z = -0.61\sigma$ ), all eight anchor Walsh seats below  $2.5\sigma$  exactly as CRT predicts. *The bound:* everything the register knows about prime gaps, every 4-prime wheel knows; nothing moves with  $c_3$  or  $\varphi_0$ ; the register’s uniqueness remains the arithmetic forcing  $\prod(p-1) = 48$  (v868). No marker move; no RH statement.

**The rounds-42–49 consolidation (the complete certified ladder, the finite Pfaffian/ $K_6$  signature, the four route decisions, the surviving positivity architecture, and the Möbius kills).** *Recorded* 2026-08-09 (a campaign note, round CVIII; fourteen frozen sandbox probes, every spec amendment fail-first on record; the promoted verification modules [verification/v887\_certified\_ladder\_complete.py], [verification/v888\_finite\_pfaffian\_signature.py], [verification/v889\_route\_decisions.py], [verification/v890\_total\_positivity\_flow.py] and [verification/v891\_moebius\_firewall\_kills.py] embed their discovery probes byte-exact). The campaign completes the certificate rollout ( $\sigma_h > 0$  proven on all 42 reachable rungs), closes the finite Pfaffian/ $K_6$  signature of the compiler (the canonical sign character; the self-hosting counting theorem; the flavor hexagon as the automorphism orbit), decides the four routes of the sum-rule program (the wall dies off-diagonally; PNT is insufficient; the diagonal input is pair-correlation class — the sum-rule contract retyped conditional), measures the surviving positivity architecture (flag positivity – all leading principal minors, Sylvester; typing corrected per the round-51 review, classical total positivity is not claimed; the wall margin as the pole distance of the  $s$ -flow; multi-factor inheritance), and honestly kills the Möbius/carrier-invariance route with mechanisms. The registrations and updates follow as the blocks below (plus the dated retypes on the sum-rule and tail contracts above). Full narrative in the prime-front paper and Paper 1; audit cards in Paper 3. No marker moves; no RH statement.

**The complete certified ladder (PRIME.PORT.CERTIFIED.LADDER.01, delivered).** *Registered and delivered* 2026-08-09 ([verification/v887\_certified\_ladder\_complete.py], 7/7, CERTIFIED-LADDER-COMPLETE). *Proven:*  $\sigma_h > 0$  on all 42 reachable rungs of the deployed window ladder — 12 rungs ( $h \leq 300$ ) by exact-rational integer Bareiss/LDL certificates, 30 rungs by validated-precision mpmath Cholesky at dps 120/200 (precision doubling; every pivot  $> 10^6 \times$  the accumulated per-pivot rounding bound at both precisions); 0 refusals, 0 out-of-reach; the module closes the two deepest rungs  $h = 859/878$  with floors  $\geq 8.352/9.187 \times 10^{-6}$  and pivot-to-bound ratios  $\sim 10^{115}/10^{195}$ ; the Epstein control refused at pivot index 10 by the identical machinery; the 40-rung parent rollout stays sandbox evidence, runtime-excluded. *Typed honestly:* the validated-precision tier is never labeled exact-rational; proven modulo the declared conservative lag-evaluation error model (validated-numerics standard); the ball-arithmetic closure and the Lean composition through krein\_cofinal\_weil remain the named steps of PRIME.PORT.TAIL.01; the asymptotic tail is untouched. *Interval note (2026-08-09, round CXV):* the error model is upgraded — [verification/v897\_certified\_interval\_ladder.py] (PRIME.PORT.BALLLADDER.01, below) re-proves all 42 rungs with rigorous interval-arithmetic shifts (15 exact-rational + 27 validated-precision; Epstein refused at pivot 10 by the identical interval machinery); the informal conservative  $\varepsilon_c$  model is retired from the ladder, and only the Lean composition remains named. *Kill:* a certified pivot failure on any rung under the frozen error model. No marker move; no RH statement.

**The diagonal separation measurement (PRIME.PORT.DIAGSEP.01).** *Registered* 2026-

08-09, measured ([verification/v889\_route\_decisions.py], 8/8, DIAGSEP-MEASURED (DIAG-INTERMEDIATE)). *Measured*: under the exact off-critical injection family the wall dies 100% through off-diagonal coherence — at the kill point the diagonal testing profile is statistically frozen ( $\max T_m = 0.94\text{--}0.99$ ; the coherence share equals the entire baseline diagonal margin); the separation ratio is 2.0–2.7 on the head rungs and  $> 10$  deep (no diagonal crossing up to  $10\times$  the wall’s kill amplitude). *The typed reading*: the  $O(1)$  diagonal margin does not track the exponentially collapsing wall margin —  $T_h \leq 1$  is a genuinely weaker finite statement; the RH-scale content stays in the coherence share. *Kill*: a rung on which the diagonal crosses before the wall. No marker move; no RH statement.

**The zone split (PRIME.CASE.ZONESPLIT.01, frozen).** *Registered and frozen* 2026-08-09 ([verification/v889\_route\_decisions.py], 12/12 under SPEC v2, ZONESPLIT-MEASURED (ZONE-FROZEN( $\theta^* = 0.700$ ))). *Measured*: outside the critical zone  $a > h^{2\theta^*}$  the testing margin  $1 - T$  is uniformly positive on all 42 rungs (deep-half infimum  $+0.0214 \geq$  the 0.02 bar; cut-stable;  $\theta = 2/3$  fails at  $+0.0191$ ); inside, the margin law is  $\sim 0.87/\log^2 h$  — the bar the arithmetic input must beat. *Typed honestly*: OUTLIERS-DANGEROUS (on kz 13 the margin minimizer coincides with a top-3 oscillation node — the density defects are not provably measure-zero for the margin). *Control*: scramble drives the outside margin to  $-1.4 \times 10^4$ . *Kill*: a rung with negative outside margin at the frozen  $\theta^*$ . No marker move; no RH statement.

**The PNT decider (PRIME.CASE.PNTGAMMA.01, decided).** *Registered and decided* 2026-08-09 ([verification/v889\_route\_decisions.py], 10/10 under SPEC v2, PNTGAMMA-MEASURED (PNT-INSUFFICIENT)). *Measured*: the four matched worlds (mass  $\times$  mask in {actual, PNT-smooth}, identical pipeline, exact tent-integral construction self-tested at  $6 \times 10^{-12}$ ) show the truth’s port correction  $\gamma > 0$  at every alias on every rung, *but* the PNT-smooth reference itself has  $\min \gamma < 0$  on every rung ( $-0.96$  to  $-7.8 \times 10^{-3}$ ) — the smooth reference violates the very inequality the diagonal theorem needs; the dominant rescue piece is the mask–mass *interaction* (positive on 4/4 deep rungs); the measured arithmetic deviation  $\varepsilon_A = 0.23\text{--}6.2$  is  $6\text{--}970\times$  larger than every gap. *The decision*: the diagonal theorem’s arithmetic input is genuinely stronger than PNT on this family — no classical zero-free-region-strength input closes it by smooth comparison. *Control*: scramble flips 6/6 port gammas. No marker move; no RH statement.

**The signed-homotopy classification (PRIME.CASE.SIGNEDHOMOTOPY.01, decided).** *Registered and decided* 2026-08-09 ([verification/v889\_route\_decisions.py], 12/12 under SPEC v2, HOMOTOPY-MEASURED (HOMOTOPY-INDEFINITE-ARITHMETIC); ward metrology only — run 1 ended honestly PIPELINE-BROKEN on an over-sharp deep ward). *Measured*: the signed homotopy identity along the exact PNT-to-truth interpolation is exact calculus ( $4.3 \times 10^{-14}$  with exact rational mask-crossing times); 66.3% of the 14047 samples have  $J < 0$  (pointwise positivity dead); the fixed-chamber quadratic kernel is concavity-locked negative-semidefinite on both frozen rungs; the polynomial response drains the gain ( $-17$  to  $-231\%$ , worsening with  $h$ ). *The classification*: the truth endpoint wins only through the signed *first-order* alignment of the arithmetic residual with the extremal-polynomial weights — the diagonal route’s required input is *pair-correlation class*, not any PSD kernel positivity; PRIME.CASE.SUMRULE.01 is retyped from unconditional candidate to conditional reduction theorem (the dated update above). No marker move; no RH statement.

**Flag positivity (Sylvester) and the pole-free  $s$ -corridor (PRIME.PORT.HIROTA.01 [O]).** *Registered* 2026-08-09, OPEN — the finite architecture measured ([verification/v890\_total\_positivity\_flow.py], 10/10 + 7/7, HIROTA-MEASURED (MINORS-POSITIVE) + SFLOW-MEASURED (CORRIDOR-CLEAR)). *Measured*: every nested leading principal minor  $\tau_h^{(k)}$  of  $I - C_h$  ( $k = 1, \dots, 12$ ) is strictly positive on all 37 full-window rungs — the section family is *flag-positive* (Sylvester: all leading principal minors positive; classical total positivity of all minors is not claimed), a strictly stronger finite statement than the wall scalar; the Desnanot–Jacobi identity is warded exact; the  $s$ -corridor  $s \in [0, 1]$  is pole-free on every rung (the controls pull their pole inside), and by Cauchy interlacing  $s_h^* = 1/\lambda_{\max}(C_h)$ , so  $s_h^* - 1 = \tau_h/(1 - \tau_h)$  — the wall margin *is* the pole distance of the integrable  $s$ -flow (ladder slope  $-2.90$ ). *Typed*

*honestly*: no naive rung-to-rung bilinear closure in raw  $h$  (median residual 0.655); the Toda form holds pointwise (0.001) but with drifting coefficients — TODA-VARYING. *The demand*: the multi-factor inheritance question as a total-positivity propagation statement, derived from the IKS generators. *Kill*: a full-window rung with a non-positive section minor. *Skeleton update* (2026-08-09, round CXI): the propagation question is answered in *bounded-margin* form by [verification/v892\_closure\_architecture.py] and [verification/v893\_relative\_margins.py] (the hermitian congruence margin  $\eta$  and the flag chain; the soft flag is the  $\tau$  scale after the kz-21 coordinate artifact) — the demand narrows to the  $O(1)$  ratio bound (PRIME.PORT.ETASOURCE.01). *Dictionary update* (2026-08-24, round 226): the TODA-VARYING finding is resolved into an exact dictionary — the bilinear Toda form holds with the source-pure coefficient  $\hat{\gamma}_n/\gamma_n$  from a scaled *signed* Stieltjes recursion, and its positivity is WALL-EQUIVALENT (quasi-definiteness of the signed defect measure  $\mu - \nu$ ; [verification/v955\_tau\_iiks\_toda\_dictionary.py], PRIME.PORT.TAU.FINITE.IKS.01 [E]); no source-pure positive cone was found (HIROTA-CONE-GO not reached — measured and closed via the Cauchy–Schwarz route). No marker move; no RH statement.

**The multi-factor determinant anatomy (PRIME.PORT.TAU.MOEBIUS.01, decided).** *Registered and decided* 2026-08-09 ([verification/v890\_total\_positivity\_flow.py], 14/14 under SPEC v2, TAUMOEBIUS-MEASURED (BRIDGE-DIFFUSE)). *Measured (the reviewer’s key run)*: the exact factorization  $Q_h = \det(W_{h+1})/\det(W_h) = \prod_i (1 + \mu_i)$  is warded at  $9.2 \times 10^{-12}$ ; the truth keeps every factor off  $-1$  on all 31 full-window steps over 24 orders of magnitude, while the frame-surviving smooth-mass world crosses factors on 22 of its 28 steps (first at  $h\ 210 \rightarrow 218$ ) and flips  $Q$  thirteen times — the arithmetic detector fires. *The decision*: the one-scalar Möbius bridge is closed — BRIDGE-DIFFUSE (bulk factor margin 0.005 vs the frozen 0.5 bar; the fitted Möbius datum  $|\alpha| \sim 10^{-3}$  is determinant-inert); positivity inheritance, if it exists, is a multi-factor/total-positivity statement, not a scalar cocycle inequality. No marker move; no RH statement.

**The Möbius kill battery (PRIME.PORT.MOEBIUS.KILLS.01, decided).** *Registered and decided* 2026-08-09 ([verification/v891\_moebius\_firewall\_kills.py], 10/10 + 13/13 + 7/7, CRFIREWALL-MEASURED (CR-DEAD) + GAUGEFW-MEASURED (GAUGE-MADE) + MOEBIUSCENSUS-MEASURED (UNDECIDED) — the verdicts are kills and are recorded as such). *Measured*: the earlier carrier-invariance finding was a normalization artifact — fit-free gauge-free cross-ratios over 8126 well-conditioned quadruples: pooled median defect 0.128 (bars  $10^{-3}/10^{-2}$ ), anchor triples disagree, the steps do not compose; none of the three source-frozen gauges is raw-invariant (two contrast-degenerate — invariance vacuous; the determinant-phase frame shows real motion); the scalar cocycle does not track the wall (corr = 0.13); the Jacobi-transfer identification is dead ( $d_P = 0.997$ ); and the six-world census shows the arithmetic separating *at the frame* — every control world including the smooth-mass ones loses window subcriticality (truth 41/41 alive, all controls 0), and scramble beats truth on the bare Möbius bars (0.0060/0.0011 vs 0.0218/0.0446): rank-2 window kinematics, not arithmetic dynamics. *The surviving remnant (the named object)*: the deep-core 8-node cross-ratio coherence ( $4.3 \times 10^{-4}$  on truth vs  $2.7 \times 10^{-2}$  smooth — prime-mass-driven); related sandbox kills recorded (the static Cauchy identification dead; dev/ $\tau$  grows at slope +1.46 — the Cauchy route can never certify the wall). No marker move; no RH statement.

**The rounds-50–53 campaign (the closure architecture of the wall, the bounded-margin skeleton, the diagonal refinement, the collectivity, and the Wick functor arc).** *Recorded* 2026-08-09 (a campaign note, round CXI; fourteen frozen sandbox probes, every spec amendment fail-first on record; the promoted verification modules [verification/v892\_closure\_architecture.py], [verification/v893\_relative\_margins.py], [verification/v894\_diagonal\_refinement.py], [verification/v895\_collective\_comb.py] and [verification/v896\_wick\_block\_functor.py] embed their discovery probes byte-exact). The campaign gives the wall its closure architecture (the exact hermitian congruence with a non-decaying inheritance margin; the flag/pivot chain positive by pure algebra; the honest two-sided Rouché kill; the wall as one fixed  $8 \times 8$  Schur core), sources both bounded margins and prints the theorem skeleton whose *one* unconditionally open statement is the  $\tau$ -sign inheritance in bounded-margin form, makes the conditional diagonal contract concrete (the exact band-limited

prime kernel, frozen in-probe), measures the collectivity of the arithmetic (the one-sided law; the deep core as the entire multiplicative von Mangoldt comb), and closes the Wick functor arc of the finite compiler at the structure level (the proved scalar obstruction; the block resolution; the seam-diagonal demand). The registrations follow as the blocks below (plus the dated skeleton update on PRIME.PORT.HIROTA.01 above). Full narrative in the prime-front paper and Paper 1; audit cards in Paper 3. No marker moves; no RH statement.

**The hermitian congruence inheritance form (PRIME.PORT.RELCONG.01).** *Registered* 2026-08-09, measured ([verification/v892\_closure\_architecture.py], 16/16, RELCONG-MEASURED). *Measured:*  $A_{h+1} = A_h^{1/2}(I + H_h)A_h^{1/2}$  exact ( $6.5 \times 10^{-14}$ ;  $H_h = A_h^{-1/2}\Delta_h A_h^{-1/2}$  exactly hermitian); the margin  $\eta_h = 1 + \lambda_{\min}(H_h)$  is *non-decaying* (min 0.0050, median 0.29, power slope +0.108 while  $\tau$  collapses at  $-2.74$ ) — the inheritance inequality is bounded-margin-shaped, not collapsing;  $\rho(W^{-1}\Delta C)$  was a misleading statistic (strongly non-normal:  $\|M\|/\rho$  up to 26,  $\kappa(V)$  up to 134); the smooth-mass world is CONE-EXITED on all 28 pairs. *Mini-theorem (verified):* base PD +  $\eta_h > 0$  for all  $h \Rightarrow$  the whole ladder (Sylvester congruence). *Source:* the  $\eta$  scale is explained by PRIME.PORT.ETASOURCE.01; the  $O(1)$  ratio is the open piece. *Kill:* a truth step with  $\eta \leq 0$ . No marker move; no RH statement.

**The flag-chain induction (PRIME.PORT.PIVOTFACTOR.01).** *Registered* 2026-08-09, measured ([verification/v892\_closure\_architecture.py], 15/15 under SPEC v2, PIVOTFACTOR-MEASURED). *Measured:* the 12 LDL<sup>T</sup> pivots of  $A_h = I - C_h$  factor rung-to-rung as  $d_{k,h+1} = d_{k,h}(1 + \mu_{k,h})$  with *every* factor positive — all 12 flag quotients  $Q_k > 0$  on all 31 truth steps (pivot = minor quotient warded  $8.3 \times 10^{-14}$ ); the ladder minimum sits at the soft flag  $k = 12$  at  $2.06 \times 10^{-14}$ , sign-decidable at  $6.4 \times 10^9 \times$  the eigenvalue noise floor (Cauchy-interlacing guard); the smooth world violates at the *bulk* flags first ( $k = 1$  on 10 steps,  $k = 3$  on 9; zero steps violate the soft flag alone); per-prime non-decompositional (leverage to 987 $\times$ ). *The shape:* the shortest induction of the program — base (v884/v887) +  $\min_k Q_{k,h} > 0 \Rightarrow$  the whole ladder by pure algebra; the soft-flag scale is  $\tau$  itself (PRIME.PORT.RELFLAG.01). *Kill:* a truth step with a non-positive flag quotient. No marker move; no RH statement.

**The Rouché route (PRIME.PORT.EULER.ROUCHE.01, dead).** *Registered and closed* 2026-08-09 ([verification/v892\_closure\_architecture.py], 14/14 under SPEC v2, ROUCHE-MEASURED — an honest kill with the mechanism documented). *Measured:* on complex contours enclosing the corridor  $s \in [0, 1]$ ,  $\sup_{\Gamma} \|A_h(s)^{-1}\Delta_h(s)\| \geq 1$  on 22/31 truth steps (max 63.4) *including* off-cap — not merely contour-limited: a two-sided contour norm bound is the wrong theorem shape for inheritance on this ladder (consistent with the one-sided avoidance law of PRIME.PORT.FACTORAVOID.01); the argument-principle machinery itself verifies exactly (integer zero counts at  $5 \times 10^{-15}$ ); the smooth world is corridor-broken on 28/28 but the mechanism is non-discriminating since truth fails the bound too. No marker move; no RH statement.

**The  $8 \times 8$  Schur-core reduction (PRIME.PORT.DEEPCORE.SCHUR.01).** *Registered* 2026-08-09, measured ([verification/v892\_closure\_architecture.py], 16/16 under SPEC v2, DEEPCORESCHUR-MEASURED). *Measured:* on the full wall operator, the block split at the fixed deep-core aliases  $\{2, \dots, 16\}$  gives Haynsworth-exact inertia and the wall *is* the fixed  $8 \times 8$  Schur core:  $\lambda_{\min}(S_h) w_{\text{core}}/\tau_h = 1 \pm 8.4 \times 10^{-8}$  with core mass  $w_{\text{core}} \geq 0.999937$  on all 41 full-core rungs — the RH-critical object is a fixed  $8 \times 8$  family, not a growing operator; the absolute exterior margin shrinks ( $\lambda_{\min}(R) \sim h^{-2.865}$ ) but  $\lambda_{\min}(R)/\tau$  is trendless at 210–2200 — the reduction is conditional on a  $\tau$ -relative exterior bound only ( $c_R = 210$  measured, PRIME.PORT.RELFLAG.01); core inheritance holds (0 crossings,  $\eta_{\text{core}}$  min 0.0315); the per-prime grain is dead at fixed dimension too. *Kill:* a full-core rung breaking the seven-digit identity. No marker move; no RH statement.

**The  $\eta$ -margin source (PRIME.PORT.ETASOURCE.01).** *Registered* 2026-08-09, explained ([verification/v893\_relative\_margins.py], 20/20 under SPEC v2, ETASOURCE-MEASURED (SCALE-MATCHED)). *Measured:* in the dangerous direction the increment  $d_h = |v^T \Delta v|$  and the operator  $a_h = v^T A v$  collapse *together* (slopes  $-3.716$  vs  $-3.456$ , residual correlation 0.974, 94% shared variance): the non-decay of the inheritance margin  $\eta$  is *explained* —  $\eta$  is scale-free because numerator and denominator are sections of one arithmetic object. *Open:* only the  $O(1)$  scatter of the ratio

(the fitted  $\tau$ -laws predict none of it); seat anatomy: the minimizer sits in the core while the binding restricted block is the 4-alias exterior (0.0056 vs core 0.0319). The round-53 sandbox continues with the mechanism decomposition. No marker move; no RH statement.

**The relative margin laws and the theorem skeleton (PRIME.PORT.RELFLAG.01).** *Registered* 2026-08-09, measured ([verification/v893\_relative\_margins.py], 18/18 under SPEC v2, RELFLAG-MEASURED). *Measured:*  $d_{12} = 1/(A^{-1})_{12,12}$  exact with the factorization  $r_{12} = (\tau'/\tau)(c'/c)$  at  $2.1 \times 10^{-16}$ ; the  $c = d_{12}/\tau$  ladder is trendless with one outlier (kz 21) diagnosed in the round-53 sandbox as a *coordinate artifact* (the kz-19 mechanism repeats; corrected spread  $124 \rightarrow 15.0$  — the soft flag is  $\tau$ -tied);  $\lambda_{\min}(R)/\tau$  trendless with ladder min  $c_R = 210$ ; COMPRESSIONS-DISTINCT (honest: the port bulk margin is an *absolute* floor while the core exterior is  $\tau$ -tied — different statements at two compressions). *The printed skeleton:* (i) base certified (v884/v887; interval-rigorous since v897); (ii)  $\lambda_{\min}(R) \geq 210\tau$  trendless; (iii) core  $\leftrightarrow \tau$  at seven digits; (iv)  $\eta > 0$  non-decaying — the *one* unconditionally open statement is the  $\tau$ -sign inheritance, reorganized from collapsing-margin form into *bounded-margin* form ( $c, c'/c, c_R, \eta$  — all  $O(1)$ ; round-55 sharpening:  $c'/c$  sensitivity range  $[0.223, 5.258]$  and the exact Christoffel identification of  $c$ , [verification/v899\_christoffel\_normsquare.py]). No marker move; no RH statement.

**The extracted conditional contract (PRIME.CASE.PAIRCORR.CONTRACT.01 [O]).** *Registered* 2026-08-09, OPEN — the conditional reduction extracted with every object machine-defined ([verification/v894\_diagonal\_refinement.py], 11/11 + 13/13 + 12/12, PAIRCORR-EXTRACTED (KERNEL-BANDLIMITED) + ENDPOINT-MEASURED (ENDPOINT-PARTIAL) + INNERLINEAR-MEASURED (INNER-EMPTY)); the route stays conditional per v889). *Measured:* the first-order homotopy gain regroups exactly ( $4.5 \times 10^{-15}$ ) into the prime-side form with the explicit node kernel  $K_{h,m}$ ; the kernel is band-limited (compact spectral support — the contract hypothesis is one-sided band-limited explicit-formula positivity at the port frequencies, pair-correlation class); the full LaTeX contract statement is frozen in-probe; all 79 (rung, alias) finite margins positive; the endpoint refinement: the global linear collapse fails (142/227) but deep small- $a$  cells pass monotonically; the inner-zone accounting localizes the irreducibly-hard part to a *non-shrinking* band ( $\sim 20\%$  of aliases, 38–56% of the mass). *Demand:* state-and-prove the band's weighted pair inequality, or accept conditionality. *Norm-square note (2026-08-09, round CXV):* the kernel's frequency weights are  $W_{f,m} \geq 0$  *exactly* and the single top-edge tent defect is an extension-convention artifact killed by the periodic full-weight fold (negative index 0; 79/79 modified margins positive; the boundary term  $\beta$  carried closed-form, max  $\beta/\text{margin}$  0.21) — the contract is now a *norm square* in the frequency weights ([verification/v899\_christoffel\_normsquare.py], PRIME.CASE.KERNEL.SOS.01 + PRIME.CASE.EDGEDEFECT.01 below); the arithmetic hypothesis itself stays conditional. *Kill:* a negative finite margin on the frozen census. No marker move; no RH statement.

**The one-sided collective law (PRIME.PORT.FACTORAVOID.01, decided).** *Registered and decided* 2026-08-09 ([verification/v895\_collective\_comb.py], 18/18 under SPEC v2/v3, FACTORAVOID-MEASURED (AVOIDANCE-STRUCTURAL)). *Measured:* the norm-bound theorem shape does not exist —  $\rho(W^{-1}\Delta C) \geq 1$  on 17/31 truth steps (max 57.3) but always via the harmless  $\mu > +1$  side; the true law is *one-sided*:  $\max(-\min \mu) = 0.9950 < 1$  on all 31 steps, which *is* PD inheritance itself; the geometry/atom split is exactly degenerate (the atom cutoff is slaved to the tent-support edge; every leaving step is atom-nil bitwise) — no separable atom channel: the increment is the collective window re-test of the full comb; per-prime leave-one-out is non-decompositional (a leverage census, not a decomposition). No marker move; no RH statement.

**The deep core is the port core (PRIME.PORT.DEEPCORE.01).** *Registered* 2026-08-09, measured ([verification/v895\_collective\_comb.py], 12/12, DEEPCORE-MEASURED). *Measured:* the 8 deepest common nodes of the carrier ladder are the fixed even aliases  $\{2, \dots, 16\}$  = the Bessel coordinates  $a_m = \pi^2 m^2$  (0.1% at depth; census 39/41 steps; coherence knee  $k^* = 10$ ); the certificate-level cross-ratio coherence (truth  $4.3 \times 10^{-4}$  vs smooth  $2.7 \times 10^{-2}$ ) is killed by *all four* surgical modifications — edge-only smoothing, interior-only smoothing, prime-power echoes removed, and the wrong- $\Lambda$  world: **the signal is the entire multiplicative von Mangoldt comb** — positions, powers, and exact  $\log p$  values all load-bearing; no smaller frozen object exists — the global rigidity of v883

at the finest coordinate. *Kill*: a surgery that preserves the coherence. No marker move; no RH statement.

**The Wick functor arc (SEAM.CFIN.WICKFUNCTOR.01 [O]).** *Registered* 2026-08-09, OPEN — structure exact, obstruction proved, block candidate constructed ([verification/v896\_wick\_block\_functor.py], 27/27 + 33/33 + 26/26, WICKFUNCTOR-PARTIAL + WICKMIX-MEASURED (WICKMIX-OBSTRUCTED) + WICKBLOCK-MEASURED (BLOCKFUNCTOR-CONSTRUCTED / SEAM-DIAGONAL)). *Exact*: the six compiler roles exist in the deployed seam ( $D_5$  slot pairs + the  $A_3$  boundary block;  $\mathfrak{so}(16)$  sector law,  $q^*$  edge grading,  $C_6$  intertwining all exact) but the deployed vacuum kernel is channel-diagonal; *the proved scalar obstruction*: the pinned  $C_6$  generator’s 2-cycle  $\{4, 5\}$  forces a duad zero — every  $C_6$ -covariant antisymmetric scalar covariance reaches at most 14/15 duads (a theorem of the deployed action); *the block resolution*: the  $C_6$ -covariant cross-block space has dimension 33; the canonical block covariance is exactly covariant, strictly CAR-positive, carries 15/15 blocks and 15/15 block Wick monomials with the canonical  $\chi$  sign law; *the decisive fact*: the deployed seam covariance has *zero* cross-blocks everywhere (SEAM-DIAGONAL). *Demand*: a physical channel-mixing mechanism (a boundary/transport-twisted state, not the vacuum) realizing the block functor’s value level — then the  $K_6$ /Pfaffian grammar of v888 would be the measured correlation grammar of the seam. *Candidate note (2026-08-09, round CXV)*: the demanded state now has a *constructed candidate* (SEAM.CFIN.KSMIX.01 below, [verification/v898\_kms\_schur\_mixing.py]); whether the *actual* seam realizes it is untouched. The [O] premise is unmoved. No marker move; no RH statement.

The rounds 54–56 block (2026-08-09; promoted in [verification/v897\_certified\_interval\_ladder.py], [verification/v898\_kms\_schur\_mixing.py], [verification/v899\_christoffel\_normsquare.py] and [verification/v900\_core\_update\_anatomy.py]) delivers the reviewer’s “immediate” list: the interval-rigorous certificate base, the channel-mixing candidate state, the Christoffel/norm-square anatomy of the diagonal contract, and the exact core-update coordinates with their honest negatives.

**The certified interval ladder (PRIME.PORT.BALLLADDER.01, delivered).** *Registered and delivered* 2026-08-09 ([verification/v897\_certified\_interval\_ladder.py]; probe verdicts BALLHEAD-PROVEN and BALLLADDER-COMPLETE, 10/10, 3960 s). *Proven*:  $\sigma_h > 0$  with *rigorous* mpmath.iv interval-arithmetic shifts on all 42 reachable rungs  $h = 142\text{--}878$  — 15 exact-rational integer Bareiss/Sylvester certificates ( $h \leq 300$ , floors  $4.5 \times 10^{-5}\text{--}2.1 \times 10^{-4}$ ) + 27 validated-precision Cholesky certificates (dps 120/200, floors  $5.2 \times 10^{-6}\text{--}8.4 \times 10^{-5}$ , pivot-to-bound ratios  $\sim 10^{115}$ ); the node-enclosure lemma, native outward-rounded interval transcendentals, exact midpoint/radius extraction and the rigorous shift  $2h \text{ rad}_{\max}$ ; **the informal  $\varepsilon_c$  lag-evaluation error model of v884/v887 is retired from the ladder**; the Epstein control refused at pivot index 10 by the identical interval machinery. Suite economy declared: the promoted module re-certifies all 15 tier-1 + 3 tier-2 rungs verbatim and freezes the full census with both probe SHAs. *Remaining named step*: the Lean composition through krein\_cofinal\_weil; the asymptotic tail stays with PRIME.PORT.TAIL.01 / PRIME.PORT.LEADING.SIGN.01. *Kill*: a refused pivot on any rung under the interval shifts. No marker move; no RH statement.

**The Christoffel pivot identity and the norm-square contract (PRIME.PORT.CHRISTOFFEL.RATIO.01 + PRIME.CASE.KERNEL.SOS.01 + PRIME.CASE.EDGEDEFECT.01).** *Registered and measured* 2026-08-09 ([verification/v899\_christoffel\_normsquare.py], 19/19 + 15/15 + 19/19, CHRISTRATIO-MEASURED (WINNER-CHRISTOFFEL / EXTSOURCE-DECOUPLED) + KERNELSOS-MEASURED (INDEX-LARGE; RMIN=1; NOT-ABSORBED) + EDGEDEFECT-MEASURED (NORMSQUARE-ACHIEVED; KILL=PERIODIC; CHAIN=CARRIED)). *Exact*:  $1/d_{12} = 1 + v_* K_\sigma(y_*)$  — the soft pivot is a deflated-Christoffel evaluation (independent route at  $2.0 \times 10^{-11}$ ), winning the boundedness exposure by two orders of magnitude; *the honesty fence, warded exactly*:  $\lambda_{\min}(I - G) = \tau$  — the positivity premise is the wall’s own PD premise in Christoffel coordinates, a coordinate change, not an independent source;  $c'/c$  sensitivity range  $[0.223, 5.258]$ ,  $r_{12} > 0$  on 31/31, exterior reserve decoupled (corr 0.454). *The norm-square form*: the pair kernel’s frequency weights are  $W_{f,m} \geq 0$  exactly; the entire negative content is one top-edge tent (an extension-convention artifact), killed exactly by the periodic full-weight fold (negative index 0; 79/79 modified margins positive; the boundary term

carried closed-form, max 0.21) — the conditional contract of PRIME.CASE.PAIRCORR.CONTRACT.01 is a norm square in the frequency weights; **the arithmetic hypothesis itself stays conditional** (pair-correlation class per v889). *Kill*: a negative modified margin on the frozen census. No marker move; no RH statement.

**The normalized core update anatomy** (PRIME.PORT.NORMALIZED.CORE.01 + PRIME.PORT.GRAPH.REGION.01). *Registered and measured* 2026-08-09 ([verification/v900\_core\_update\_anatomy.py], 20/20 + 23/23, NORMCORE-MEASURED (FAMILY-BOUNDED / UPDATE-EXACT / REGION-OPEN) + GRAPHREGION-MEASURED (COUPLING-FREE / GRAPH-REGION-OPEN / SHADOW-LAWLESS)). *Exact*: the normalized deep core  $X_h = S_h/\tau_h$  obeys  $X' = c(X + U)$  at  $8.3 \times 10^{-16}$  with  $c = \tau/\tau'$ ; the family is bounded (diameter 2533) and the input effectively two-dimensional ( $c$  + one matrix mode at 97.5% energy). *The registered honest negatives* (boundary conditions for any invariant-region theorem): no linear coupling law  $u \sim F(X)$  (both  $R^2$  below the 200-shuffle null  $q_{95}$ ); the rank-1 truncation fails — the positivity is carried by the 2.5%-energy off-mode; the naive product and graph-tube invariant sets are refused. *The registered direction*: keep the exact  $X + U$  with its sign-dependent off-mode structure — no mode may be truncated. *Kill*: an update defect above the ward, or a coupling law beating the null. No marker move; no RH statement.

**The two dead readings** (SEAM.CFIN.MIXING.NORMALIZATION.01 + SEAM.CFIN.MINIMAL.MEDIATOR.01) — **registered negatives**. *Registered as first-class negatives* 2026-08-10, rounds 57–58 ([verification/v903\_seam\_rp\_exclusion.py], 23/23 + 20/20 on these two claims, SEAMNORM-MEASURED + MEDIATOR-MEASURED). *Dead reading 1*: ‘ $t = 1/8 = 2\pi c_3/|Z_2|$  is a compiler value’ is a *decoration* — the three factors are individually deployed but the combination is free: all 55 scanned  $t$  pass all frozen v898 gates (the allowed set is mathematically  $(0, \infty)$ ;  $u \in \{0, \dots, 2\}$  and  $\beta \in \{1/2, \dots, 4\}$  pass too), no frozen functional distinguishes  $1/8$  (PIN-ABSENT), and the anti-numerology census finds four equally cheap deployed-constant readings; the kept structure fact: the channel mixing is *thermal* (it dies to the ground state, 0/15 cross-blocks at  $\beta = 100$ ). *Dead reading 2*: ‘ $N_{\text{fam}} = 3$  is the minimal boundary-mediation rank’ is *refuted* at theorem grade for this finite size — one boundary pair suffices ( $\det B_x = \gamma_{ab}\gamma_{cde}$  symbolically for every symplectic-rank-1 mediator; the explicit integer witness passes the full requirement 10/10; seeded census 2000/2000), the deployed mediator has rank 2 *exactly* ( $S = 1_5 \otimes 3J$ ), and what 3 really pins is boundary *purity* = dimension bookkeeping ( $\dim(A_3) = 6 = |Z_2|N_{\text{fam}}$ ), not a mediation theorem. Negative results carry their own ledger rows, not footnotes. *Kill* (for the deadness itself): a frozen v898-gate functional that pins  $t = 1/8$ , or a physical demand that excludes the rank-1 witness. No marker move; no RH statement.

The rounds 60–63 block (2026-08-10/11; promoted in [verification/v905\_bffloor\_ideal\_certificate.py], [verification/v906\_tail\_cartography.py], [verification/v907\_halfgap\_registered\_target.py] and [verification/v908\_seam\_equilibrium\_wiring.py]) closes the round-60 seam questions on the compiler side and delivers, on the RH side, the certified  $B$ -half surface floor, the tail mechanism map and the registered half-gap target with its first blind holdout (the three RH-side items live in the prime-front document; the seam wiring is registered here). The rounds 64–71 block (2026-08-11/12; promoted in [verification/v909\_finite\_wall\_closure.py], [verification/v910\_finite\_zero\_transfer.py] and [verification/v911\_wiring\_freedom.py]) closes the *finite* composed wall census and its transfer law on the prime front (PRIME.WALL.FINITE\_CLOSURE.01, PRIME.WALL.FINITE\_ZERO\_TRANSFER.01; both in the prime-front document) and, on the seam side, answers the wiring selector below.

**The Reed–Muller identification** (PRIME.PACKET.RM14.01). *Registered* 2026-08-06 ([verification/v819\_prime\_packet\_rm14.py], 21/21 incl. the three must-fire controls, RM14-EXACT; Lean companion TfftCarrier/PackageRM14.lean, 30 theorems). *Certified*: the 15-label channel row hull is the punctured  $\text{RM}(1, 4)^* = [15, 5, 7]$  (enumerator  $1 + 15z^7 + 15z^8 + z^{15}$ );  $C_0 = [15, 4, 8]$  simplex;  $\chi_{\text{NSR}}$  one codeword (7+8 its two sides,  $C_2$  the affine bit); the self-reproduction cycle closes with exactly 1344 equivalences; the CSS code  $[[15, 1, 3]]$  with dual  $[15, 10, 4]$ , logical distances (7, 3), triorthogonality  $8/4/2$ . *Typed  $[H]$* : the  $\zeta_8$ /T-gate phase pointer (v798); the TFPT fingerprints ( $14 = \dim G_2$ ,  $3 = N_{\text{fam}}$ ,  $4 = |\mu_4|$ ), audit-only. *Kill*: any census identity failing on the frozen frame.

No marker move; no RH statement.

**The canonical check rule and the plane frames (PRIME.KRAUS.RM24.01).** *Registered* 2026-08-06, carrying PRIME.C2.PLANEFRAMES.01 ([verification/v820\_prime\_kraus\_rm24.py], 16/16 + 14/14, KRAUS-RM24-RULE-CANONICAL + PLANEFRAMES-EXACT; one honest coset-census fix carried). *The honest double kill:* the literal leg-support claim is dead twice (weights  $1/2/3 < \text{distance } 4$ ;  $Sp(4, 2)$  orbits  $15+90$  vs  $45+60$  — no equivariant bijection exists). *Certified:* the canonical  $\omega$ -built rule identifies the rule-A protocol set with the 105 weight-4 checks of  $[15, 10, 4]$  exactly,  $Sp(4, 2)/\sigma/\text{deck}/\text{KMS-covariant}$ ; the six Gram identities entrywise;  $A_3 = 12 / B_3 = 16$  as per-point plane completion counts ( $\sigma_3(3) = 28$ ,  $a_3 = -4$ , independent eta-product rebuild); the graded frame intertwiner  $V^*V = \sigma_3(n)I$ ,  $V^*\Gamma V = a_n I$  exact for  $n = 3, 5, 7$ . *Kill:* an equivariance mismatch of the  $\omega$  rule under the frozen generators. No marker move; no RH statement.

**The vacuum completion (PRIME.VACUUM35.01).** *Registered* 2026-08-06 ([verification/v821\_prime\_vacuum35.py], 22/22 incl. the three must-fire controls, VACUUM35-EXACT). *Certified:*  $140 = 105 + 35$ ; the Grams  $22I+6J / 6I+J / 28I+7J$  entrywise with spectra  $\{112, 22^{14}\} / \{21, 6^{14}\} / \{133, 28^{14}\}$ ; the  $E_7$  completion  $112 + 21 = 133$  as a matrix identity (the  $E_7 \times A_1$  corpus reading a [C] fingerprint, anchor v170); Weitzenböck  $H_{105}^T H_{105} = 2B^2 + 14I$  as the canonical label-side left factor (v758's continuum wall explicitly untouched); the  $\pm B$  cancellation inside one Gram; the tight frame constant 28 on mean-zero; RM(2, 4) dual-containing. *Typed [H]:* the vacuum/continuum reading — decided negatively by the closure row below. No marker move; no RH statement.

**The vacuum-route closure (PRIME.CONTINUUM.UNSHORTEN.01, carrying PRIME.VACUUM.DILATION.01).** *CLOSED* — *honest double kill* (registered and decided 2026-08-06; [verification/v822\_prime\_vacuum\_dilation.py], 17/19 + 12/14 with the four FAILs being the preregistered PSD deciders themselves, controls  $3/3 + 3/3$ , UNSHORTEN-DEAD + DILATION-DEAD; first runs after freeze, no repairs). *The dead reading:* the archimedean/pole continuum as the RM(2, 4) vacuum completion of the 15-label window objects — dead in *both* transcriptions: the one-slot bordering by the v758 mechanism verbatim and on scale (with the shape theorem: one slot induces only a rank-one pure- $J$  coupling; the demanded  $6I+J$  needs 35 separate rows), and the corrected target itself — the 35-row Stinespring dilation, shape fixed exactly,  $\times 13.1$  distribution gain — still failing both sectors under both frozen normalizations. *The frozen fence (the round's quantitative product):* any continuum re-entry through the RM check geometry is PSD iff its coupling satisfies  $c \leq c^* = (2/21) \lambda_{\text{pencil}}(T_{\text{dep}}, T_{\text{cont}}^2)$ ; measured deficits  $\times 3 \times 10^7$  (GL1) /  $\times 2 \times 10^2$  (f8) against the mass calibration — the reading dies *with the calibration*, not the geometry: the continuum's completion role is already exhausted inside  $T_{\text{dep}}$  (the explicit-formula balance), and the path-state construction is not buildable on this gating. The RM integer layer stays certified. No marker move; no RH statement.

**The PD-persistence universal form (PRIME.PD.PERSISTENCE.01).** *OPEN* (registered 2026-08-04): the universal-theorem form of the PD-persistence wall ([verification/v727\_l1\_identification.py], [verification/v734\_s1\_canonical.py]) with the important correction  $\lambda_{\min} \geq 0$  instead of  $\geq \delta$  — the Ihara lesson (FTR.IHARA.01): true RH analogues can have *exact zero modes after rank saturation*, so a uniform margin is stronger than necessary and the wrong instrument. *Target form:* cofinally infinitely many PD windows, established via a state / Gram representation / cone-preserving recursion — *not* via eigenvalue bounds. *The review blacklist, confirmed as permanent kills:* the raw symbol minorant, norm triangles, perturbation theory, position-blind estimates, the direct Sonin transfer ([verification/v732\_strat3\_sonin\_prolate.py]) and naive reflection positivity ([verification/v729\_strat2\_rp\_universality.py]) stay retired; further zero-matching records do *not* change the quantifier. A research contract, *not* a claim; no marker move; no RH statement. [2026-08-04 update (the incidence programme): the contract now carries a REGISTERED INDUCTION FORM, a CLOSED route, and a NAMED open object. (i) The induction form: the ramified Schur ladder on the canonical tower ([verification/v755\_simpler\_schur\_recursion.py], 13/13, SCHUR-RECURSION-ALIVE) —  $S_k \succeq 0$  measured on all 8 ram-odd stages ( $n_k = 2^{2k+1}$ ,  $\lambda_{\min}$  flat, long-double-hardened), with basis, Douglas range condition, Albert/Haynsworth step and cofinality

all exact: the open proof need is exactly the structural Gram certificate  $S_k = R_k^T R_k$  per step; the measured boundary data for any candidate  $R_k$  are recorded (the label part enters inverted,  $\rho_k < 0$  with  $|\rho_7|$  within factor 1.34 of 49/4 — cross-term structure, not an added block — and continuum vs atom sides are order-1 comparable). Honestly: for  $\zeta$  data this is equivalent to tower PD — the wall in recursion shape. (ii) The closed route: the review's assembled certificate [continuum root $\oplus$ Kraus] is not mountable ([verification/v758\_simpler\_certificate.py], 8/8, CERT-CONTINUUM-ROOT-DEAD, the keystone negative): the arch+pol block is PSD on no stage (Levinson break at cell 47, exactly between the  $n = 2$  and  $n = 3$  cells, ladder-wide), the continuum Schur complement is no Haynsworth object, the additive split fails on every stage (the order-1 balance sits in the cross-term), and the rescue cascade is strictly cell-ordered — the wall IS the missing left factor. (iii) The named open object: the corrected target shape is the cell-wise cascade factorization  $R = \prod(\text{cell factors})$  — rescue order  $n = 2, 3, 4, 5, \dots$  across the channels — with the exact building blocks of the incidence programme ([verification/v752\_projective\_hamming\_incidence.py], [verification/v753\_ramified\_polarity.py], [verification/v754\_ramodd\_twostep.py], [verification/v756\_kms\_incidence\_stinespring.py]) as the label-side material. No marker move; no RH statement.]

**The diagonal Gram closure theorem (PRIME.GRAM.DIAGONAL.01).** *CLOSED* (registered 2026-08-04, the diagonal gram round; closed 2026-08-05 per its own frozen stakes — an honest negative adjudication, not a marker upgrade; the full closure cascade is the dated note at the end of this block). The target form, kept verbatim for the record: construct, for the countable dense test family  $\{\phi_i\}$  of the canonical  $(D, X)$  tower ([verification/v762\_dense\_weil\_core.py], 26/26, DENSE-CORE-CANONICAL: frozen enumeration, explicit stage map, admissibility located per class), PSD source matrices  $G_n$  without any target decomposition, zero data, or positivity test of the target, satisfying

$$\max_{i,j \leq n} |(G_n)_{ij} - W(\phi_i * \tilde{\phi}_j)| \leq 2^{-n};$$

then every fixed  $m$ -corner of  $W$  is a PSD limit of PSD matrices — Weil positivity on the rational span of the family, and the density step closes by the classical approximation argument. *The round's measured decision state:* Gate 1 positive — the handoff error is boundary/cutoff, not quadrature ([verification/v759\_handoff\_fixed\_window\_resolution.py], RESOLUTION-BOUNDARY); Gate 2 positive — the anti-alias identities are exact with sharp count  $N_{f,\min} = 2M-1$  ([verification/v760\_antialias\_exact.py], ANTIALIAS-CORRECTED) and the paired atom+pole remainder carries the explicit unconditional vanishing bound  $\eta_K$  with no RH-strength input ([verification/v761\_atom\_pole\_abel.py], combined ABEL-PAIRED-BOUND); Gate 4 positive — the family is canonical ([verification/v762\_dense\_weil\_core.py]); Gate 3 negative on the canonical class — the battery near-kernel pairing  $q_f$  is non-Cauchy at honest maximal depth and survives orthogonalization by structural subordination ([verification/v763\_yosida\_handoff.py], QF-DEAD): the  $q_f$  wall — density and near-kernel avoidance collide on the same function class — is the named wall object; Gate 5 (the assembled diagonal certificate, the planned v764) was not executed — its precondition (theorem-capable bounds from Gates 2 and 3) failed via Gate 3; the red team is green ([verification/v765\_handoff\_redteam.py], REDTEAM-GREEN-COMBINED, with the binding note that certificates must always carry both gate families — symbol/PSD and rate/ratio/identification). *Kill criteria that did NOT fire:* no RH-strength input was needed anywhere (the v761 circularity flag and the v763 KILL audit are machine-checked: no  $1/\varepsilon$  bound, no PD-margin assumption); no spectral gap was assumed (exact zero modes stay admissible — the Ihara lesson). *The binding stop-list* (ten items, frozen): no atom/prime-layer positivity induction; no separate absolute atom/pole bounds; no 16-dimensional channel square sums; no exact neighboring-window embeddings; no unbounded  $A^{-1}$  observables; no Ward-drift signs; no rate fits as proof substitutes; no zero-matching records as progress metric; no positive continuum root as certificate factor; no cellwise PSD demands. A research contract, not a claim; no marker move; no RH statement. [2026-08-05 closure note — the route is CLOSED per its own frozen stakes. The final decision cascade, full tally: Gates 1/2/4 positive from the diagonal gram round (above); Gate 3 DEAD AT THE CORNER LEVEL — the audit pair ([verification/v769\_qf\_contract\_necessity.py],

combined QF-GAUGE/GATE3PRIME-DEAD) first proves the full-ladder Cauchy demand on  $q_f$  was stronger than the contract needs ( $q_f$  is representation gauge: a deterministic  $\varepsilon$ -net subsequence of density 0.144 makes all 14 profiles Cauchy at the parent bars, and the contract's own  $E_Y$  corner objects are blind to  $q_f$  jumps, corr +0.038, median  $\rho = 1.13$ ), and then the corner-native weakened Gate 3' dies on its own predeclared  $1.6 \times 10^7$  comb ( $X = 16.5625$ ): the corner increments fall to  $X \sim 12.9$  and then rise sustained (+0.10..+0.34 per  $X$  unit at  $\varepsilon = 10^{-3}$ , 0/56 cells) — the v764 precondition is refuted; Gate 5 / v764 therefore was never executed. The  $q_f$  offensive's own line: gauge  $\rightarrow$  corner rise ([verification/v769\_qf\_contract\_necessity.py])  $\rightarrow$  settled-positive levels ([verification/v770\_qf\_spectral\_bundle.py], QF-SETTLES-POSITIVE: the band weights plateau at positive per-function constants, R2 family 0.225–0.358, R1 family 0.008–0.079; zero-bar depth  $X^* \sim 227 = \text{never}$ )  $\rightarrow$  no fixed-d object ([verification/v772\_qf\_feshbach\_effective.py], FESHBACH-PARTIAL/EDGE-KEEPS-MOVING: the moving spectral edge — avoided crossing 0.0039 at  $M = 1240$  on the  $10^9$  comb, widening non-regular entry cadence  $\Delta X = 1.625/1.8125/2.625$ , spread 0.381)  $\rightarrow$  cocycle domain-only ([verification/v773\_qf\_cell\_cocycle.py], COCYCLE-DOMAIN-ONLY: exact identity and zero-breach domain transport, 0/6 convergence cells — structure without limit). What survives as the theorem-shaped facts, each typed in its own module row: the exact anti-alias theorem ([verification/v760\_antialias\_exact.py]); the unconditional paired Abel bound ([verification/v761\_atom\_pole\_abel.py]); the canonical dense family ([verification/v762\_dense\_weil\_core.py]); the exact source-native domain-preserving monotone Möbius/Redheffer cell cocycle through mode entries ([verification/v773\_qf\_cell\_cocycle.py]); the Kato bundle frame ([verification/v770\_qf\_spectral\_bundle.py]); the Herglotz-certified exact Feshbach reduction ([verification/v772\_qf\_feshbach\_effective.py]). NO new variants; the ten-item stop-list stays binding; the sharpened PRIME.Z1.OPERATOR.01 takes over with the measured handover content [2026-08-05: executed as the v780 trilogy — compactness carried, selection the one shared open object, merged target Z1-COMPACTNESS]. Lean legacy (2026-08-05): the exact cell-cocycle core of the closed route is machine-checked as general Mathlib theorems over  $\mathbb{R}$  — TfptCarrier/CellCocycle.lean, 18 declarations: the exact Banachiewicz/bordered-Schur update identity, PD cell block  $\Rightarrow$  PSD corner correction  $\Rightarrow$  monotone Loewner flow with domain preservation, bounded-monotone entrywise matrix convergence (the theorem names the missing analytic ingredient precisely — a uniform Loewner upper bound, which the finite ladder did not certify and which is not claimed), and the Redheffer/linear-fractional composition laws over arbitrary commutative rings. No RH statement.]

**The structural upgrade: the canonical nested tower replaces the historical window family (PRIME.TOWERNEST.01).** The window quantifier of the PD-persistence formulation reduces exactly to the canonical nested  $(D, X)$  tower ([verification/v749\_simpler\_tower.py], 8/8, TOWERNESTED-REDUCES): in  $X$  an ascending chain —  $X$ -nesting exact to  $6.7 \times 10^{-18}$ , so limit positivity is one limit object per chain — and in  $D$  dyadic downward inheritance, exact: PD fine forces PD coarse; every deployed window is a tower member at deviation 0.0. The window quantifier was an artifact of the historical selection: box step functions are dense, so tower-PD for  $D \rightarrow 0$ ,  $X \rightarrow \infty$  is equivalent to the Weil criterion (classical density step). PRIME.PD.PERSISTENCE.01 is therefore restated with the canonical truncation in place of the historical window family. *Honest limits:* incommensurable root- $D$  chains connect only by approximate interpolation (the residual does not close the PD transfer), and the margins fall along the tower — the wall stands, in canonical nested form. *Companion diagnostics of the simplification trilogy:* the gluing identity is an identity-level off-line detector — it sees  $\delta = 0.01$ , sharper than the PD break at  $\delta^* = 0.02$ , and the envelope slope measures  $\delta$  — while the rigidity thesis is honestly false (the identity holds unconditionally; [verification/v750\_simpler\_gluing.py], 5/5, GLUING-DETECTOR); and the component running is deterministic in direction with the exact  $\mp u/2$  drift anatomy from  $p^2 \equiv 1 \pmod{4}$  and a slope-Ward cancellation at  $4 \times 10^{-4}$  — the margin inherits zero deterministic drift — but the balance sign is zero content (budget  $2.9 \times 10^4 \times$  the margin): the induction route is RH-circular at the balance level ([verification/v751\_ward\_monotone.py], 15/15, WARD-MONOTONE-MIXED). The cone-preserving layer recursion is retired by a one-liner: cone-preserving iff layer-positive, and no prime layer is ( $c_0 = 0$ ; [verification/v743\_schur\_cone\_recursion.py], 7/7, SCHUR-CONE-DEAD). No marker move;

no RH statement.

**The four fixed-size angle instruments (PRIME.ANGLE.INSTR.01).** The instrument cores of the discovery parts T155 and T157, recomputed on the same declared surface as PRIME.RITZ.CEIL.01 (the 12 deepest in-cap frame-A windows,  $n = 29 \dots 139$ ,  $m = 96 \dots 291 \leq 300$ ; a no-go size ladder  $m = 48 \dots 288$ ; parity and Dirichlet controls at  $m = 64, 128, 256$ ; each statement with at least one mutation control that must fail loudly) — none of the four closes an angle, and none is promoted as anything but a per-instance instrument. (1) The  $12 \times 12$  complement-floor certificate:  $\min_{v \perp W} v^T L_P v \geq \mu_{13}^P - \lambda_{\max}(M^{1/2}(I - GG^T)M^{1/2})$  with  $M = \text{diag}(\mu_{13}^P - \mu_k^P)$  and  $G = T_{12}Q_W$  — an identity plus *one* certified Rayleigh bound on the high block, a theorem for every  $m$  and every subspace, direction checked: the certificate never exceeds the exact size- $m$  complement bottom (12/12) and agrees to  $1 - \text{ratio} \leq 1.01 \times 10^{-3}$  against the declared cap  $2 \times 10^{-3}$  on this small surface; exactness controls both ways ( $W = \text{span}\{t_1..t_8\} \rightarrow \mu_9^P$ ; the *worthless* configuration  $W = \text{span}\{t_2..t_{13}\} \rightarrow \mu_1^P$ ; both to  $2.0 \times 10^{-11}$  — the instrument reports its own worthlessness); a corrupted overlap row moves the certificate by  $\geq 0.549$ , only downward. (2) The sine-block confinement: from the certified pencil floor  $A \succeq tL_P$  and the certified ladder ceiling  $\lambda_1 \leq S\mu_1^P$  *alone*,  $\|\gamma_H\|^2 \leq \lambda_1/(t\mu_{17}^P) \leq (S/t)/\rho_{17} = 0.0153\text{--}0.0245 < 1$  on 12/12 — the bottom eigenvector lives 97.5–98.5% inside the first 16 parity sines by a per-instance theorem (measured tail  $6.7 \times 10^{-6}\text{--}3.0 \times 10^{-5}$ ; the discovery part’s measured “98 per cent” replaced by one line); negative control: the pencil ceiling  $\kappa$  in place of the ladder  $S$  is vacuous ( $17.6\text{--}477 > 1$  on 12/12); mutation: the inflated floor  $\kappa L_P$  is refuted at  $e_1(A)$  by  $3.3 \times 10^3\text{--}1.0 \times 10^5$ . (3) The Schur-entry identity:  $s = \mu_1^P t_1^T A^{-1} t_1 = (B^{-1})_{11}$  exactly ( $1.1 \times 10^{-11}$ ) and  $(B^{-1})_{LL} S_L = I$  to  $5.8 \times 10^{-11}$  with  $S_L = B_{LL} - B_{LH} B_{HH}^{-1} B_{HL}$  the *same* fixed-size  $16 \times 16$  Schur complement the chain already forms, with the literal-entry direction  $1/s \leq (S_L)_{11}$  on 12/12:  $1/s = 4.0460\text{--}5.8962$  is an  $O(1)$  number carried by one entry, so the  $R2''$  loss ceiling  $r \leq 1/(Ls)$  is a statement about *one* number; negative control: the inversion-free Cauchy–Schwarz ceiling is valid and useless, off by  $250\text{--}7.27 \times 10^3$  (the cancellation in  $b^T B_{HH}^{-1} b$  is nearly complete); mutation: a one-index shift breaks the identity by  $1.55\text{--}2.15$ . (4) The adaptive Lipschitz ceiling: interval bisection with the *local* Lipschitz constant certifies  $\inf f^{\text{arch}}/(4 \sin^2(\theta/2)) \geq 0.25$  on 12/12 windows in 519–1201 symbol evaluations (certified floor  $0.2505\text{--}0.3062$ , never above the fine-grid minimum; the global constant is documented insufficient — a uniform grid would need up to  $61.4\times$  more); mutation: fed a target above the true grid minimum it *refuses* on 12/12. Stress: on the T145 no-go family the Schur floor survives ( $t = 0.05$ ) while the confinement goes vacuous ( $7.93\text{--}413, \times 52.0$ ) and  $(S_L)_{11}$  explodes ( $1185 \rightarrow 7.11 \times 10^4, \times 60.0$ ) against the flat  $4.05\text{--}5.90$  real band; parity control exact to  $4.9 \times 10^{-14}$ ; the Dirichlet residual hits the predicted  $2/\sqrt{N}$  to  $7.4 \times 10^{-5}$ . **Named limits as load-bearing content:** per instance on small windows — a finite list with an explicit maximum; nothing uniform in the zone index or in  $D$ , and explicitly *no* statement for all  $D$ ; *nothing about the angles is closed* — the two open terms after T157 (an  $m$ -free lower bound on  $\hat{g}_1^2$ , equivalently an  $m$ -free upper bound on the one Schur diagonal entry  $1/s$ ; an  $m$ -free  $R1''$  domination with a non-shrinking margin — certified per window with a  $7.3 \times 10^{-4}$  margin at the largest window and a shrinking trend) stay open, typed open; every floor  $t$  and ceiling  $S$  consumed is certified per window and its  $m$ -freeness is exactly what remains open; PRIME.RITZ.CEIL.01’s ceiling and PRIME.ODD.SECTOR.IDENT.01’s ladder stay as stated and nothing is re-opened. Fences: Kac–Murdock–Szegő 1953 / Schur 1917–Haynsworth 1968 / Courant–Fischer 1920–Cauchy 1829 / Cauchy–Schwarz (refuted as a ceiling, wired as a control) / Kantorovich 1948 / Temple 1928–Kato 1949 (floor devices, not used here) / Sylvester 1852–Bunch–Kaufman 1977 / Wilkinson 1968–Higham 2002 / Chebyshev 1852 / Szegő 1915–Widom 1958 (a named address, never a licence) named classical; Weil 1952 cited, never used as a criterion; zero-firewall AST-checked; no marker moves. Machine-checked: [verification/v552\_angle\_instruments.py]. [E] (per-instance identities and certified inequalities)

**The exact-form identities (PRIME.EXACT.FORM.IDENT.01).** The theorem cores of the discovery parts T158 and T159, recomputed on the same declared surface as PRIME.RITZ.CEIL.01 and PRIME.ANGLE.INSTR.01 (the 12 deepest in-cap frame-A windows,  $n = 29 \dots 139$ ,  $m = 96 \dots 291 \leq 300$ ; no-go ladders  $m = 48 \dots 288$  on *both* T145 forms; parity and Dirichlet controls at  $m =$

64, 128, 256; each statement with at least one mutation control that must fail loudly) — the algebra of the sixteen-form itself, and it closes neither open term. (1) The Thomson dual form and the positive Cholesky ladder:  $s = (B^{-1})_{11} = \max_x (2x_1 - x^\top Bx)$  (Maz'ya 1985 / Miclo 1999) with the direction *executed*, not asserted — 8 random trials per window never exceed  $s$ ; the ladder  $g_K = \sum_{j \leq K} y_j^2$ ,  $y = L^{-1}e_1$  from one Cholesky of the leading  $16 \times 16$  block (nested Schur complements): all terms strictly positive, partial sums strictly monotone,  $1/s \leq 1/g_{16} \leq 1/g_1 = B_{11}$  literally on 12/12, tight to 1.0642–1.1522 against the declared cap 1.5; mutation: a 5% corruption of one off-diagonal entry destroys positive definiteness — the Cholesky *refuses* on 12/12. (2) The two-dimensionality:  $\text{span}\{t_1, A^{-1}Lpt_1\}$  attains  $s$  exactly ( $g/s = 1$  to  $9.0 \times 10^{-13}$ ) for the theorem reason  $Lpt_1 = \mu_1^P t_1$  (KMS 1953); negative controls: the two-sine span misses by 1.42–4.8, the corrupted Green column by  $255\text{--}7.7 \times 10^3$ . (3) The gauge identity and the Toeplitz-minus-Hankel signature: the reflection-odd section annihilates constant lag vectors *entrywise with zero tolerance* ( $\text{odd\_toeplitz}(\mathbf{1}) = 0$  exactly), hence  $\sum_d w_d = 0$  to  $1.1 \times 10^{-14}$  of  $\sup |w|$  — the form is blind to the lag mass of  $c$ , exactly where the archimedean and atom halves are  $h^2$ -sized and cancel; the signature  $A_{r,s} - A_{r+1,s+1} = f(r+s)$  holds to  $1.4 \times 10^{-16}$ ; mutations: a one-index Hankel shift breaks the identity by  $2.7 \times 10^{-7}\text{--}4.2 \times 10^{-5}$ , and the T145 form has signature defect  $7.4 \times 10^{-2}$  — not a parity section, the identity is not even defined for it. (4) The exact lag sum and the two closed scalars:  $x^\top B_{LL}x = \sum_d c_d w_d$  ( $w_d = T_d - H_d$ ) against the direct assembly to  $2.7 \times 10^{-16}$  of the absolute scale  $\sum_d |c_d w_d|$ , with the honest price *measured* (relative to the cancelled total only  $1.4 \times 10^{-11}$  on this small surface — that depth is the  $h^2$  obstruction);  $w_0 = \sum_k x_k^2 / \mu_k^P$  to  $4.2 \times 10^{-16}$  and  $2w_0 - w_1 = \|x\|^2$  to  $7.4 \times 10^{-16}$ ; mutations: a 1% lag perturbation and a one-index  $\mu$  shift both break loudly. (5) The Z-matrix law and the Collatz–Wielandt floor:  $B_{HH}^{\text{arch}}$  is a symmetric Z-matrix *raw* — every off-diagonal entry  $\leq 0$  with no tolerance on 12/12; the closed sign-based floor  $\min_k (B_{HH}^{\text{arch}} \mathbf{1})_k = 0.9093\text{--}1.1284 \geq 0.25$  with the theorem direction verified per instance against the exact  $\lambda_{\min} = 0.9138\text{--}1.1432$  (Perron 1907 / Frobenius 1912 / Collatz 1942 / Wielandt 1950), beating the sign-free norm route by 0.69–8.8; the honest limit as content: the atom block obeys no sign law (0.44–0.53, a coin toss), the full block is not a Z-matrix and its row-sum floor is negative (−48.3... − 0.79) — admissible and vacuous, R1'' stays open; mutation: the checkerboard similarity (an isometry,  $\lambda_{\min}$  moves by 0 exactly) destroys the raw law — an entrywise fact, not a spectral one. Stress: the T145 no-go breaks on three axes of this module's structure; on the  $c(l) = 1/(1+l)$  reconstruction the ladder theorem survives while  $1/g_{16}$  explodes ( $1187 \rightarrow 7.11 \times 10^4$ ) — the instrument reports the collapse instead of hiding it; parity: for  $A = L_P$  the machinery returns  $s = 1$  with every ladder partial sum constant in  $K$ ; Dirichlet controls exact. **Named limits as load-bearing content:** per instance on small windows — a finite list with an explicit maximum; nothing uniform in the zone index or in  $D$ , and explicitly *no* statement for all  $D$ ; the two open terms after T159 (the one explicit pairing inequality for R2'' —  $\sum_d (\Delta c)_d W_d^1$  with  $\|\Delta c^{\text{arch}}\|_1 \leq 6$ , needing a correlation bound instead of sizes; a third device for the atom off-band block for R1'', neither norm nor sign) stay open, typed open; the pairing bound and  $m$ -freeness are explicitly *not* claimed. Fences: Maz'ya 1985 / Miclo 1999 / Schur 1917–Haynsworth 1968 / KMS 1953 / Dirichlet 1829 / Abel / Fejér 1915 / Perron 1907–Frobenius 1912 / Collatz 1942–Wielandt 1950 / Gantmacher–Krein 1950 named classical; Weil 1952 cited, never used as a criterion; zero-firewall AST-checked; no marker moves. Machine-checked: `[verification/v553_exact_form_identities.py]`. [E] (per-instance identities and certified inequalities)

**The sampling/harmonics identities (PRIME.SAMPLING.HARM.01).** The closed theorem cores of the discovery parts T160 and T161, recomputed on the same declared surface as PRIME.EXACT.FORM.IDENT.01 (the 12 deepest in-cap frame-A windows,  $n = 29 \dots 139$ ,  $m = 96 \dots 291 \leq 300$ ) plus a declared  $\Lambda$ -battery of six log-spaced cut-offs  $X = 10^3 \dots 3.2 \times 10^5$  (finite von-Mangoldt sums only, no matrix anywhere; each statement with at least one mutation control that must fail loudly) — what the sixteen-form *pairs against*, and it closes no open term. (1) The sampling identity: the atom half of the pairing equals  $-\frac{1}{2} \sum_{n \leq X} (2\Lambda(n)/\sqrt{n}) \hat{w}(\log n)$  (plus one reflected term below the grid;  $\hat{w}$  the linear interpolant of the lag weights), exact to  $5.6 \times 10^{-15}$  of the absolute atom scale on 12/12 — hence identically a finite combination of sums

$\sum_n \Lambda(n)n^{-1/2} \cos(t \log n)$  at 32 explicit frequencies; mutation: a one-index grid shift breaks by  $5.1 \times 10^{-5}$ – $3.1 \times 10^{-3}$ . (2) The moment laws and the total-variation theorem:  $m_0 = 0$  ( $1.0 \times 10^{-16}$  of  $\|w\|_1$ ),  $m_1 = -[S_0^2 + 2 \sum_j P_j^2]$  ( $9.2 \times 10^{-16}$ ) and  $m_2 = -[2S_1 - (M-1)S_0]^2$  ( $3.4 \times 10^{-14}$ , a perfect square), both strictly negative on every window; U3 as a theorem with *checked* hypotheses ( $c^{\text{arch}} \leq 0$  and monotone on  $d \geq 1$ , checked 12/12):  $\text{TV}(c^{\text{arch}}) = 3.9817$ – $4.8991 \leq |c_0| + 2 \sup < 6$ , closed and window-independent; mutations: a prefix shift, an  $(M-2)$  mis-normalisation, and the full atom-including kernel all fail loudly. (3) The harmonic-frequency theorem and the PNT main term:  $t_j(2\alpha) = 2\pi j$  *exactly* ( $2.8 \times 10^{-14}$ ) — the 32 frequencies are precisely the Fourier harmonics of the log-window  $[0, 2\alpha]$ ; at those frequencies the partial-summation main term collapses to the closed parameter-free prediction  $(\sqrt{X} - 1)/(\frac{1}{4} + t^2)$  (de la Vallée Poussin 1896), matching the finite sums to 0.0007–0.097 of the trivial mass against the declared cap 0.12 on the battery; mutations: a wrong Mellin pole worsens the aggregate residual by  $\times 1.8$ – $\times 95.8$ , and a mass-matched *uniform* fake measure misses by 0.23–0.51 — the agreement is a property of the von-Mangoldt measure. (4) The head split and the scale-free Bernstein rate:  $c_d^{\text{arch}} = \Psi_d + D\hat{G}(dD, D)$  with  $\Psi_d = -\frac{1}{2}[(d+1)\log(d+1) - 2d\log d + (d-1)\log(d-1)]$  closed,  $D$ -free and  $m$ -free (the triangle-smeared simple pole; the  $\log D$  cancels identically), exact to  $3.2 \times 10^{-16}$  at every lag of every window;  $\rho^* = (3 + \sqrt{5})/2 = 2.618034$  is the root of  $x^2 - 3x + 1$ , returned exactly by the  $D \rightarrow 0$  limit of the peeled interval and invariant under  $\alpha \times 10$  (Bernstein 1912); mutations: dropping  $\Psi$  misses by 0.99–1.00, the wrong interval ratio moves the limit to 3.0. (5) The off-diagonal fraction bound and the sign refutation: the gauge identity licenses one boundary-free Abel step, the decomposition  $Q^{\text{arch}} = \sum_{k,l} a_k a_l \sum_d (\Delta c^{\text{arch}})_d R_{kl}^1(d)$  is exact ( $2.0 \times 10^{-13}$ ), and the certified per-window inequality is  $|\text{off-diagonal}| \leq \frac{1}{4}|Q^{\text{arch}}|$  (0.171–0.212 on 12/12) — with the *refuted* sign law wired as the negative control (arch half negative, off-diagonal block positive on 12/12, 176/240 pairs violate the per-pair form): nothing single-signed is promoted; mutation: a one-index  $R$ -kernel shift breaks by  $\geq 6.8 \times 10^{-5}$ . Stress: the head split misses the T145 no-go reconstruction by 2.39 of  $\sup |c|$ ; the Dirichlet cosine-sum identity (including the degenerate branch), the KMS eigenpairs and odd  $\text{toeplitz}(c^L) = L_P$  (zero tolerance) are exact. **Named limits as load-bearing content:** per instance on small windows — a finite list with an explicit maximum; nothing uniform in the zone index or in  $D$ ; finite  $\Lambda$ -sums are allowed and used, *no* RH statement is made, and the  $\delta$  triage of T161 (required cancellation depth against RH strength) is a documentation item, *not* a claim of this module; the open terms after T161 (R-A' the prime-free log-moment  $\sum_d \Psi_d w_d$ , outside the polynomial ladder; R-B' the  $m$ -free  $16 \times 16$  Gram fraction bound, Fejér 1915 / Schur 1917; R-C' *one* split of the pairing with  $\delta_{\text{eff}} < \frac{1}{2}$ , T161 refuting the arch/atom and arch+smooth/fluctuation splits with numbers; R-D a fifth R1'' device) stay open, typed open. Fences: Dirichlet 1829 / Abel 1826 / Fejér 1915–Schur 1917 / Bernstein 1912 / Chebyshev 1852–Mertens 1874 / de la Vallée Poussin 1896 / Clausen–Lerch / KMS 1953 named classical; Weil 1952 / Bombieri 2000 cited, never used as a criterion; zero-firewall AST-checked; no marker moves. Machine-checked: [\[verification/v554\\_sampling\\_harmonics\\_identities.py\]](#). [E] (per-instance identities and certified inequalities)

**The Pareto/total-variation identities (PRIME.PARETO.TV.01).** The closed theorem cores of the discovery parts T162 and T163, recomputed on the same declared surface as PRIME.SAMPLING.HARM.01 (the 12 deepest in-cap frame-A windows,  $n = 29 \dots 139$ ,  $m = 96 \dots 291 \leq 300$ ), with the one arithmetic input  $\kappa = 0.038821$  *verified* at every jump point of  $\psi(x)/x$  on the von-Mangoldt table in  $[100, 3.2 \times 10^5]$  (reproduced to  $4.2 \times 10^{-7}$ ; each statement with at least one mutation control that must fail loudly) — the *price structure* of the pairing, and it closes no open term beyond the negative closure that the floor inequality itself is. (1) The exchange law and the crossing criterion:  $\delta_{\text{bnd}}(x) = \frac{1}{2} + \log(2\kappa g_{16} \text{TV}(x)/P(x))/\log X$  is an identity at all 612 grid points of the declared Fejér knob (50 geometric  $\sigma$  on  $[1, 4096]$  plus  $\infty$ ;  $5.1 \times 10^{-16}$  relative, admissibility  $Q(x) > 0$  everywhere), and  $\delta_{\text{bnd}} < \frac{1}{2} \iff P > 2\kappa g_{16} \text{TV}$  holds as an exact equivalence at every point — price and demand are two coordinates of *one* object; mutation: the wrong constant ( $\kappa$  for  $2\kappa$ ) moves the law by  $6.5 \times 10^{-2}$ –3.4. (2) The total-variation floor — the theorem that carries the T163 verdict: (T1)  $w_0(x) = \|a(x)\|^2$  ( $4.9 \times 10^{-16}$ ; orthonormal parity

rows); (T2)  $\text{TV}(x) \geq |w_0(x)|$  by telescoping with  $w_M = 0$ ; (T3)  $x_1 = 1$  forces  $\|a\|^2 \geq 1/\mu_1^P$ , hence  $\mu_1^P \text{TV}(x) = 5.00\text{--}28.84 \geq 1$  on all 708 trial vectors built (51 knob settings plus 8 random admissible sixteen-vectors per window), the closed floor  $1/\mu_1^P = 1/(4\sin^2(\pi/N)) = 943.6\text{--}8609.6 \sim h^2$ ; (T4) every sub- $\frac{1}{2}$  grid point pays  $P > 2\kappa g_{16}/\mu_1^P = 15.4\text{--}142.8$  (all 54 respect it) — a flat-price sub- $\frac{1}{2}$  demand is impossible in the parity sector on this surface, the inequality that closes R-C''' negatively; negative control: the *inadmissible* vector  $0.01x^*$  ( $x_1 \neq 1$ ) does undercut the floor ( $\mu_1^P \text{TV} = 2.5\text{--}2.9 \times 10^{-3}$ ) and is recognised as inadmissible — the floor is a property of the normalisation meeting the smallest parity eigenvalue (Kac–Murdock–Szegő 1953). (3) The chain-derived price axis:  $\sigma = 1$  gives  $x = e_1$  exactly with  $P = g_{16}/g_1$  (the  $K = 1$  rung of the T158 Cholesky ladder, the T157 route-(0) certificate;  $4.1 \times 10^{-16}$ ),  $\sigma = \infty$  gives  $g_{16}Q(x^*) = 1$  (the  $K = 16$  top rung;  $1.2 \times 10^{-12}$ ); the ladder carries all 16 terms strictly positive with monotone partial sums, and the knob is monotone in both coordinates on every window — the knob curve *is* the Pareto front; mutation: a corrupted off-diagonal sixteen-block entry makes the Cholesky *refuse* on 12/12. (4) The Mellin cell moments and the Abel step: the  $k$ -th rung  $\sum_d w_d C_d(-(2k + \frac{1}{2}))$  with  $C_d(s)$  the closed tent integral (Mellin 1896) agrees with an independent per-cell Gauss–Legendre quadrature to  $2.8 \times 10^{-13}$  on every rung  $k = 1 \dots 4$  (shifted tent centres break by  $\geq 1.3 \times 10^{-2}$ ); the boundary-free Abel identities hold at levels  $k = 1 \dots 3$  ( $8.3 \times 10^{-17}$ ; Abel 1826), the dropped-boundary level-1 form agreeing because the gauge identity kills the boundary term exactly; the optimal level is *one* for T162's closed reason ( $32\pi/\alpha = 40.6\text{--}59.4 > 1$ , measured level-2/level-1 ratio  $30.8\text{--}69.1 > 1$ ), and a gauge violation breaks the dropped form by its closed prediction  $v_0 c_0 M$  to  $1.2 \times 10^{-15}$  — the positive control. (5) The Lerch/Frullani log-moment:  $L = \sum_d \Psi_d w_d$  agrees on three independent routes on every window — direct; double-Abel *with* the Wronskian boundary term  $-\frac{1}{2}M \log M w_{M-1}$  ( $2.8 \times 10^{-14}$ ; dropping it breaks by  $2.2 \times 10^{-4}\text{--}1.3 \times 10^{-2}$ ); the Lerch/Frullani integral with the  $d = 1$  term peeled closed ( $1.8 \times 10^{-15}$ ; the peel constant  $2 \log 2$  to  $2.2 \times 10^{-16}$ ,  $\Psi_1 = -\log 2$  exact; Frullani 1828 / Lerch 1887 / Clausen 1832). Stress: the monotone no-go staircase — coarsening over  $\nu = 4/2/1/0.5$  degrades the arch/atom cancellation monotonically in the median ( $1.00 \rightarrow 3.49 \rightarrow 10.81 \rightarrow 24.98$ , the  $\nu = 4$  column an exact null control); the Dirichlet cosine-sum identity ( $2.4 \times 10^{-14}$ , including the degenerate branch), the KMS eigenpairs ( $1.4 \times 10^{-14}$ ) and  $\text{odd\_toeplitz}(c^L) = L_P$  (zero tolerance) are exact. **Named limits as load-bearing content:** per instance on small windows — a finite list with an explicit maximum; nothing uniform in the zone index or in  $D$ ; the T163  $h$ -exponents ( $P_{\text{aff}} \sim h^{3.05}$ ,  $P_{\text{cross}} \sim h^{1.91}$ , the TV exponent  $\sim h^{2.00}$ ) are sandbox fits and are *not* consumed — what is consumed is the closed per-window floor; finite  $\Lambda$ -sums are allowed and used, *no* RH statement is made ( $\frac{1}{2}$  appears in exactly one role: as the strength of a classical statement a required depth is compared against); the open terms after T163 (R-E the successor, both arms prime-free — arm A: a growing  $1/s$  ceiling for the downstream T159/T160 chain; arm B: a sector whose lowest eigenvalue does not vanish like  $h^{-2}$ ; R-B''' the  $h$ -uniform positivity margin; R-D a fifth R1'' device) stay open, typed open. Fences: Fejér 1915 / Mellin 1896 / Abel 1826 / Chebyshev 1852–Rosser–Schoenfeld 1962 (verified on the table) / Legendre / Dirichlet 1829 / KMS 1953 / Frullani 1828–Lerch 1887–Clausen 1832 / Schur 1917–Haynsworth 1968 named classical; Weil 1952 / Bombieri 2000 cited, never used as a criterion; zero-firewall AST-checked; no marker moves. Machine-checked: [\[verification/v555\\_pareto\\_tv\\_identities.py\]](#). **[E]** (per-instance identities and certified inequalities)

**The gauge/ $P_{\text{pr}}$  identities (PRIME.GAUGE.PPR.01).** The closed theorem cores of the discovery parts T164 and T165, recomputed on the same declared surface as PRIME.PARETO.TV.01 (the 12 deepest in-cap frame-A windows,  $n = 29 \dots 139$ ,  $m = 96 \dots 291 \leq 300$ ), with the one arithmetic input  $\kappa = 0.038821$  *verified* at every jump point of  $\psi(x)/x$  on the von-Mangoldt table in  $[100, 3.2 \times 10^5]$  (each statement with at least one mutation control that must fail loudly) — the *gauge structure* of the pairing, and it does not close the one remaining quantifier. (1) The gauge theorem (T164):  $x_1 = 1$  fixes  $a_1 = 1/\sqrt{\mu_1}$ , i.e. only the *scale* of  $a$ ;  $Q$  and  $\text{TV}$  are homogeneous of degree two in  $a$ ; hence the demand ratio  $Q/\text{TV}$  of the same  $a$ -direction — and with it  $\delta_{\text{bnd}}$  — is the same number in *every* sector, unchanged to  $2.9 \times 10^{-11}$  on 2520 sector-by-trial-vector combinations

(5 weight families  $\times$  7 shifts  $\times$  6 Fejér taper widths on 12 windows), with  $\text{TV} \geq |w_0|$  in every sector and the pure shift scaling the certified value by exactly  $\theta = \mu_1^P/(\mu_1^P + \kappa_s)$  ( $2.2 \times 10^{-11}$  on 432 points); the full space is settled in closed form and strictly worse ( $\mu_0 = 0$  exactly, the lowest remaining eigenvalue  $\leq \mu_1^P$  on 12/12); mutation: the wrong transport map (the  $x$ -coordinates carried unchanged instead of the  $a$ -direction) moves the ratio by  $2.0 \times 10^2$ – $5.4 \times 10^3$ . (2) The  $h^2$  conservation law (T164): floor  $\times$  transfer  $= 1/(\mu_1^P + \kappa_s) \cdot (1 + \kappa_s/\mu_1^P) = 1/\mu_1^P$  *identically* at every shift on the declared grid (72 window-by-shift points,  $2.2 \times 10^{-16}$  of  $1/\mu_1^P$ ) — the  $h^2$  is conserved, never removed; it lives in the ratio, not in the floor; mutation: the wrong transfer misses by 0.9995 at the largest shift on every window. (3) The power-one spend as an identity (T164): the  $r$ -ceiling  $r \leq U/L$  with  $L = \lambda_1(A)/\mu_1^P = 1.1019$ – $1.5845$  (computed per instance from the full window) is linear in the entry ceiling, so  $d \log r / d \log U = +1$  exactly (within  $1.1 \times 10^{-14}$  on 12/12) — an  $h^\varepsilon$  ceiling costs  $h^{-\varepsilon}$  of  $r$ -margin for every  $\varepsilon > 0$ , no free tolerance in the consumption step (the measured power of the full T156 loss margin stays sandbox); mutation: the quadratic consumption reads slope 2.0000 and is recognised. (4) The  $P_{\text{pr}}$  identity and the predicate equivalence (T165):  $P_{\text{pr}} = g_{16} \cdot R \cdot (\text{TV}/(t_1 v)^2)/\mu_1^P$  with no inequality anywhere — four factors, each a named object: the R-E-A quantifier  $g_{16} = 0.1559$ – $0.2227$ , the demand  $R = Q/\text{TV}$ , the T163 floor  $\text{tvf} \geq 1$ , and the KMS 1953 scale  $1/\mu_1^P = 943.6$ – $8609.6 \sim h^2$  — exact to  $4.3 \times 10^{-16}$  on all 192 battery vectors (6 knob settings, 6 random sixteen-vectors and 4 random *full-window* vectors per window), gauge-invariant under  $v \rightarrow tv$  to  $1.5 \times 10^{-11}$  against the declared cancellation horizon; the predicates  $R > 2\kappa$  and  $P_{\text{pr}} > 2\kappa g_{16} \text{tvf}/\mu_1^P$  agree as booleans on all 192 (no search is run — T165's ascent stays sandbox), so the demand half of R2'' and the gauge-invariant crossing condition are *one* predicate — R-F is strictly stronger than the R2'' demand and its  $m$ -uniform form *is* R-E-A — and  $\text{tvf} \geq 1$  pointwise, hence every crossing vector pays  $P_{\text{pr}} > 2\kappa g_{16}/\mu_1^P = 15.4$ – $142.8$ , above the preregistered tolerance  $\Theta_{\text{TOL}} = 10$  on 12/12 windows: the R-F conjunction is *empty* on this surface by an identity plus a floor; mutation: the price with  $(t_1 v)$  unsquared scales like  $t$  and moves by  $1.0$ – $1.0 \times 10^4$  under the same round trips. (5) The Schur-cascade direction and the exponent closure (T164/T165):  $s = \mu_1^P t_1^T A^{-1} t_1 = 0.1696$ – $0.2472 \geq g_{16} \geq g_1$  window by window ( $1/s \leq 1/g_{16} \leq 1/g_1 = B_{11}$ ; Schur 1917–Haynsworth 1968),  $P_K = \hat{a} s = 254.8$ – $7.7 \times 10^3 \geq 1$  (Kantorovich 1948); the exponent closure is an identity of the lists — per instance  $\log(1/g_1) = \log(g_{16}/g_1) + \log(1/g_{16})$  to  $1.6 \times 10^{-15}$ , the least-squares fits closing to  $1.3 \times 10^{-15}$  by linearity (the exponent *values* are sandbox fits, not consumed — what is promoted is that the whole certification lives in the cascade gain  $g_{16}/g_1$ , the object of the one open quantifier); mutations: the  $g_8$  closure breaks by 0.10–0.24, a corrupted sixteen-block makes the Cholesky *refuse* on 12/12. Stress: the monotone no-go staircase ( $1.00 \rightarrow 3.49 \rightarrow 10.81 \rightarrow 24.98$ , the  $\nu = 4$  column an exact null control); the Dirichlet cosine-sum identity ( $2.9 \times 10^{-14}$ , including the degenerate branch), the KMS eigenpairs ( $1.4 \times 10^{-14}$ ) and  $\text{odd\_toeplitz}(c^L) = L_P$  (zero tolerance) are exact. **Named limits as load-bearing content:** per instance on small windows — a finite list with an explicit maximum; nothing uniform in the zone index or in  $D$ ; the T164/T165  $h$ -exponents and the measured anatomy (the free ascent, the  $\eta(\text{Cap})$  profile, the decoupled  $\nu$  surface with its zone-depth finding, the retired constant  $U_{\text{ref}} = 4.90 \rightarrow 5.7327$ ) stay in the diary and the paper; finite  $\Lambda$ -sums are allowed and used, *no* RH statement is made ( $\frac{1}{2}$  appears in exactly one role: as the strength of a classical statement a required depth is compared against); **the one open object after T165 stays open, typed open:**  $\inf_m g_{16}(m) > 0$ , equivalently a lower bound on the sixteen-step Schur-cascade gain  $g_{16}/g_1$  uniform in  $m$  — a quantifier over a certified flat list; R-B''' (an independent  $\alpha/h$  surface) stays open as narrowed by T164. Fences: Schur 1917–Haynsworth 1968 / KMS 1953 / Abel 1826 / Chebyshev 1852–Rosser–Schoenfeld 1962 (verified on the table) / Kantorovich 1948 / Fejér 1915 / Dirichlet 1829 named classical; Weil 1952 / Bombieri 2000 cited, never used as a criterion; zero-firewall AST-checked; no marker moves. Machine-checked: [verification/v556\_gauge\_ppr\_identities.py]. [E] (per-instance identities and certified inequalities)

**The cascade/vector identities (PRIME.CASCADE.VECT.01).** The closed theorem cores of the discovery parts T166 and T167, recomputed on the same declared surface as PRIME.GAUGE.PPR.01 (the 12 deepest in-cap frame-A windows,  $n = 29 \dots 139$ ,  $m = 96 \dots 291 \leq 300$ ), with the one arithmetic input  $\kappa = 0.038821$  *verified* at every jump point of  $\psi(x)/x$  on the von-Mangoldt table

in  $[100, 3.2 \times 10^5]$  (each statement with at least one mutation control that must fail loudly) — the *anatomy of the cascade* and the *exactness of the closed vector at  $K = 2$* , and it does not close the one remaining scalar. (1) The cascade identities (T166): the prefix property — one Cholesky of the sixteen-block yields all sixteen rungs,  $g_K = e_1^\top Q_K^{-1} e_1$  at every  $K$  (192 comparisons,  $1.1 \times 10^{-12}$ ); the minor form  $1/g_K = \det G_{[1..K]} / (\mu_1^P \det G_{[2..K]})$  on the raw arithmetic Gram block  $G = TAT^\top$  ( $2.7 \times 10^{-12}$  on 36 (window,  $K$ ) pairs,  $K \in \{2, 6, 16\}$ ; Schur 1917–Haynsworth 1968); the regression form  $g_K/g_1 = 1/(1 - R_K^2)$  with  $R_K^2$  the squared multiple correlation of mode 1 on modes  $2 \dots K$  ( $8.8 \times 10^{-13}$ ) — the open object *is* a near-collinearity statement; mutation: the minor with the wrong row/column deleted misses by 0.99–1.45. (2) The diagonal invariance (T166): the gain  $g_K/g_1 = B_{11}(Q_K^{-1})_{11}$  is unchanged under  $B \rightarrow DBD$  on 144 random positive diagonals ( $2.7 \times 10^{-12}$ ) — the cascade does not see the KMS  $1/\sqrt{\mu}$  normalisation, the gain is a property of the arithmetic Gram block alone; mutation: a non-diagonal congruence moves the gain by 2.7–3.8. (3) The exact  $K = 2$  identity (T167): the closed vector  $u = (1, -Q_{21}/Q_{22})$  attains the constrained minimum,  $u^\top Q_2 u = 1/g_2 = Q_{11}(1 - r_{12}^2)$  to  $1.1 \times 10^{-12}$ , and  $B_{11}g_2 = 1/(1 - r_{12}^2)$  identically to  $1.2 \times 10^{-12}$  (measured  $1 - r_{12}^2 = 2.8 \times 10^{-4}$ – $5.9 \times 10^{-3}$  on this surface, the  $h$ -trend a sandbox fit): at  $K = 2$  the fast-null-vector needs no approximation — the third dress collapses onto the second, one scalar built from three closed lag sums; mutation: the wrong-sign vector pays the relative excess  $4r_{12}^2/(1 - r_{12}^2) = 6.7 \times 10^2$ – $1.5 \times 10^4$ . (4) The pivot identity and the unification (T167): for  $u = x_K^* + \delta$  with  $\delta_1 = 0$ ,  $u^\top Q_K u = 1/g_K + \delta^\top Q_K \delta$  exactly ( $2.0 \times 10^{-12}$  on 252 evaluations across three  $\delta$  scales) — every relative excess is an exact number; the Rayleigh envelope holds with its  $\delta_1 = 0$  ceiling attained exactly ( $1.1 \times 10^{-15}$ ) and Cauchy-interlaced; the sign-aligned entrywise perturbation attains  $|x^{*\top} E x^*| = \varepsilon S_K$  exactly ( $3.8 \times 10^{-16}$ ), so the entry threshold  $\rho_{\text{carry}}(1/g_K)/S_K$  is exact — and the unification:  $(1/g_K)/S_K$  is the one cancellation ratio of all three dresses, at  $K = 2$  literally  $(1/g_2)/S_2 = (1 - r_{12}^2)/(1 + 3r_{12}^2)$  ( $1.2 \times 10^{-12}$ ; the closed factor  $1 + 3r_{12}^2 \rightarrow 4$  is T166’s independently derived quarter): *one* inequality, not three; mutation:  $\delta_1 = 0.1$  breaks the pivot identity by exactly the closed cross term  $2\delta_1 = 0.2$ . (5) The trial theorem and the two must-breaks (T166/T167): all 288 random trials with  $u_1 = 1$  satisfy  $u^\top Q_K u \geq 1/g_K$  (min ratio 15.14) — every closed candidate is a certified *lower* bound (Schur 1917); the  $L_P$  no-go: the normalised block is the identity ( $1.6 \times 10^{-15}$ ) and the cascade gain exactly 1 (deviation 0.0) at  $m = 64/128/256$ ; the scramble type-change: uniform random atom positions at the same total mass make  $B_{11}$  itself lose positivity on 49/60 declared draws and on a majority of draws on 12/12 windows — at  $K = 2$  the collapse under scrambling is a change of *type*, not a factor. Stress: exact Dirichlet/parity controls ( $2.3 \times 10^{-14}$  /  $1.4 \times 10^{-14}$  / zero tolerance). **Named limits as load-bearing content:** per instance on small windows — a finite list with an explicit maximum; nothing uniform in the zone index or in  $D$ ; the T166/T167  $h$ -exponents ( $1 - r_{12}^2 \sim h^{-2.921}$  on frame A, the certified gain exponent  $h^{+2.921}$ , the threshold monotonicity in  $K$ , the  $K = 6$  separation  $h^{+1.138}$ ) and the measured anatomy (the rung profile, the polarisation split, the Kato radius, the headroom) stay in the diary and the paper; finite  $\Lambda$ -sums are allowed and used, *no* RH statement is made ( $\frac{1}{2}$  appears in exactly one role: as the strength of a classical statement a required depth is compared against); **the one open object after T167 stays open, typed open:** R1, the single scalar — an  $m$ -free upper bound on  $1 - r_{12}^2 = 1 - Q_{12}^2/(Q_{11}Q_{22}) \leq C h^{-3+\varepsilon}$ , three closed lag sums and nothing else; the routes T167 closed off (perturbation theory of every order around the KMS bottom mode — the Kato series converges to the *wrong* object; position-blind surrogates) stay sandbox negatives; R-B''' stays open as narrowed. Fences: Schur 1917–Haynsworth 1968 / KMS 1953 / Rayleigh 1877–Courant–Fischer 1920 / Cauchy 1829 / Kato 1949 (cited as the closed-off route, never used) / Chebyshev 1852–Rosser–Schoenfeld 1962 (verified on the table) / Dirichlet 1829 named classical; Weil 1952 / Bombieri 2000 cited, never used as a criterion; zero-firewall AST-checked; no marker moves. Machine-checked: [verification/v557\_cascade\_vector\_identities.py]. [E] (per-instance identities and certified inequalities)

**The bilinear/rank identities (PRIME.BILINEAR.RANK.01).** The closed theorem cores of T168 (contract LAGRANGE.MINORS), T169 (contract TSTAR.RATIO) and T170 (contract BILINEAR.SIEVE) as one deliberately narrow module (27 checks,  $\sim 0.2$  s), recomputed on the declared small frame-A

surface (12 deepest in-cap zones,  $m = 96 \dots 291 \leq 300$ ;  $\kappa = 0.038821$  verified at every jump point of  $\psi(x)/x$  on the table) — the *exact shape of the one remaining scalar*, and it does not close that scalar. (1) The Lagrange/Wronskian anatomy (T168): the lag-sum identity  $\hat{a}_{ij} = t_i^\top A h t_j = \sum_d c_d W_d^{ij}$  with the closed correlation weights ( $1.6 \times 10^{-15}$ ; the  $1/\sqrt{\mu}$  factors cancel, T169-TH3), Euclidean orthogonality  $t_1 \cdot t_2 = 0$  ( $4.4 \times 10^{-16}$ ) — the near-parallelism is created *entirely* by the arithmetic metric — the Wronskian telescope  $(\mu_1 - \mu_2)u_r v_r = W_{r-1} - W_r$  with  $W_{-1} = W_{h-1} = 0$  ( $1.9 \times 10^{-17}$ ) and the maximal-minor identity  $\frac{1}{2} \sum M^2 = 1/(\mu_1 \mu_2)$  exactly ( $4.6 \times 10^{-16}$ ; Lagrange 1773); mutation: the wrong eigenvalue pair misses by 1.66–1.67 relative (the closed miss  $(\mu_2 - \mu_3)/(\mu_1 - \mu_2) \rightarrow 5/3$ ). (2) The exponent account and the Gershgorin certificate (T168/T169):  $1 - r_{12}^2 = \nu_1 \nu_2 / (\hat{a}_{11} \hat{a}_{22})$  identically ( $2.5 \times 10^{-13}$ ; measured  $1 - r_{12}^2 = 2.8 \times 10^{-4} - 5.9 \times 10^{-3}$  on this surface, the  $h$ -trend a sandbox fit);  $\nu_1 \leq \max(\hat{a}_{11}, \hat{a}_{22}) + |\hat{a}_{12}|$  (Gershgorin 1931, unconditional, closed in the three entries) on 12/12 with loss 1.1885–1.1947; mutation: dropping  $|\hat{a}_{12}|$  fails on every window ( $\nu_1 - \max \text{diag} = 0.30 - 1.17 > 0$ ). (3) The det-collapse theorem (T169-TH7): on the Rayleigh family  $u(t) = (t, -1)$  the candidate  $K7 = \sqrt{\hat{a}_{22}/\hat{a}_{11}}$  collapses in closed form,  $R(K7) = 2\sqrt{\hat{a}_{11}\hat{a}_{22}} \det \hat{A} / [(\hat{a}_{11} + \hat{a}_{22})(\sqrt{\hat{a}_{11}\hat{a}_{22}} + \hat{a}_{12})]$  ( $8.7 \times 10^{-13}$ ) — it meets the threshold precisely because it reintroduces  $\det \hat{A}$ : the loop T167 scalar  $\rightarrow$  T168 factor  $\rightarrow$  T169 candidate closes as an identity; usable form  $\nu_2 \leq 2 \det \hat{A} / (\hat{a}_{11} + \hat{a}_{22}) = 2\nu_1 \nu_2 / (\nu_1 + \nu_2)$ , tight to 1.9980–1.9999  $< 2$  exactly (Rayleigh 1877); mutation: the closed constant  $1/\sqrt{2}$  overshoots  $\nu_2$  by  $17.1 - 2.3 \times 10^2$ . (4) The exact bilinear von Mangoldt form (T170 TH1–TH3): the exact split  $\hat{A} = B - S$  with  $S_{ij} = \sum_{n \leq X} (\Lambda(n)/\sqrt{n}) X_n[ij]$  against the direct block ( $1.7 \times 10^{-14}$ ); the polarisation  $\det \hat{A} = \det B - D(B, S) + \det S$  ( $1.0 \times 10^{-10}$ ); the double sum  $\det S = \sum_{n,m \leq X} \Lambda(n)\Lambda(m)K(n,m)/\sqrt{nm}$  with  $K = \frac{1}{2}D(X_n, X_m)$  through the *explicit* kernel matrix ( $7.5 \times 10^{-16}$ ; Cauchy–Binet 1812–1815); the wedge reading to the measured rank-one defect  $4.6 \times 10^{-3} - 3.1 \times 10^{-2}$  (honest certificate); mutation: the wrong cross-coefficient breaks by 64–430. (5) The rank-3 theorem and the Type II collapse (T170 TH4–TH5): the polarisation matrix on  $(X_{11}, X_{22}, X_{12})$  has eigenvalues  $\{-2, -1, +1\}$  exactly — rank 3, signature (1, 2) — so the kernel  $K = PMP^\top$  has  $\sigma_4/\sigma_1 \leq 1.6 \times 10^{-16}$  with nonzero spectrum =  $\text{spec}(MP^\top P)$  ( $1.0 \times 10^{-15}$ ): the bilinear form *is* the rank-3 polynomial  $S_{11}S_{22} - S_{12}^2$  in three linear  $\Lambda$ -sums; Vaughan’s identity coefficient-exact for every  $U = V \in \{3, 4, 5, 6\}$  (residual 0.0; Vaughan 1977), Type II kernel 99.9%-energy rank  $[2, 2, 2, 2]$  against generic  $[20, 15, 11, 9]$  ( $4.5 \times 10^{-10}$ ; Montgomery–Vaughan 1973 cited as the priced route). (6) The R4-free chain and the two discriminators (T170-TH6):  $1 - r_{12}^2 = \det \hat{A} / (\hat{a}_{11} \hat{a}_{22})$  is an identity (0.0) — no positivity of  $A_h$ , the Weil fence never approached; the  $L_P$  no-go:  $\hat{a}_{12} = 1.1 \times 10^{-17}$  identically, no decay; the scramble discriminator: rank-3 *untouched* on all 60 draws while the value of  $1 - r_{12}^2$  moves by a median factor 834.8 (range  $69 - 1.8 \times 10^4$ ) — the arithmetic sits entirely in the joint values of  $(S_{11}, S_{22}, S_{12})$  against  $B$ . Stress: exact Dirichlet/parity controls ( $1.2 \times 10^{-14}$  /  $1.4 \times 10^{-14}$  / zero tolerance). **Named limits as load-bearing content:** per instance on small windows — a finite list with an explicit maximum; nothing uniform in the zone index or in  $D$ ; the T168–T170  $h$ -exponents ( $1 - r_{12}^2 \sim h^{-2.92}$  on the deep frame-A surface, the route-table deltas — large sieve +0.669, MV −0.384, exact split +0.996, Vaughan +0.669 against the measured truth 2.747 — the five-order cancellation and the  $h^{-2.711}$  row-angle trend) and the measured anatomy stay in the diary and the paper; finite  $\Lambda$ -sums are allowed and used, *no* RH statement is made ( $\frac{1}{2}$  appears in exactly one role: as the strength of a classical statement a required depth is compared against); **the one open object after T170 stays open, typed open:** R1, now finally *classified* but not closed — an  $m$ -free unconditional certificate that the two rows of  $\hat{A}$  become collinear at rate  $h^{-3+\epsilon}$ , equally a joint near-degeneracy of the three explicit linear  $\Lambda$ -sums  $(S_{11}, S_{22}, S_{12})$  against the archimedean block — a rate statement the module neither makes nor approaches; the T168/T169/T170 verdicts (MINORS-RESIST / RATIO-RESISTS / SIEVE-RESISTS) are sandbox verdicts. Fences: Lagrange 1773 / Cauchy–Binet 1812–1815 / Gershgorin 1931 / Rayleigh 1877–Courant–Fischer 1920 / Montgomery–Vaughan 1973–Vaughan 1977 (the priced route) / KMS 1953 / Chebyshev 1852–Rosser–Schoenfeld 1962 (verified on the table) / Dirichlet 1829 named classical; Weil 1952 / Bombieri 2000 cited, never used as a criterion; zero-firewall AST-checked; no marker moves. Machine-checked: [verification/v558\_bilinear\_rank\_identities.py]. [E] (per-instance

identities and certified inequalities)

**The phase-2 capstone (PRIME.PHASE2.CAPSTONE.01).** Promotion of discovery T171 (contract FINAL.MAP, verdict MAP-COMPLETE) as one module (35 checks,  $\sim 0.3$ s): *the capstone verifies the map and claims nothing new* — the whole sixteen-link reduction chain of phase 2 as one machine-checked pass, the eight-route no-go battery, and the precision ledger, recomputed on the declared small frame-A surface (12 deepest in-cap zones,  $m = 96 \dots 291 < 300$ ;  $\kappa = 0.038821$  verified at every jump point of  $\psi(x)/x$  on the table; the  $16 \times 16$  low block positive definite on 12/12 — a numerical fact, never routed through the Weil criterion). (1) The chain, 13 THEOREM / 3 CERT: T152 Schur two-block (KMS  $1.3 \times 10^{-15}$ ; holds at  $0.5\lambda_{\min}$ , fails at  $1.02\lambda_{\min}$  — sharp; Haynsworth  $3.8 \times 10^{-15}$ ); T151 Sobolev  $b_k \leq C_S k$ ,  $C_S = (2\pi + 2)/\sqrt{N}$  (worst ratio 0.4816); T153  $\Psi$ -collapse  $\Psi = \sum_d c_d w_d = x^T G x$  ( $6.3 \times 10^{-16}$ ;  $\sum_d w_d = 0$  to  $9.0 \times 10^{-17}$ ); T154 Ritz ceiling (interlacing;  $g_{32}/s = 0.9520\text{--}0.9849$  from below); T155  $12 \times 12$  floor [CERT] (holds at  $(1 - 10^{-9})\lambda$ , fails at  $(1 + 10^{-4})\lambda$ ;  $t = 0.2295\text{--}0.3719$ ); T156  $F(P, r)$  [CERT] ( $P_K \geq 1$ ,  $r_K \geq 0$ ,  $L(16) = 6.6\text{--}15.6$ ); T158 Thomson/ladder ( $3.1 \times 10^{-12}$ ; increments strictly positive; 36/36 trials below); T160 sampling (closed Dirichlet weights = correlation weights  $1.1 \times 10^{-15}$ ; atom half = sampled  $\Lambda$ -mass  $9.9 \times 10^{-15}$ ); T163 TV floor/swap (Abel  $1.8 \times 10^{-11}$ ; overshoot  $\times 5.2 \times 10^4\text{--}1.1 \times 10^6$ ); T164 gauge/quantifier (invariance  $2.2 \times 10^{-11}$ ;  $\Psi(x^*) = 1/g_{16}$  to  $7.0 \times 10^{-12}$ ; 48/48 perturbed larger); T165  $P_{\text{pr}}$  ( $1.1 \times 10^{-16}$ ; = 1 at  $x^*$ ); T166 cascade  $1/(1 - R_{16}^2)$  ( $1.8 \times 10^{-12}$ ; diagonal-invariant); T167  $K = 2$  exactness ( $6.8 \times 10^{-13}$ ; milder  $\times 1.62\text{--}5.20$ ); T169 det collapse R4-free ( $\leq 4.2 \times 10^{-12}$ ; no positivity of  $A_h$  consumed, the Weil fence never approached); T170 rank-3 ( $J$ -eigenvalues  $\{-2, -1, +1\}$  exact;  $\sigma_4/\sigma_1 \leq 1.6 \times 10^{-16}$ ); T170 R1 shape [CERT] (the  $2 \times 2$  Lagrange identity  $3.6 \times 10^{-16}$ ;  $\sin \angle = 1.169 \times 10^{-4}\text{--}2.467 \times 10^{-3} \leq 10^{-2}$  per window; rate and quantifier OPEN). (2) The no-go battery, 8/8 failing as classified: size budget  $\times 2.7 \times 10^3\text{--}1.6 \times 10^4$ ; symbol infimum  $-47.6\text{--}-21.3 < 0$  on 12/12; sector change an invariance ( $\leq 4.2 \times 10^{-12}$ ); perturbation  $\times 9.8 \times 10^2\text{--}2.0 \times 10^4$  off scale (Kato 1949, the closed-off route); atom band dominates  $\times 1.71\text{--}2.83$ ; sieve operator-norm on  $\det S$  alone  $\times 8.8 \times 10^3\text{--}7.5 \times 10^4$  over (by the rank-3 theorem a bound on  $\det S$  alone cannot see the three-term cancellation; Montgomery–Vaughan 1973 / Vaughan 1977 the priced route); scramble  $\times 679$  median (range  $68\text{--}2.3 \times 10^6$ ) with rank-3 *untouched* on all 60 draws;  $L_P$  gain = 1 to 0.0. (3) The precision ledger: needed  $2.752 \times 10^{-4}$  at the deepest window ( $m = 285$ ) — finer than the RH yardstick  $h^{-0.5}$  by  $215\times$  and the best unconditional  $h^{-0.996}$  by  $13\times$  on this surface; both self-similarity identities ( $6.8 \times 10^{-13}$  /  $3.1 \times 10^{-13}$ , the  $t^*$  overshoot  $\geq 0$  everywhere, sign-agnostic). Mutations fail loudly (wrong regression row, wrong Thomson entry, off-by-one swap, wrong polarisation coefficient — the  $10^{-2}$  bar missed by orders). **Named limits as load-bearing content:** per instance on small windows — a finite list with an explicit maximum; nothing uniform in the zone index or in  $D$ ; the T171 trend fits (pooled sin-angle  $h^{-2.33}$ , frame-split  $\nu = 4$ :  $h^{-1.86}$  /  $\nu = 6$ :  $h^{-2.87}$ , the no-go growth exponents,  $\delta = +1.044$  against 3.0, the full-surface precision numbers  $2.284 \times 10^{-7}$  /  $1.2 \times 10^5 \times$  /  $3.1 \times 10^3 \times$  at  $h = 1444$ ) are sandbox numbers on T171's two-frame surface and are *not* consumed; **zero new uniform-in- $m$  statements are added — which is exactly why the capstone is promotable; the one open phase-2 object stays open, typed open:** R1 — an  $m$ -free unconditional certificate that the two rows of  $\hat{A}$  become collinear at rate  $h^{-3+\varepsilon}$ , a joint near-degeneracy of the three explicit linear  $\Lambda$ -sums against the archimedean block; R2/R3 booked (T169), R4 off the path (T170-TH6); the T171 verdict MAP-COMPLETE is a verdict about the *map*, not about R1, and stays a sandbox verdict ( $\frac{1}{2}$  appears in exactly one role: as the strength of a classical statement a required precision is compared against). Fences: Schur 1917–Haynsworth 1968 / KMS 1953 / Cauchy 1829–Courant–Fischer 1920 / Rayleigh 1877 / Abel 1826 / Dirichlet 1829 / Lagrange 1773 / Cauchy–Binet 1812–1815 / Gershgorin 1931 / Kato 1949 / Montgomery–Vaughan 1973–Vaughan 1977 / Chebyshev 1852–Rosser–Schoenfeld 1962 / Fejér 1915 named classical; Weil 1952 / Bombieri 2000 cited, never used as a criterion; zero-firewall AST-checked; no marker moves. Machine-checked: [verification/v559\_phase2\_capstone.py]. [E] (per-instance identities and certified inequalities)

**The frame/deficit identities (PRIME.FRAME.DEFICIT.01).** Promotion of the theorem cores of discoveries T172 (contract FRAME.BEYOND, verdict PARTIALLY-PORTABLE), T173 (FRAME.RATE,

DEFICIT-VARIES) and T174 (CANCEL . IDENTITY, PARTIAL-CANCEL) as one module (24 checks,  $\sim 0.5$  s): *the gauge route of phase 2, closed as algebra* — the companion to the capstone above, certifying what T172/T173/T174 established *about* the map’s frame dependence, everything recomputed on a small declared surface (a  $4 \times 3$  (anchor,  $h$ ) rectangle of gap-blind frame-B cells, prime-power anchors  $n = 41 \dots 557$  crossed with the declared ladder  $h = (160, 226, 320)$ , comb complete on every cell, positive definite on 12/12; one non-prime-power instance  $n = 210$ ; one truncated-comb stress instance  $n = 1009$ ;  $\kappa = 0.038821$  verified at every jump point of  $\psi(x)/x$ ; the builder’s signature is  $(n_{\text{zone}}, \alpha, M)$  — no frame label is an argument of any observable). (1) The identity transfer + the sieve-horizon discriminator: the KMS/Schur pair ( $3.2 \times 10^{-16}$ ; quadratic-form pair  $5.2 \times 10^{-15}$ ), the lag-sum / correlation identity ( $1.4 \times 10^{-16}$  of the larger half; Abel 1826), the R4-free det collapse (at 0.020 of the conditioned bar  $64\varepsilon/|1 - r_{12}^2|$ ) and the additive split ( $7.6 \times 10^{-16}$ ), re-derived *exactly* on a gap-blind frame-B instance *and* a non-prime-power-anchor instance — the links use no property of the frame rule or the anchor arithmetic, which is why they transfer; the truncated-comb instance turns *indefinite* ( $\lambda_{\min} = -9.6 \times 10^4$ ) against 12/12 positive definite complete-comb cells, and the identity links keep holding there (none uses positivity): indefiniteness sits at the declared sieve horizon, not at a frame or zone. (2) The  $q = 1 - s$  identity + the calibration:  $hD = \alpha$  per cell ( $8.9 \times 10^{-16}$ ) forces the OLS slopes to obey  $q = 1 - s$  to  $2.9 \times 10^{-14}$  (OLS is linear in its response); the gauge-degree account comes out as integers —  $G2 = t/\lambda_{\max}(B_{LL})$  and the delivery  $1 - r_{12}^2$  both  $(d_{\text{amp}}, d_{\mu}) = (0, 0)$ ,  $G1$  at  $d_{\text{amp}} = 1$ ,  $G3$  at  $d_{\mu} = 1$  — so  $G2$  is the unique doubly gauge-invariant arithmetic member of T173’s family;  $G3 = \mu_1^P G2$  identically ( $2.2 \times 10^{-16}$ ; exponent  $h^{-1.9955}$ , the offset-2 explained);  $(D/\alpha)^3 h^3 = 1$  a tautology ( $2.2 \times 10^{-16}$ ). (3) The blindness theorems + the additive limit: the amplitude  $c \rightarrow \kappa c$  is a gauge of GAP, of  $1 - r_{12}^2$  and of  $R$  over twelve decades (worst 0.026 of the per-cell conditioned bar; rebuilt from scratch 0.098); the scalar  $\mu_1^P = 4 \sin^2(\pi/(2h + 1))$  — *the entire  $h^{-2}$  channel of the demand* — cancels exactly on both sides (0.001 of bar); the delivery is invariant under the full positive diagonal group (48 congruences spanning  $e^{\pm 8}$ , worst 0.036 of bar) while a non-scalar diagonal moves the demand by 0.23–1.40 (the asymmetry is the anatomy); neither the arch term nor the comb term alone has a positive definite low block on a single cell (0/12 + 0/12 against 12/12 for the sum): no factorisation  $\hat{A} = (\text{frame}) \times (\text{arithmetic})$  exists, so T173’s P6 is *not* promoted to a theorem — only its multiplicative half. (4) The KMS shape band (unconditional, declared  $\varepsilon = (K\pi/(2h + 1))^2/6$ ):  $|\sqrt{\hat{S}_k/k^2} - 1| \leq \varepsilon_{\text{cert}}$  on 12/12 and the  $k^2$ -swap moves  $\log \text{GAP}$  by  $\leq 0.312$  of the certified band  $2 \log((1 + \varepsilon)/(1 - \varepsilon))$  (Rayleigh 1877; Schur 1917) while the delivery does not move at all. (5) The placement discriminator: the value-multiset scramble destroys the positive low block on 24/24 declared draws and moves  $|R|$  by a median factor  $10^6$ . Mutations fail loudly (the  $d = 0$  halving dropped, the  $N = 2h$  ladder, the off-by-one  $h$ , the comb-only rescale, the wrong universal profile). **Named limits as load-bearing content:** per instance on a small declared surface — a finite list with an explicit maximum; nothing uniform in the zone index or in  $D$ ; **the frame-free deficit is a measurement, not a claim of this row:** T174’s  $+0.1111 \pm 0.0222$  ( $5.0\sigma$ , 151 anchor clusters, cluster-robust;  $0.4\sigma$  from T173’s  $+0.155 \pm 0.102$ ) stays a sandbox MEASURED number, as do all T172/T173/T174  $h$ -exponents, the density anatomy and the  $\nu$ /density collinearity; **the one open phase-2 object stays open, typed open:** R1, now stated frame-free — an  $m$ -free unconditional certificate that two explicit finite  $\Lambda$ -sum vectors become collinear at the demanded rate; the residual per-anchor scatter ( $\chi^2/\text{dof} = 3.02$ ) and the sparse corner are sandbox items under test in T175 ( $\frac{1}{2}$  appears in exactly one role: as the strength of a classical statement a required precision is compared against). Fences: Schur 1917 / KMS 1953 / Dirichlet 1829 / Rayleigh 1877 / Abel 1826 / Chebyshev 1852–Rosser–Schoenfeld 1962 / Wilkinson 1968–Higham 2002 named classical; Weil 1952 / Bombieri 2000 cited, never used as a criterion; zero-firewall AST-checked; no marker moves. Machine-checked: [verification/v560\_frame\_deficit\_identities.py]. [E] (per-instance identities and certified inequalities)

**Lemma 1** (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*

$\int \langle \Theta F, F \rangle_{\kappa_g} d\mu_{\text{Cal}, \chi} \geq 0$  ( $G_4$ ).  **$G_4$  is the decisive step:**  $RP$  after integrating over dynamical metrics, not merely at fixed background.

**Lemma 2** ( $G_5$  — gap dominance [E]).  $2\|V_{\text{metric}}\|_{\text{rel}} < \Delta = 6 \log \frac{3}{2}$ , 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).

## Externalization Contract — PRIME.RESIDUE.EXTERNAL.01 (the prime-front terminal residue, packaged for external specialists)

*Typing.* The third hand-off document: the *terminal open statements* of the prime-front programme (diary and status record: `tfpt_prime_front`), packaged self-containedly for outside specialists in analytic number theory — large sieve, sampling and frame theory, de Branges spaces, explicit-formula methods. It *documents*, it does not close: the contract identifier names this document, not a ledger row; every prime-front ledger row behind it (`PRIME.*`) is an [E]-typed NOT-RH-EVIDENCE identity, finite theorem or honest kill record, and none moves here. **This contract offers open problems. It asserts no progress toward the Riemann Hypothesis in either direction; nothing below is evidence for or against it, and no submitted solution is solicited as such evidence.** Throughout, “ $\lambda$ -uniform” is house shorthand for *uniform over the whole rung ladder* (the infinite quantifier), as opposed to per-rung, finite, or classical-per-census legs; “census” always means the finite root set defined below, never a zero set of  $\zeta$ .

**(1) The finite objects, self-contained.** Fix an integer rung  $h \geq 4$  and set  $a = \frac{1}{2} \log h$ ,  $K = \lceil 1.25 h \log h \rceil$ , modes  $\omega_k = k\pi/a$  ( $k = 0..K-1$ ), poles  $b_k = \omega_k^2$ , and the block height  $T_z = 2\pi h$ . The *source data* are the von Mangoldt atoms  $(u_q, w_q) = (\log q, \Lambda(q)/\sqrt{q})$  over the prime powers  $q \leq h$  — the only arithmetic input anywhere. The *wall matrix*  $M_h = M_{\text{pole}} + M_{\text{arch}} - M_{\text{prime}}$  is the  $K \times K$  matrix of the Weil-explicit-formula quadratic form in the normalised even trigonometric mode basis  $\{\cos(\omega_k u)\}$  on the window  $[0, 2a]$ :  $M_{\text{pole}}$  is the rank-one pole block,  $M_{\text{arch}}$  the archimedean ( $\Gamma$ -term) block, and  $M_{\text{prime}}$  the *only* atom-dependent block — *linear* in exactly two atom transforms per mode,  $P_k = \sum_q w_q \sin(\omega_k u_q)$  off-diagonal and  $\tilde{P}_k = \sum_q w_q [(a - u_q/2) \cos(\omega_k u_q) - \sin(\omega_k u_q)/(2\omega_k)]$  on the diagonal (the **v935** per-mode linearity; the exact assembly is frozen in the builder cited in (6) and rebuildable from the explicit formula for this window class). Let  $\tau_h = \lambda_{\min}(M_h)$  and let  $d$  (the *wall argmin ray*) be its unit eigenvector; the *source coefficients*  $c_k$  are the de-normalised components of  $d$ . From these:  $A_0 = \sum_k (-1)^k c_k$ ,  $A_2 = \sum_{k \geq 1} (-1)^k c_k b_k$ ,  $y_t = |A_2/A_0|$ ,  $B_1 = \sum_{k \geq 1} b_k = (\pi/a)^2 (K-1)K(2K-1)/6$ ,  $\kappa = B_1/y_t$ ; the rational function

$$F(y) = c_0 + \sum_{k \geq 1} (-1)^k c_k \frac{y}{y - b_k},$$

whose numerator  $N(y)$  (degree  $K-1$ , leading coefficient  $A_0$ ) is the *census polynomial*; its roots  $y_j$  are *the census*. The moment–Laurent normal form is  $(y/y_t) F(y)/A_0 = \Phi(z) = z - 1 + \sum_{m \geq 1} J_{m+1} z^{-m}$  at  $z = y/y_t$  (**v924**). The *jet mass* at rung  $h$  is

$$\delta_h = \frac{\sum_{\gamma} \sin^2(a\gamma) F(\gamma^2)^2 / \gamma^2}{A_0^2 \sum_{\gamma} \sin^2(a\gamma) / \gamma^2},$$

the sum over verified ordinates  $\gamma \in (T_z + 6, \gamma_{7000}]$  (caches in (6)); the step-A identity  $\delta_h = |J_h|_{G_h}^2$  with the jet vector  $J = d/A_0$  and  $G_h$  the Gram of  $\psi_k(\gamma) = \gamma^2/(\gamma^2 - b_k)$  under the weight  $\sin^2(a\gamma)/\gamma^2$  is *exact*, remainder 0 (**v939**). The three per-rung hypotheses:

- **H1**( $h; c^*$ ): no census root with  $\text{Re } y \geq c^* y_t$ , with  $c^* \in [1.10, 1.15]$  (grid values, flat over all built rungs). H1 is *certified source-pure* per rung ( $h \leq 26$ ): one exact evaluation of the monotone envelope,  $\text{ENVJ}(c^* y_t) < |A_0|$  — no cache, no zeros (**v931**).

- **H2( $h$ ):** the census is complete-real and nonnegative (all  $K-1$  roots real,  $\geq 0$ ). Landed at all 25 built rungs including the degree-116 census at  $h = 28$ ; finite classical per rung, hard  $h \leq 13$  (v931).
- **H3( $h$ ):**  $y_t \leq 0.155 T_z^4$ . Source-pure certificate, certified 26/26 with margin  $\geq 1.99$ , *refutable* (an explicit inflation witness breaks it) and not norm-continuous (v932).

**(2) The three minimal open statements.** These are the programme's terminal residue in the canonical form (note DII of the discovery record), stated as clean problems about the finite family above.

- **(R1) The triple.** *Prove  $\{H1 \wedge H2 \wedge H3\}$ -cofinal: for every sufficiently large  $k$  there is a rung  $h$  in the dyadic block  $(2^k, 2^{k+1}]$  at which  $H1$ ,  $H2$  and  $H3$  hold simultaneously — one rung per block, all three at the same  $h$ .* Only this cofinal form is asked; the limsup form is available only modulo the measured quantifier defect  $D = 0.0042$  (block-average + pigeonhole is proven exact yet subsidy-empty, v933). Everything downstream of the triple is finite algebra or classical-per-census: the product floor and its assembly are theorems conditional *exactly* on it (v931).
- **(R2) The loop characterization.** *The census-forall- $k$  statement — the same census schedule quantified over all blocks at once — is identified as RH-equivalent-in-currency by four independent machine-detected proof-route cycles:* census universalization == the RH loop (v928); the pinning-supply route == the same loop's costume (v929); BANDPOS on the unbounded ladder  $\iff$  RH (v930); and the de Bruijn–Newman backward-persistence leg == the  $[\Lambda \leq 0]$  class (v938). Stated honestly: within the corpus vocabulary this statement and RH price each other; no proof of equivalence in the literature sense is claimed. The external value is either a genuinely *new independent sign source* for it, or a *proof that the loop is essential* (that census-forall- $k$  is RH-equivalent), which would close the programme's hardness question honestly.
- **(R3) The H-pin.** *The one  $\lambda$ -uniform edge of the counting pair  $\{L1, WPD\}$  (v937):* L1 = TAIL is proven and the H-pin is open; the weak positional dominance  $WPD(a < \gamma_1^2)$  follows from the H-pin for every integer rung  $x \geq 3$ . The pin splits as  $\Omega\text{-a} \oplus \Omega\text{-b} \oplus \text{H123M}$  counting:  $\Omega\text{-a}$  (EPS-LOCK) is the variational quotient  $\bar{\varepsilon} \leq \text{poly}(x) |A_0| \sqrt{G}$  — numerator trial-priceable, denominator floor == the flagged  $A_0$ -triangle loop;  $\Omega\text{-b}$  (SPACING-REMAINDER) is the QSUBGAP kernel. The exact place the floor escapes linearity is machine-exhibited at *two* maps (v937): the eigenvector map ( $\lambda_{\min}$  is not additive —  $\lambda_{\min}(A+B) = -\sqrt{2} \neq \lambda_{\min}(A) + \lambda_{\min}(B) = -2$  on an exact  $2 \times 2$  instance) and the census-root map (not affine,  $d^2 y^* / dw_1^2 \neq 0$  symbolically); the bridge atoms are  $c$ -free while the pin floor  $m_j = 4|A_0 \sin(a\mu_j)| \text{PR}_j$  is  $c$ -loaded — the floor lives on the vector side of the mode-level bridge, and the conditioning wall of v935 applies verbatim.

**(3) What is proven and citable (complete typed hypothesis list).** PROVEN-FINITE = machine-checked in the suite (module cited); CITED-CLASSICAL = a published theorem, used as stated, never re-proved; OPEN = the targets themselves.

| Type            | Hypothesis                                                                                                                                                                                                                                                         | Witness       |
|-----------------|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|---------------|
| PROVEN-FINITE   | Theorem PF: the product floor $F/A_0 \geq (1 - c^*/z)^{(1+\kappa)/c^*}$ , exponent = the trace, conditional <i>exactly</i> on {H1 + H2 per rung}; H1 source-pure (one envelope evaluation per rung); witness modus tollens                                         | v931          |
| PROVEN-FINITE   | the rate dictionary $a = p/2 - 1$ proven at $p = 3, 4, 5$ ; TR-CAP proven <i>and</i> typed circular-for-TOPROOT; the H3 certificate (26/26, margin $\geq 1.99$ , refutable)                                                                                        | v932          |
| PROVEN-FINITE   | block-average + pigeonhole exact yet subsidy-empty ( $D = 0.004183$ ); the witness boundary a theorem boundary, DK-scoped                                                                                                                                          | v933          |
| PROVEN-FINITE   | the unconditional Landau/Gonek pricing: Landau 1912 + Gonek 1985/1993 identified unconditional; Gonek 1984 machine-checked <i>not an ancestor</i> (hidden loop flagged, never consumed); envelope closed form sympy-exact, census-depth-insensitive                | v934          |
| PROVEN-FINITE   | the mode-level Landau bridge: the first linear per-mode classical identity ( $M_{\text{prime}}$ linear in two atom transforms per mode); the rate dictionary a proven bijection (transports every $\theta_\infty$ , pins none)                                     | v935          |
| PROVEN-FINITE   | the tau-free DK exclusion theorem: $\sin \theta \leq (r_0/\tau)/(FG - a)$ with $\tau$ cancelling, proven $h \leq 13$ at the $10^{-25}$ bar + the exact required-bar schedule $b^*(h)$ ; the BA3 scale-gauge theorem                                                | v936          |
| PROVEN-FINITE   | WPD's $\lambda$ -uniform content == the H-pin edge (AND-fire at the declared edges); the restatement licenses; the two linearity-escape exhibits                                                                                                                   | v937          |
| PROVEN-FINITE   | the exact census semigroup: the de Bruijn–Newman flow on the derived census polynomial is $e^{-tT}$ , $T = 4y \partial^2 + 2 \partial$ , closed form, trace law at measured deviation 0.0; the factor-4 normalisation lemma; the dictionary threefold-inequivalent | v938          |
| PROVEN-FINITE   | step A exact: $\delta_h =  J_h _{G_h}^2$ with remainder 0, generic; the cluster normal form (basis-change determinant = the spacing product); Montgomery–Vaughan carried at $3.15\pi$ with exact margin 1/21                                                       | v939          |
| PROVEN-FINITE   | the sub-dof composition chain typed exactly (four exact legs; the floor <i>value</i> is the measured leg); the complete unconditional lower-class pricing table                                                                                                    | v940          |
| PROVEN-FINITE   | the Gram is source-free (proven three ways); the ray-to-law map projective-linear and $A_0$ -free ( $s_i A_0$ linear in $d$ ); the one shared nonlinear link is the wall eigenmap                                                                                  | v941          |
| CITED-CLASSICAL | Landau 1912 (the classical zero-counting/sum form over all nontrivial zeros, unconditional)                                                                                                                                                                        | v934          |
| CITED-CLASSICAL | Gonek 1985/1993 (unconditional exponential-sum bounds; the RH-conditional Gonek 1984 family is <i>flagged</i> , never consumed)                                                                                                                                    | v934          |
| CITED-CLASSICAL | Montgomery–Vaughan 1974 (the generalized Hilbert inequality; killed at linkage 1, carrying at $3.15\pi$ )                                                                                                                                                          | v939          |
| CITED-CLASSICAL | Ortega-Cerdà–Seip 2002 (Fourier frames, Ann. of Math. <b>155</b> , 789–806: the Carleson box condition is equivalent to the <i>upper</i> inequality alone — wrong direction for a floor)                                                                           | v940          |
| CITED-CLASSICAL | Nazarov–Turán 1993 (explicit unconditional lower bound, exponential in the term count — the classical mirror of the measured occupancy law)                                                                                                                        | v940          |
| CITED-CLASSICAL | Rodgers–Tao 2020 ( $\Lambda \geq 0$ ), consumed as a <i>cited ceiling</i> only — its proof consumes pair correlation, which is exactly the flagged loop class                                                                                                      | v938          |
| CITED-CLASSICAL | Polymath 15 2019 ( $\Lambda \leq 0.22$ ), consumed as a <i>cited ceiling</i> only                                                                                                                                                                                  | v938          |
| OPEN            | <b>R1–R3 above</b> — the cofinal triple, the loop characterization, the H-pin                                                                                                                                                                                      | this contract |

(4) **What is measured (honest).** Measured legs are carried as measurements, never as theorems: the unscoped sub-block floor holds with margins  $\geq 1.405$  at all built rungs and the source-only Bertrand-cofinal selector ( $\hat{h}_k = 7/13/31/\dots$ ) clears the SF2 demand with margins 2.14/3.05

(v939/v940); the alignment law is a frozen measured *triple* — the half-power envelope ( $\alpha$  inside the *descriptive* band (0.38, 0.54) at all 14 rungs, never a prediction), the quarter-depth wall lock ( $\mu_m = 0.2372 \log_{10} g_{\min}$ ,  $R^2$  0.9916) and the deep excess (5.64–7.58 dex) (v941); the classical-evaluation face  $\theta_\infty$  sits in the measured band [0.0766, 0.0977] and is typed OPEN-NONPERTURBATIVE-VARIATIONAL — deeper census improves the data  $\sim 1/T$  but provably never the limit (v933/v935); and the full-space refuter ladder is *always cited paired*: the measured first-order minimum sits 30–33 orders above the  $10^{-25}$  bar, while the *proven* schedule is  $b^*(h) = 23.8/21.7/15.7/4.4$  orders at  $h = 4/5/8/13$  and *none* at  $h \geq 16$  (v936). The quartic growth law behind H3 stays MEASURED (v926); the quantifier defect  $D = 0.004183$  is a measurement.

**(5) The kill atlas (so external solvers do not rebuild corpses).** The campaign’s obstruction record is machine-checked: an 11-node implication DAG with exactly one open edge and a candidate map complete over the corpus vocabulary — 24 candidate sign sources, 0 pass (tfpt\_prime\_front, kill-atlas section; sandbox, typed editorial). The named dead route classes, each with its one-line death mechanism:

| Route class                                      | Death mechanism                                                                                                                                                                 | Round / witness            |
|--------------------------------------------------|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|----------------------------|
| pointwise symbol minorants                       | the decisive eigenvalue is a smooth-direction cancellation invisible to any pointwise symbol minorant                                                                           | diary, wall-operator arc   |
| $\ell^1$ / absolute-value majorants              | the target is a five-order cancellation of three $\sim 200$ -sized pieces; the $\ell^1$ majorant overshoots by $10^3$ – $10^6$                                                  | diary, bilinear arc (T170) |
| phase-discarding additive source forms           | the source-only additive (Weyl-split) route is vacuous by 10.9–54.6 dex: the cancellation <i>is</i> the arithmetic                                                              | r166, v927                 |
| Vieta pinch / power sums                         | $1/A_0$ appears in every symmetric function — the resummation never decouples from the Connes scale                                                                             | r172, v932                 |
| moment / trace caps                              | TR-CAP proven yet circular-for-TOPROOT; finite- $P$ moment models undershoot by 1.4–4.5 dex (rate-blind trace-cap vacuity)                                                      | r162/172, v926/v932        |
| Riccati / half-gap induction, manifold transport | transport relabels, never reduces: frame legs exact, source legs non-transporting, the level is cheaper than the increment — the infinite quantifier does not fall by induction | r177, v936                 |
| de Bruijn–Newman pinch                           | theta-flow-inert ( $\geq 4021\times$ the unconditional window); wall fragility 8–51 orders below the window; the only closing leg is the flagged $[\Lambda \leq 0]$ class       | r179, v938                 |
| global CBJ frame inequality                      | the dof kill ( $n = 474$ coefficients vs Landau time-bandwidth rank 128), seam coupling 0.994, occupancy collapse                                                               | r180, v939                 |
| sub-dof scoping                                  | the jet mass leaves the well-conditioned cell at +0.999 dex per rung; the scoped margin dies by $\sim 125\times$ at $\hat{h}(B3) = 13$                                          | r181, v940                 |
| Carleson embedding as a floor                    | wrong direction: the box condition is equivalent to the <i>upper</i> inequality alone; every unconditional lower class misses for a named reason                                | r181, v940                 |
| naive floor-power composition                    | overshoots the honest composed floor by 3.41/10.14 dex; the level needs a per-mode lower bound on the known variational-quotient wall                                           | r182, v941                 |

**(6) Available finite models (runnable).** All models live in the repository and run standalone (Python 3, mpmath/sympy/numpy; python verification/run\_all.py runs the whole suite):

- [verification/v931\_jetmass\_floor\_theorems.py] — theorem PF, H1/H2, the assembled floor
- [verification/v932\_toproot\_theta\_statement.py] — TOPROOT, the rate dictionary, H3
- [verification/v933\_h3\_cofinal\_adjudication.py] — the cofinal adjudication,  $D$ ,  $\theta_\infty$
- [verification/v934\_gonek\_pricing\_unconditional.py] — the unconditional classical pricing
- [verification/v935\_thetainf\_landau\_bridge.py] — the mode-level Landau bridge

- `[verification/v936_manifold_invariance_exclusion.py]` — the DK exclusion, the refuter ladder
- `[verification/v937_l1wpd_closure_reduction.py]` — the H-pin split and linearity exhibits
- `[verification/v938_dbn_heatflow_census.py]` — the exact census semigroup
- `[verification/v939_cbj_frame_adjudication.py]` — step A, the Gram, the selector
- `[verification/v940_cbj_subdof_blockfloor.py]` — the sub-dof kill, the lower-class table
- `[verification/v941_alignment_law_factorization.py]` — the alignment-law factorization

The substrate modules `v922–v930` carry the spacing/jet sum rules, the razor enclosure, the moment–Laurent ladder and the block-average/transfer/sigma-floor algebra. The frozen discovery probes (deterministic, SPEC-hashed, in `experiments/tfpt-discovery/`):

- `jetmass_floor_probe.py` (r171), `toproot_theta_probe.py` (r172)
- `h3_cofinal_probe.py` (r173), `gonek_pricing_probe.py` (r174)
- `thetainf_pin_probe.py` (r175), `manifold_invariance_probe.py` (r177)
- `l1wpd_closure_probe.py` (r178), `dbn_heatflow_probe.py` (r179)
- `cbj_frame_probe.py` (r180), `cbj_subdof_probe.py` (r181)
- `alignment_law_probe.py` (r182)

In the same directory: the frozen instrument builder `radius4_an_probe.py` (`build_cell`; worlds `MAIN/SMOOTH/SCRARITH/EPSTEIN`) and the verified ordinate caches `verified_zeros_big.npy` and `verified_zeros_n7000.npy` (with `_meta.json` sidecars; the big cache holds  $2 \times 10^7$  verified ordinates and tops out at  $T(2 \times 10^7) = 9\,499\,220.48$ ). The verified height horizon is cited as PT21 ( $T_{PT} = 3\,000\,175\,332\,800$ ) with HSW22 Cor. 1.2 as the envelope form.

**(7) Formalization surface (honest).** The only prime-front Lean 4 content is exact finite algebra: the spacing product for generic finite  $K$  over an arbitrary field, the jet sum rules, and the cofinal-wall composition/predefinition audits (`TfptCarrier/SpacingProduct.lean`, `JetSumRules.lean`, `CofinalWeil.lean`, `CofinalEnvelope.lean`, `CofinalPredefinition.lean`, `WallCofinalComposition.lean`, `WallCertifiedHead.lean`; lake build clean, no sorry). **Nothing of R1–R3 is formalised:** H1/H2/H3, the census objects,  $\delta_h$ , the product floor and the pin have *zero* Lean content — no analysis, no zeros, no operator statements. An external solver should treat the Lean layer as contributing only the finite spacing/jet algebra audit.

**(8) Must-fail controls.** A submitted solution is accepted only if it respects the following (each already realised as a machine-checked negative control in the probes above):

- **MF-P1 (Epstein).** Any proposed sign source or positivity argument must fail, or demonstrably lose its constant, on the EPSTEIN world (the  $x^2 + 5y^2$  lattice atoms: an  $L$ -function class with off-line zeros). An argument that goes through unchanged there proves too much.
- **MF-P2 (Scramble).** It must fail on SCRARITH (the golden-ratio permutation of the atom weights on the same support): the law must consume the joint sign-and-magnitude structure of the source, not its histogram.
- **MF-P3 (Smooth).** It must fail on the atom-free SMOOTH world: the conclusion must be arithmetic, not archimedean bookkeeping.
- **MF-P4 (no-zero test).** The proof must consume *no* zero tables: machine ancestry of the delivered chain must be zero-free (the G44-class no-zero gate). The caches of (6) are for *verification* of finite claims, never inputs to a submitted proof.
- **MF-P5 (predefinition test).** Clusters, weights and selectors must be predefined and sealed *before* any sign evaluation (the sealed-pre-ward discipline of r180); no sign-mined selection.

- **MF-P6 (loop firewall).** The proof must not consume the four machine-flagged loop classes: census-forall- $k$  (R2 itself); the  $A_0$ -triangle; zero-verification-as-hypothesis (including the RH-conditional Gonek 1984 family); and the RH-conditional second moments (Montgomery pair correlation / Goldston–Montgomery). Consuming any of them re-imports the hypothesis.

**(9) Acceptance tests.** An external result counts against this contract exactly when it delivers one of:

- **(a) an explicit all- $h$ , all- $a$  bound** — a  $\lambda$ -uniform inequality (over the whole rung ladder, uniform in the window parameter) that implies R1 or the H-pin edge of R3 under the hypothesis list of (3), with any additional hypothesis flagged explicitly and the must-fail controls of (8) verifiably respected;
- **(b) a genuinely new sign source** — one that *separates the worlds and independently orients* (the frozen gate the 24 corpus candidates all fail): it must break on Epstein/Scramble/Smooth while carrying an unconditional orientation on the true source; or
- **(c) a proof that one of R1–R3 is RH-equivalent** — which would close the programme’s hardness question honestly and is recorded as exactly that, never as progress toward RH.

Partial results that close a single leg (e.g. the  $\Omega$ -b kernel alone, or one dyadic block class) narrow the target and are recorded as such. Any marker move happens in `status_ledger.csv` by the house process, never in this document. **No statement about the truth, provability or falsity of the Riemann Hypothesis is made here, and none is implied by the acceptance of a submission.**

**The local-Clifford kill and the mod-8 completion (HECKE.LOCAL.CLIFFORD.01).** *Registered* 2026-08-06, the analytic ordering *dead* as frozen (`[verification/v806_hecke_local_clifford.py]`, 16/18 with the two FAILs the preregistered decider gates, controls 4/4, LOCAL-CLIFFORD-BREAKS; one declared bar re-freeze). All four  $p \bmod 8$  class sectors fail PSD (signed class indicators carry no completed  $L$ -function; the trivial mod-8 component imprimitive); the margin ordering does *not* track the proven activation depths (Spearman +0.40); class 1 the unique positivity-completable class (Euler repair exact). The arithmetic multirate theorem untouched; the Clifford-activation reading decorative analytically. *Positive side-result:*  $\chi_8/\chi_{-8}$  are two new PSD instances of the  $\Gamma_{\mathbb{R}}$  conductor rule (exact parities) — the rule map covers the full mod-8 dual group. *Next falsifier:* primitive sector data, not class indicators. No marker move; no RH statement.

A third hand-off document in the same format, **PRIME.RESIDUE.EXTERNAL.01**, packages the prime-front’s terminal *open statements* for external specialists in analytic number theory — it offers open problems only, asserts no progress toward the Riemann Hypothesis in either direction, and moves nothing.

*Separation (binding).* RH and the **PRIME.\*** programmes stay separate mathematical programmes: a green RH probe would be neither necessary nor sufficient for the 4D physics, and this architecture yields no RH statement. Status: programme section; every constructive contract in it stays **[O]**; the only verified piece is the conditional selector below.

; the **v817** packet numerator equals the code Walsh transform at every level  $\leq 16$ , and  $\hat{m}_2(n) = -1/15$  holds as an integer identity at every odd  $n \leq 16500$  — the per-prime measurement is the prime restriction of a character theorem, typed spectral-only