

The window quadratic form of the TFPT prime-front program is a Galerkin discretization of Suzuki’s localized Weil operator

Stefan Hamann*

August 2, 2026

Abstract

Suzuki (arXiv:2606.09096) realizes the Weil quadratic form localized to $[-a, a]$ as $Q_W^a(v) = \langle B_a v, v \rangle$ with $B_a = D^* G_a D$ on $H_0^1(-a, a)$, where G_a is the integral operator of the screw function g associated with $\zeta(s)$, and identifies the Connes–Consani–Moscovici operator A_a as the Friedrichs extension of B_a . Independently, the prime-front program of Topological Fixed-Point Theory (TFPT) — a machine-checked verification suite in which every quoted number is printed by a named script — measures positivity of a finite “window” matrix A_h assembled from tent (hat) reads of the Weil explicit-formula measure on a declared lattice. We prove, at the level of measures and with every layer either symbolic or certified to 20–50 digits, that these are the same object: the window matrix is a hat-function Galerkin discretization of B_a , with a *one-scalar* dictionary $c_{\text{TFPT}}(d) = D^{-1}(c_{\text{gal}}(d) - \text{pole}_K(d))$ up to a derived, window-independent second-order law $R(d) \rightarrow 1 + 1/(6d^2)$ and an exact finite-rank pole block. The L_0^2 projection in Suzuki’s construction imposes no condition on the Galerkin coefficients (the mean-zero condition lives on the differentiated objects and is automatic on H_0^1), and the TFPT parity projection selects an odd subspace of the form domain, so the bookkeeping closes trivially. On the declared window ($h = 184$) the two quadratic forms agree on ten odd test vectors to relative deviation $\leq 1.3 \times 10^{-10}$; the dictionary, frozen, transports to fresh windows $h = 285, 540, 997$ with only the explicit lattice constant D changing. No positivity claim and no statement about the Riemann Hypothesis is made: the implication chain beyond the identification (form density, spectral limits, positivity transfer) is stated as open.

MSC 2020. 11M26 (primary); 47B25, 65N30 (secondary).

Keywords. Weil explicit formula; screw function; localized Weil operator; Galerkin discretization; Riemann zeta function; Hurwitz–Lerch zeta function.

1 The two objects

1.1 Suzuki’s localized Weil operator

Let Λ be the von Mangoldt function, ψ the digamma function, $\Phi(z, s, a) = \sum_{n \geq 0} (n + a)^{-s} z^n$ the Hurwitz–Lerch zeta function, and let

$$g(t) = -4(e^{t/2} + e^{-t/2} - 2) + \sum_{n \leq e^{|t|}} \frac{\Lambda(n)}{\sqrt{n}} (|t| - \log n) - \frac{|t|}{2} \left(\psi\left(\frac{1}{4}\right) - \log \pi \right) - \frac{1}{4} \left(\Phi\left(1, 2, \frac{1}{4}\right) - e^{-|t|/2} \Phi\left(e^{-2|t|}, 2, \frac{1}{4}\right) \right) \quad (1)$$

*Working note of the verification program of Topological Fixed-Point Theory (TFPT), <https://fixpoint-theory.com>. Every number quoted in this note is produced by a machine-checked script named in the text (repository file `verification/v563_paper2_readouts.py`, the suite modules of Remark 4, and the discovery probes `experiments/tfpt-discovery/w1_*.py`, `w2_form_density_probe.py`); no number is quoted that was not printed by an executed run.

be the screw function associated with $\zeta(s)$ [1, eq. (1.3)]. The final bracket carries the prefactor $\frac{1}{4}$ on *both* terms; this reading is locked by the paper’s own data (see Remark 3). In particular $g(0) = 0$ and, near the origin [1, Sec. 2.2],

$$g(t) = \frac{1}{2}|t| \log |t| + A|t| + \sum_{n \leq e^{|t|}} \frac{\Lambda(n)}{\sqrt{n}} (|t| - \log n) + O(t^2), \quad A = \frac{1}{2}(\log 2\pi - \psi(2)) = 0.7075 \dots \quad (2)$$

With $(Gu)(x) = \int g(x-y)u(y)dy$, $L_0^2(-a, a) = \{u : \int u = 0\}$, P_a the orthogonal projection onto $L_0^2(-a, a)$, and $G_a = P_a G P_a$, Suzuki defines $B_a = D^* G_a D$ on $\mathcal{D}(B_a) = H_0^1(-a, a)$, where $D = i d/dx$ with Dirichlet boundary conditions, and proves that the self-adjoint operator A_a representing the localized Weil form Q_W^a is the Friedrichs extension of B_a [1, Thm. 1.1]; a first numerical study of the Hilbert space derived from the Weil distribution appears in [2]. For $v \in H_0^1(-a, a)$,

$$Q_W^a(v) = \langle B_a v, v \rangle_{L^2} = \int_{-a}^a \int_{-a}^a g(x-y) v'(x) v'(y) dx dy. \quad (3)$$

1.2 The TFPT window form

The TFPT prime-front program is the part of the TFPT verification suite that measures positivity surfaces of the Weil explicit formula on declared finite windows; every step below is frozen in the verification module `v563.paper2.readouts.py`. It fixes a prime-power zone n_0 , sets $\alpha = \log n_0$, chooses a lattice of M cells of width $D = 2\alpha/M$ (“frame A, $\nu = 4$ ” names the suite’s declared window convention; a reader outside the suite may take the explicit values of α, M, D quoted below as given), and assembles two lag vectors from tent functions $\tau_d(t) = (1 - |t - dD|/D)_+$:

$$c_{\text{ar}}[d] = - \int_0^\infty \tau_d(t) \rho(t) dt \quad (d \geq 1), \quad \rho(t) = \frac{e^{-t/2}}{1 - e^{-2t}}, \quad (4)$$

$$c_{\text{at}}[d] = - \sum_{n \leq e^{2\alpha}} \frac{\Lambda(n)}{\sqrt{n}} (\tau_d(\log n) + \tau_d(-\log n)), \quad (5)$$

with the $d = 0$ cell of (4) regularized by the declared near-cell scheme with δ_0 -weight $-(\gamma + \log \pi)$ and regularizer $e^{-3w/2}$. The window matrix is the odd (parity) compression of the Toeplitz form of $c = c_{\text{ar}} + c_{\text{at}}$,

$$(A_h)_{rs} = c_{|r-s|} - c_{(M-1)-r-s}, \quad r, s = 0, \dots, h-1, \quad h = M/2, \quad (6)$$

and the measured surface is $Q_T(v) = v^\top A_h v$. On the declared derivation window, $h = 184$, $M = 368$, $\alpha = \log 16 = 2.772588722240$, $D = 0.015068416969$, with 70 atoms $\log n \leq 2\alpha$.

Reading a coefficient vector $V \in \mathbb{R}^M$ as the hat expansion $\Phi_V = \sum_j V_j \varphi_j$ on the half-integer lattice $p_j = (j - \frac{M-1}{2})D$, the odd sector $V_{M-1-j} = -V_j$ of (6) is exactly the space of odd piecewise-linear functions in $H_0^1(-\tilde{a}, \tilde{a})$, $\tilde{a} = \alpha + D/2$.

2 The identification

2.1 The projection bookkeeping closes trivially

Lemma 1. (i) For every $v \in H_0^1(-a, a)$ one has $\int_{-a}^a v' = v(a) - v(-a) = 0$; hence $u = Dv \in L_0^2(-a, a)$ automatically, $P_a(Dv) = Dv$, and $\langle G_a Dv, Dw \rangle = \langle G Dv, Dw \rangle$ for all $v, w \in H_0^1(-a, a)$. In particular Suzuki’s L_0^2 constraint imposes no linear condition on the coefficient vectors of any H_0^1 -conforming Galerkin family, and the Krein-kernel replacement $g(x-y) \mapsto$

$g(x-y) - g(x) - g(-y) + g(0)$ leaves the Galerkin matrix unchanged. (ii) The TFPT parity sector consists of odd functions $\Phi_V \in H_0^1(-\tilde{a}, \tilde{a})$ with $\int \Phi_V = 0$: an extra mean-zero condition on the test function itself, selecting a parity subspace of the form domain (compatible with $\lambda_a = \min(\lambda_a^+, \lambda_a^-)$, [1, Sec. 4.5]).

The point of (i) is that the mean-zero condition lives on the *differentiated* objects ($D : H_0^1 \rightarrow L_0^2$ is a bijection, [1, Sec. 8.2]), where it holds identically; the point of (ii) is that oddness is a *v*-side condition Suzuki does not need but which does no harm. Both statements are verified computationally in `w1_theorem_probe.py` (P1.1–P1.3): the Krein correction evaluates to $-2(\int gu)(\int u) + g(0)(\int u)^2$ with $\int u = 0$ exactly on the sector (must-break control: for the non-mean-zero $u \equiv 1$ the two kernels differ by 2.8357).

2.2 Three identities and the theorem

Write g'' for the distributional second derivative and define the *Weil measure*

$$\mathcal{W} := g'' + (e^{t/2} + e^{-t/2}) = \underbrace{\rho(t) dt \ (t \neq 0)}_{\mathcal{W}_{\text{sm}}} + \underbrace{\text{origin block} + \sum_{n \geq 2} \frac{\Lambda(n)}{\sqrt{n}} (\delta_{\log n} + \delta_{-\log n})}_{\mathcal{W}_{\text{at}}}, \quad (7)$$

where the origin block is the Pf/δ_0 content of the $\frac{1}{2}|t| \log |t| + A|t|$ head of (2) (total δ_0 mass $2A + 1 = \log 2\pi + \gamma$ in the Pf scheme, [1, Sec. 2.5]).

Theorem 2 (W1, measure level). *Let g be (1) and $\mathcal{W}$ as in (7). Then, on the declared window ($\alpha = \log 16$, $M = 368$) and portably in D (Sec. 3):*

- (i) $g'' = -(e^{t/2} + e^{-t/2}) + \mathcal{W}$ as tempered distributions; the pole block $-(e^{t/2} + e^{-t/2})$ (the ζ -poles $s = 0, 1$) is exactly finite-rank after parity compression (measured rank 1).
- (ii) **Tent reads.** The TFPT arch lags are the point-route tent reads of $-\mathcal{W}_{\text{sm}}$ at every lag, with no constant and no scalar: $c_{\text{ar}}[d] = -\langle \mathcal{W}_{\text{sm}}, \tau_d \rangle$ (verified at $d = 1, 2, 3, 10, 50$ to relative 3.4×10^{-52}); the candidate origin constant $\kappa := c_{\text{ar}}[0] + \frac{1}{D} [g_{\text{sm}}(-D) - 2g_{\text{sm}}(0) + g_{\text{sm}}(D)]$ vanishes ($|\kappa| < 10^{-53}$; equivalently the closed identity $\frac{2}{D} \int_0^D t \rho dt + 2 \int_D^\infty \rho dt + \frac{f(D) - \Phi(1, 2, 1/4)}{2D} = 0$ with $f(t) = e^{-t/2} \Phi(e^{-2t}, 2, \frac{1}{4})$): the v563 near-cell scheme is Suzuki's origin bookkeeping. The atom lags (5) are the literal tent reads of $-\mathcal{W}_{\text{at}}$ (exact).
- (iii) **B-spline reads.** The hat-Galerkin lags of $B_a = D^* G_a D$, i.e. $c_{\text{gal}}(|i - j|) = \iint g(x - y) \varphi'_i(x) \varphi'_j(y)$, satisfy $c_{\text{gal}}(d) = \text{pole}_K(d) - \langle \mathcal{W}, K_d \rangle$, where K is the hat autocorrelation (cubic B-spline with knot weights $(1, 4, 1)/6$, mass D^2) and

$$\text{pole}_K(d) = \frac{32 \cosh(dD/2) (e^{D/2} + e^{-D/2} - 2)^2}{D^2}.$$

The smooth density read $-\langle \rho, K_d \rangle$ is exact for $d \geq 3$ (verified to 3.7×10^{-31} at $d = 5$); the atom layer has the closed B-spline form $-\sum_n \frac{\Lambda(n)}{\sqrt{n}} [K(u_n - dD) + K(u_n + dD)]$ (verified against kink-split double integrals to 4.2×10^{-17}); the boundary cells carry three exact computable constants $\tilde{B}(0) = 1.295327231917$, $\tilde{B}(1) = -0.671894595706$, $\tilde{B}(2) = 0.017023960540$.

- (iv) **One-scalar dictionary.** Consequently

$$c_{\text{TFPT}}[d] = \frac{1}{D} (c_{\text{gal}}(d) - \text{pole}_K(d))$$

up to (a) the derived second-order law $R(d) = 1 + \frac{D^2}{12} \frac{\rho''}{\rho}(dD) + O(D^4) \rightarrow 1 + \frac{1}{6d^2}$ (window-independent) on the smooth layer, (b) the measure-level re-binnings (atoms: tent vs. B-spline reads of the same data; smooth near cells: exact read ratios of the same density), and (c) the exact constants of (ii)–(iii). On the common odd sector the two quadratic forms agree on ten 16-mode test vectors to per-vector relative deviation $\leq 1.28 \times 10^{-10}$ (median 3.0×10^{-11} ; 16-mode block operator-norm ratio 1.34×10^{-11}), and the identification is sign-compatible with positivity: the measured positive-definiteness of A_h and $Q_W^a \geq 0$ sit on the same side.

Proof of the smooth structure in (i) (five lines). Term-wise, $e^{-t/2} \Phi(e^{-2t}, 2, \frac{1}{4}) = \sum_{n \geq 0} (n + \frac{1}{4})^{-2} e^{-(2n+\frac{1}{2})t}$ and $\frac{d^2}{dt^2} e^{-(2n+\frac{1}{2})t} = (2n + \frac{1}{2})^2 e^{-(2n+\frac{1}{2})t} = 4(n + \frac{1}{4})^2 e^{-(2n+\frac{1}{2})t}$, so

$$\frac{d^2}{dt^2} \left[e^{-t/2} \Phi(e^{-2t}, 2, \frac{1}{4}) \right] = 4 \sum_{n \geq 0} e^{-(2n+\frac{1}{2})t} = \frac{4e^{-t/2}}{1 - e^{-2t}} = 4\rho(t) \quad (t > 0) : \quad (8)$$

the Hurwitz–Lerch block is a geometric series in disguise. With the prefactor $\frac{1}{4}$ of (1) and $[-4(e^{t/2} + e^{-t/2} - 2)]'' = -(e^{t/2} + e^{-t/2})$, off the atoms $g''(t) = -(e^{t/2} + e^{-t/2}) + \rho(t)$. $\square$

The δ_0 closure in (ii) rests on the duplication identity $\psi(\frac{1}{4}) = -\gamma - 3\log 2 - \pi/2$ and the scheme conversion $k = \int_0^\infty \rho(w)(1 - e^{-3w/2}) dw = \frac{1}{2}(\psi(1) - \psi(\frac{1}{4})) = \frac{3\log 2 + \pi/2}{2}$, whence $-(\gamma + \log \pi) - 2k = \psi(\frac{1}{4}) - \log \pi$: the two δ_0 conventions are one object in one scheme (`w1_boundary_closure_probe.py`, B1–B2, symbolic + 30-digit certificates). All quantitative claims of Theorem 2 are checks in `w1_theorem_probe.py` (C0, P1–P3; 11/11 pass, verdict W1-THEOREM).

Remark 3 (Convention lock and verbatim transfer). The reading of the final bracket of (1) is fixed by Suzuki’s own Section 2.2: the printed constant $A = \frac{1}{2}(\log 2\pi - \psi(2)) = 0.707546\dots$ of (2) holds *only* if the Lerch term carries the coefficient $+\frac{1}{4}$ (machine check: with coefficient $+\frac{1}{4}$ the origin read reproduces A to 1.2×10^{-9} at $t = 10^{-8}$; with coefficient -1 it is off by ≈ 56 and $g(0) = -\frac{5}{4}\Phi(1, 2, \frac{1}{4}) = -21.4967 \neq 0$). An earlier tranche of machine verifications (probe files `w1_dictionary_probe.py`, `w1_boundary_closure_probe.py`, `w1_frozen_windows_probe.py`, `w1_matrix_identity_probe.py`) used the kernel $\tilde{g} := g - \frac{5}{4}e^{-|t|/2}\Phi(e^{-2|t|}, 2, \frac{1}{4})$ (Lerch coefficient -1). Since $\tilde{g}_{\text{sm}}'' = -4\rho$ while $g_{\text{sm}}'' = +\rho$, one has $c_{\text{gal}}^{\text{sm}}(\tilde{g}) = -4c_{\text{gal}}^{\text{sm}}(g)$ *exactly* (checked to 3.7×10^{-12}), so every measured ratio of that tranche transfers verbatim: $c_{\text{gal}}(\tilde{g})/(-4Dc_{\text{ar}}) \equiv c_{\text{gal}}(g)/(+Dc_{\text{ar}})$. All numbers quoted below from those probes are therefore valid for the true dictionary; only the labels change ($-4D$ two-layer $\mapsto +D$ one-scalar, and the sign of the identification becomes positivity-compatible).

Remark 4 (Suite provenance). The probe files quoted in this note have been promoted to the permanent, machine-checked verification suite of the TFPT repository [5] as the modules `v630` (first contact: the atom layer is literally Suzuki’s prime measure), `v631` (per-lag dictionary), `v640` (boundary closure), `v641` (frozen-window portability), `v642` (matrix identity), `v643` (the measure-level theorem of Sec. 2, including the convention lock of Remark 3 as check C0.1), and `v644` (the W2 slices of Sec. 4). The `v631/v640–v642` chain was written in the $\tilde{g}$ -normalization of Remark 3; the convention was corrected the same day, see the dated erratum in the suite (erratum blocks in `v631` and `v640–v642`, machine evidence `v643` check C0.1). By the exact transfer identity above, every measured number of that chain remains valid verbatim under the corrected convention.

2.3 The second-order law

The tent and B-spline reads of a smooth density f differ by derived moment brackets (all coefficients symbolic): $\int \tau_0 t^{2k} = D, \frac{D^3}{6}, \frac{D^5}{15}, \frac{D^7}{28}, \frac{D^9}{45}$ and $\int K t^{2k} = D^2, \frac{D^4}{3}, \frac{3D^6}{10}, \frac{17D^8}{42}, \frac{31D^{10}}{45}$

for $2k = 0, \dots, 8$, giving

$$R(d) := \frac{\langle \rho, K_d \rangle}{D \langle \rho, \tau_d \rangle} = 1 + \frac{D^2}{12} \frac{\rho''}{\rho}(dD) + O(D^4) \longrightarrow 1 + \frac{1}{6d^2} \quad (\text{window-independent}). \quad (9)$$

On the declared window the exact ratios at $d = 5, 9, 16$ are 1.006766, 1.001959, 1.000585; the D^6 -truncated bracket explains the excess $R - 1$ to 3×10^{-4} , and the derived D^8 order accounts for 77%–99% of the residual at $d = 3, \dots, 12$ (`w1_theorem_probe.py` P2.5): the remainder of the dictionary is typed bracket truncation, not an unexplained conversion.

3 Numerical verification

3.1 Portability across windows

The dictionary was derived on $h = 184$ and then *frozen* (all formulas closed in D , no knob) and executed unchanged on three fresh frame-A windows (`w1_frozen_windows_probe.py`, pre-registered kill test, 10/10 pass, verdict W1-DICTIONARY-PORTABLE; ratios transfer to the true dictionary by Remark 3):

window	M	D	smooth ratio $c_{\text{gal}}^{\text{sm}}/(Dc_{\text{ar}})$ at $d = 3, 5, 9, 16$					$ R(16) - 1 $	atoms
$h = 285$	570	0.01731394	1.0211	1.0067	1.0019	1.0006	5.8×10^{-4}	1.000000	D^2
$h = 540$	1080	0.00936342	1.0213	1.0069	1.0020	1.0006	6.1×10^{-4}	1.000000	D^2
$h = 997$	1994	0.00603181	1.0214	1.0069	1.0020	1.0006	6.2×10^{-4}	1.000000	D^2

Here “atoms” is the per-atom Galerkin mass over the TFPT tent mass ($\log 2$ and $\log 3$ agree to $< 10^{-14}$): the atomic re-binning constant is exactly the B-spline/tent mass ratio D^2 , i.e. the same $+1/D$ per-lag scalar as the smooth layer. The moment prediction (9) explains the excess to 0.6% on all three windows; only D changes.

3.2 Form equality on the common odd sector

On $h = 184$, with ten odd 16-mode test vectors (fixed seed) and the dictionary applied in stages (`w1_theorem_probe.py` P3.1), the 16-mode block operator-norm ratios $\|T(A_T - A_S)T'\|/\|TA_T T'\|$ are

stage	norm ratio
one-scalar $1/D$	4.37×10^{-3}
+ $\tilde{B}(0.2)$ (exact boundary)	3.00×10^{-3}
+ moment law (9)	2.72×10^{-3}
+ measure-level atom re-binning	4.62×10^{-6}
+ measure-level smooth re-binning (near cells)	1.34×10^{-11}

and the per-vector deviations $|Q_T - Q_S|/|Q_T|$ reach at most 1.28×10^{-10} (median 3.0×10^{-11}). With the derived moment law alone (no exact near-cell ratios) the worst vector sits at 1.02×10^{-4} , entirely the typed D^8 truncation at $d = 3, 4$ hitting near-cancellation denominators. Nothing is fitted at any stage.

4 Scope: what is and is not claimed

This note establishes an *identification of finite-dimensional quadratic forms* and its exact measure-level skeleton. It does *not* establish positivity of anything, and it makes no statement about the Riemann Hypothesis. The honest implication map is:

step	statement	status
W1	window form = Galerkin read of B_a (explicit dictionary)	closed at exact/certified level (this note; <code>w1_theorem_probe.py</code> , 11/11).
W2	nested window spaces are form-dense in the odd form sector of A_a ; window spectra converge	started (<code>w2_form_density_probe.py</code>): FEM density with classical rates verified (H^1 rate 1.000, L^2 rate 2.000 over $M = 92\text{--}736$); Rayleigh–Ritz eigenvalue convergence and λ_a -profile measured; Mosco/norm-resolvent convergence <i>open</i> .
W3	control in a : continuity/limits of λ_a , the $a \rightarrow \infty$ strong-resolvent limit of [1, Sec. 7]	<i>open</i> ; [1, Thm. 1.3] gives continuity of λ_a (no monotonicity claimed there).
W4	positivity transfer: window positivity measurements $\Rightarrow$ Weil positivity	<i>not claimed</i> ; would require W2 + W3 + uniformity in the window family. Weil positivity for all a is equivalent to RH [3, 4]; nothing here touches that equivalence.

The W2 slices measured in `w2_form_density_probe.py` (suite module v644, verdict w2-SLICES-PASS, 7/7) are, beyond the FEM rates of the table: Rayleigh–Ritz on the nested dyadic refinements $M = 92/184/368/736$ at fixed $a = \log 16$ gives lowest-three eigenvalues decreasing monotonically from above in all six sequences (full and odd sector), with Cauchy gaps shrinking at every step and the parity identity $\lambda_{\text{full}} = \min(\lambda_{\text{even}}, \lambda_{\text{odd}})$ holding to 10^{-15} per refinement. The λ_a scan across five frame-A windows $a = 2.77, \dots, 6.01$ at fixed $M = 368$ gives λ_a between $+10^{-12}$ and $+2 \times 10^{-10}$ with per-window discretization gaps $\sim 10^{-9}$: the honest reading is $\lambda_a = 0^+$ within $\sim 10^{-9}$ — consistent with the continuity statement of [1, Thm. 1.3] and *not* a sign statement — with ground-state parity even at every measured a (reported; [1, Thm. 1.4] covers only small a).

Within W1 itself the remaining declared freedoms are exactly: the finite-rank pole block (rank 1 after parity compression), the typed bracket truncation of (9) beyond D^8 , and the floating-point/quadrature certification level stated with each check.

References

- [1] M. Suzuki, *Weil’s quadratic form via the screw function*, arXiv:2606.09096 (2026).
- [2] M. Suzuki, *On the Hilbert space derived from the Weil distribution*, Canadian J. Math. (to appear), arXiv:2607.24830.
- [3] A. Weil, *Sur les “formules explicites” de la théorie des nombres premiers*, Comm. Sém. Math. Univ. Lund (1952), Tome Supplémentaire, 252–265.
- [4] H. Yoshida, *On Hermitian forms attached to zeta functions*, in: Zeta functions in geometry (Tokyo, 1990), Adv. Stud. Pure Math. 21, Kinokuniya, Tokyo (1992), 281–325.
- [5] TFPT repository (<https://fixpoint-theory.com>): verification modules `v563_paper2_readouts.py`, `v630`, `v631`, `v640–v644`, and discovery probes `w1_dictionary_probe.py`, `w1_boundary_closure_probe.py`, `w1_frozen_windows_probe.py`, `w1_matrix_identity_probe.py`, `w1_theorem_probe.py`, `w2_form_density_probe.py` (2026).