

The window quadratic form as an off-line zero detector: a structure theorem with quantified thresholds

Stefan Hamann*

August 2, 2026

Abstract

The finite window quadratic form of the TFPT prime-front program — identified in the companion note [2] as a hat-function Galerkin discretization of Suzuki’s localized Weil operator — is shown to satisfy an exact structure theorem: the deployed odd-Toeplitz form is the Cantoni–Butler odd-sector compression of the full Toeplitz matrix of the Weil lag vector, its generalized spectrum is sandwiched between the periodized-symbol minimum and the discrete-sine-mode reads (the naive “DST-diagonal” reading is *false*; the measured parity defect is $O(1)$), and every mode read is, unconditionally by the Weil explicit formula, a sinc²-damped *alias comb* over the nontrivial zeros plus an explicit nonnegative pole layer, in which a hypothetical off-line zero quadruple at $\frac{1}{2} + \delta + i\gamma$ enters with the amplification $\cosh(\delta u) \leq \cosh(2\alpha\delta)$ per lag u — the form is an off-line zero detector. The detector is validated in two exact laboratories: on the Ihara side (proven Ramanujan and non-Ramanujan graphs) the machinery is positive semidefinite on every proven-true control and breaks on the proven violators with the detection law $K^* \cdot s \approx 2\text{--}3$, and at the Epstein zeta function of $x^2 + 5y^2$ a measured argument-principle census of 26 off-line zeros (12 in $[0.6, 1.4] \times [2, 100]$) predicts the measured form break quantitatively ($\lambda_{\min}$ ratio 0.803 inside the preregistered factor-2 gate). On the deployed 67-window family the measured positivity becomes a certificate with a quantified threshold map $s_{\min}(w, \gamma)$ (median $2\alpha s_{\min} = 1.4\text{--}2.6$ on the resolved band $\gamma \leq \pi/D \leq 868$), and the $C = 1$ contraction constant of the model/real ladder is identified as a quadrature statement: the zero-free density integral of the lock profile predicts the measured constant with median ratio 1.005. The Fejér density identity $s_{\text{tot}} = 2\pi (F_a * dN)$, verified to 2.9×10^{-6} and carrying an unconditional smoothed-density lower bound with $c_0 = 0.306$ (Riemann–von Mangoldt counting plus Trudgian’s $|S(t)|$ bound, ten finite machine-checked certificates), is the lemma connecting all three results. Uniform positivity over the full window family is the conjecture — by Weil’s criterion equivalent to the Riemann Hypothesis [3, 4, 5] — and nothing in this note claims it.

MSC 2020. 11M26 (primary); 47B35, 65F15 (secondary).

Keywords. Weil explicit formula; Toeplitz quadratic form; alias positivity; off-line zeros; Epstein zeta function; Ihara zeta function; Riemann zeta function.

1 Setup

This note is the successor of [2], which identified the TFPT window quadratic form as a Galerkin discretization of Suzuki’s localized Weil operator B_a [1] at the level of measures.

*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 (the discovery probes `experiments/tfpt-discovery/w3.structure_theorem.probe.py`, `fejerdensity-bound.probe.py`, `c1.mechanism.probe.py`, `ihara_ground_truth.probe.py`, `epstein.firewall.probe.py`, all on the frozen window machinery of `verification/v563.paper2.readouts.py`; `zero_gap.theorem.probe.py` for the addendum of Sec. 7; the suite promotions of these probes are listed in Remark 4); no number is quoted that was not printed by an executed run.

We keep its conventions and recall only what is needed. A *frame- A window* is a triple $(\alpha, M = 2h, D = 2\alpha/M)$; the deployed family has 70 windows of which 67 are *complete* ($e^{2\alpha}$ inside the frozen atom table; the three truncated windows $h = 1219, 1292, 1445$ are excluded throughout). The lag vector is $c_d = c_{\text{ar}}(d) + c_{\text{at}}(d)$, $d = 0, \dots, M-1$: c_{ar} is the archimedean layer read by a width- D tent at lag dD , and

$$c_{\text{at}}(d) = - \sum_{n \leq e^{2\alpha}} \frac{\Lambda(n)}{\sqrt{n}} g_d(\log n), \quad g_d(u) = \tau_D(u - dD) + \tau_D(u + dD), \quad (1)$$

with $\tau_D(u) = (1 - |u|/D)_+$. The window matrix is the odd (parity) compression

$$(A_h)_{jk} = c_{|j-k|} - c_{(M-1)-j-k}, \quad j, k = 0, \dots, h-1, \quad (2)$$

and G_h is the analogous compression of the hat mass vector $(g_0, g_1) = (2D/3, D/6)$. All spectral statements below are generalized eigenvalue statements for the pencil (A_h, G_h) ; mode reads are Rayleigh reads $R_A(k) = u_k^\top A_h u_k$ of the discrete sine (DST) vectors $u_k[j] = \sin(\theta_k(j - h + \frac{1}{2}))$, $\theta_k = 2\pi k/M$. Every form and every mode read in this note is assembled from *primes and the digamma function only*; Riemann zeros enter only the verification of identities (through a Turing-certified cache of the first 2000 zeros, live-checked to 2.7×10^{-12}) and the Epstein census is a measured output of the argument principle, never an input table.

Three results are established: **(A)** the structure theorem of Sec. 2 — window positivity is *exactly* alias positivity of the zero symbol, and off-line zeros enter with a quantified cosh gain; **(B)** a two-laboratory validation and a threshold map (Secs. 3–4), typing the measured 67-window positivity as a certificate with quantified reach; **(C)** a quadrature identification of the $C = 1$ contraction constant (Sec. 5), with the Fejér density identity as the connecting lemma (Sec. 6). Section 7 gives the honest scope, the circularity ledger, and the implication map.

2 The structure theorem

All checks of this section are from `w3.structure_theorem_probe.py` (22 checks, 0 failures, verdict `W3ST-STRUCTURE-THEOREM`), run on the three verification windows $h = 142, 859, 1433$ (smallest complete, mid-family, largest complete) with preregistered bars.

2.1 Toeplitz $\Leftrightarrow$ symbol: the sandwich, not the naive diagonalization

With $f \in \mathbb{R}^M$ the odd extension of $x \in \mathbb{R}^h$ ($f_j = x_j$, $f_{M-1-j} = -x_j$) one has exactly $x^\top A_h x = \frac{1}{2} f^\top T_M(c) f$ with $T_M(c)_{ij} = c_{|i-j|}$: every lag $0 \leq d \leq M-1$ is present. T_M is persymmetric, and the reflection sectors are invariant: $\text{spec } T_M = \text{spec } A_{\text{even}} \cup \text{spec } A_{\text{odd}}$ (Cantoni–Butler [6]), verified eigenvalue-by-eigenvalue on all three windows to relative deviation 1.5×10^{-15} .

The honest correction. It is tempting to read the compressed form as “DST-diagonal with symbol eigenvalues”. That statement is **false** for the deployed lag vector: only the reflection-symmetrized part E_h (the $\mathbb{Z}_M$ -circulant compression of $e_d = (c_d + c_{(M-d) \bmod M})/2$) is exactly DST-diagonal, with eigenvalues the alias-symbol samples $\sigma(2\pi k/M)$, $\sigma(\theta) = c_0 + 2 \sum_{d \geq 1} c_d \cos d\theta$ (verified to 3.0×10^{-13}). The parity defect Ω is *measured and not small*: $\|\Omega\|_2/\text{rad} = 0.916, 1.424, 1.551$ on $h = 142, 859, 1433$. The correct finite statement is the two-sided sandwich (Grenander–Szegő congruence [7] plus the exact mode reads):

$$\min_{\theta} \frac{\sigma_A(\theta)}{\sigma_G(\theta)} \leq \lambda_{\min}(A_h, G_h) \leq \min_{1 \leq k \leq h} \frac{R_A(k)}{R_G(k)}, \quad \sigma_G(\theta) = \frac{2D}{3} + \frac{D}{3} \cos \theta > 0, \quad (3)$$

with the exact closed-form mode weights, derived and verified to 9.1×10^{-13} against the deployed code path,

$$R_A(k) = \sum_{d=0}^{M-1} c_d w_d(k), \quad w_d(k) = \left(h - \frac{d}{2}\right) \cos(\theta_k d) + \frac{\sin(\theta_k d)}{2 \sin \theta_k} \quad (1 \leq k < h), \quad (4)$$

$w_0 = h/2$; at Nyquist $k = h$: $w_0 = h$, $w_d = (-1)^d(2h - d)$. The tent factor $h(1 - u/2\alpha)$ at $u = dD$ in (4) is exactly the tent test-pair weight of the continuum symbol of Sec. 6 — the first bridge between the finite window and the Fejér object. The measured sandwich (3) on the three windows:

window	$\min \sigma_A / \sigma_G$	$\lambda_{\min}(A, G)$	$\min_k R_A / R_G$	$\ \Omega\ _2 / \text{rad}$
$h = 142$	-5.775×10^3	$+1.131 \times 10^{-2}$	$+7.225$	0.916
$h = 859$	-1.466×10^5	$+2.632 \times 10^{-3}$	$+18.52$	1.424
$h = 1433$	-4.699×10^5	$+1.516 \times 10^{-3}$	$+30.20$	1.551

The lower bound is valid but far from tight (the periodized symbol dips deeply negative); the informative side is the mode-read upper bound, which the theorem below turns into a zero statement.

2.2 The symbol is the alias density of the zeros

Each g_d of (1) is an admissible even Weil test pair; its transform $h_d(r) = 2 \cos(dDr) D \operatorname{sinc}^2(rD/2)$ is entire of exponential type. The explicit formula ([8, Thm. 5.12]; no Riemann Hypothesis input) gives, for *every* lag,

$$c_d = \sum_{\gamma > 0} h_d(\gamma_\rho) - \text{pole}_d, \quad \text{pole}_d = 2D \cosh\left(\frac{dD}{2}\right) \left(\frac{\sinh(D/4)}{D/4}\right)^2, \quad (5)$$

the sum over all nontrivial zeros $\rho = \beta + i\gamma$ with $\gamma_\rho = (\rho - \frac{1}{2})/i$ (complex if and only if ρ is off the critical line; h_d is entire, so the formula needs no RH). The archimedean layer of (5) is verified against digamma quadrature at lags $d \in \{0, 1, 3, 17, 100, 283\}$ (deviations 7.0×10^{-6} down to 3.2×10^{-9} , within the declared tail-limited bars), and the full identity is verified against the certified zero comb plus the smooth Riemann–von Mangoldt tail at seven lags with printed $S(T)$ -blind residuals $\leq 8.5 \times 10^{-3}$.

Theorem 1 (structure theorem, machine-checked). *Let (α, M, D) be a window of the family and $x \in \mathbb{R}^h$ arbitrary with odd extension f , and set $F_x(\varphi) = \sum_{j=0}^{M-1} f_j e^{ij\varphi}$. Then, unconditionally,*

$$x^\top A_h x = \sum_{\gamma > 0} T_x(\gamma_\rho) + P(x), \quad T_x(r) = D \operatorname{sinc}^2\left(\frac{rD}{2}\right) F_x(Dr) F_x(-Dr), \quad (6)$$

where $P(x) = -\frac{1}{2}(T_x(i/2) + T_x(-i/2)) \geq 0$ is the closed pole layer. On the real axis $T_x(r) = D \operatorname{sinc}^2(rD/2) |F_x(Dr)|^2 \geq 0$ and $|F_x|^2$ is $2\pi/D$ -periodic: T_x is the sinc^2 -damped alias comb of the mode profile. Consequently $A_h \succeq 0$ if and only if the aliased zero side $\sum_{\gamma > 0} T_x(\gamma_\rho) + P(x)$ is nonnegative on the window's test cone — window positivity is alias positivity of the zero symbol. A zero quadruple $\{\rho, 1 - \rho, \bar{\rho}, 1 - \bar{\rho}\}$ with $\delta = \beta - \frac{1}{2} \in (0, \frac{1}{2}]$ at height γ contributes exactly $2 \operatorname{Re} T_x(\gamma + i\delta)$, with per-lag transform

$$h_d(\gamma + i\delta) = 2[\cos(dD\gamma) \cosh(dD\delta) - i \sin(dD\gamma) \sinh(dD\delta)] D \operatorname{sinc}^2\left(\frac{(\gamma + i\delta)D}{2}\right) : \quad (7)$$

the $\cosh(\delta u)$ amplification at lag $u = dD$, up to $\cosh(2\alpha\delta)$ at the deepest lag — the detector gain.

Both directions of the equivalence are trivial given the identity; the nontrivial content is (6) itself. The trigonometric master identity is verified to relative 5.1×10^{-13} (bar 10^{-11}), the pole layer carries the closed square form $T_x(\pm i/2) = -D \operatorname{sinc}^2(iD/4) e^{(M-1)D/2} (\sum_j f_j e^{-jD/2})^2 \leq 0$, and the mode-level reconstruction $R_A(k) = \sum_{\text{comb}} T_k(\gamma) + \text{RvM tail} + P_k$ closes at symbol-scale deviations 0.014/0.0001/0.031 for modes at $t \approx 20/50/100$ (the residual is genuine $S(T)$ -blindness of the smooth tail). The off-line closed form (7) is verified against the entire continuation to 5.8×10^{-16} , its $\delta \rightarrow 0$ limit and the envelope $|\text{OFF}| \leq (1 + e^{2a\delta})/(\pi a(x^2 + \delta^2))$ hold on the declared grids.

3 Two-laboratory validation

A detector claim needs a laboratory where the target provably exists and one where it provably does not. The program runs both; we regard this two-laboratory methodology as a contribution in its own right.

3.1 Ihara: the positive control and the margin calibration

The Ihara zeta function of a finite $(q+1)$ -regular graph satisfies the Riemann Hypothesis analogue if and only if the graph is Ramanujan [14], and both classes have small proven members. `ihara_ground_truth_probe.py` (22 checks, 0 failures, verdict IHARA-GT-CONFIRMED) rebuilds the analogue of the window machinery over the path-count atoms of five cubic graphs. On the three proven Ramanujan controls (Petersen, Heawood, coupled cliques) the window form is positive semidefinite for *every* depth $K = 2..40$ with exact rank saturation $\min(K, F)$ (F the spectral support size). On the two proven violators (the prisms $C_{16} \times K_2$, $C_{24} \times K_2$) the form breaks at $K^* = 15$ and $K^* = 11$. Two calibration lessons transfer:

- (i) *True positivity has no uniform margin.* On the Ramanujan side the margin $\delta(K) = \lambda_{\min}/\lambda_{\max}$ is $O(1)$ only while $K \leq F$ and is *identically zero* beyond (measured $|\lambda_{\min}| \leq 7.5 \times 10^{-14}$, the float floor). The honest positivity target is $\lambda_{\min} \geq 0$, not $\lambda_{\min} \geq \delta > 0$.
- (ii) *The detection law.* With violation exponent $s = \operatorname{arccosh}(\max |\lambda_{\text{nt}}|/2\sqrt{2})$ ($s = 0.117/0.270$ for the two prisms), the measured onset obeys $K^* \cdot s = 1.75$ and 2.97 : a violation of strength s needs window depth $\sim (2-3)/s$ to become visible.

3.2 Epstein: the negative control and the quantitative form-break prediction

The Epstein zeta function $E(s) = \frac{1}{2} \sum'_{(x,y)} (x^2 + 5y^2)^{-s}$ (discriminant -20 , class number 2) has no Euler product and violates RH [13]. The probe rebuilds its arithmetic from lattice data: the genus identity $E = \zeta L_{-20} + L_{-4} L_5$ holds coefficient-wise exactly to $n = 10^5$, and the von Mangoldt analogue Λ_E (the $-E'/E$ coefficients) has 1351 negative sites (first at $n = 36$) and 4849 support points off the prime powers — the Euler structure is measurably gone (`epstein_firewall_probe.py`, 13 checks, 0 failures). Feeding Λ_E atoms into the *unchanged* deployed machinery breaks the form on all three preregistered picks, roughly $9\times$ deeper than the Euler-true twin $\zeta_K = \zeta L_{-20}$ (same completed functional equation, so archimedean and pole layers cancel exactly in the difference):

pick	$\lambda_{\min}(\zeta_K)$	$\lambda_{\min}(\Lambda_E)$	$\lambda_{\min}$ predicted	ratio	band corr.
$h = 277$	-6.43×10^2	-5.90×10^3	-5.41×10^3	0.916	0.830
$h = 454$	-1.31×10^3	-7.90×10^3	-7.56×10^3	0.957	0.813
$h = 1077$	-2.52×10^3	-4.37×10^4	-3.51×10^4	0.803	0.917

The prediction column is built *from the zeros alone*: an argument-principle census (adaptive winding walk with Muller polish, $|E(\rho)| \leq 10^{-8}$ enforced, all boxes resolved) localizes 12 off-line zeros in $[0.6, 1.4] \times [2, 100]$ — e.g. $\rho = 0.9330 + 15.6682i$ ($\delta = +0.433$) — and 14 more in $[0.6, 1.4] \times [100, 180]$; the control census of ζ_K on the same box winds to exactly 0. Inserting the 26 censused quadruples into (7) and adding the resulting lag increment to the measured ζ_K matrix predicts the measured Λ_E break within the preregistered factor-2 gate on the deepest pick: ratio 0.803, with band-limited mode-read correlation $0.917 \geq 0.90$ (modes $t_k \leq 0.9 \times$ census top; the raw lag correlation 0.720 is printed as the honest out-of-band residual, since the sinc^2 reach extends past the censused band). The detector mechanism of Theorem 1 is therefore quantitative, not narrative: where off-line zeros provably exist, the same formulas that certify the ζ family predict the magnitude of the break they cause.

4 The threshold map: what the certificate excludes

On the ζ side the measured base is: $\min_k R_A(k) > 0$ on all 67/67 complete windows (global minimum 7.199; assembly zero-free). By Theorem 1, if all zeros except one additional hypothetical quadruple at (γ, δ) lie on the line, then $R_A(k) \geq 2 \text{Re } T_k(\gamma + i\delta)$ for every mode k . This yields a two-level certificate on the resolved band $\gamma \leq \pi/D_w$:

- (i) *Counting level (strength-blind)*. An additional quadruple of even minimal strength $\delta = 5 \times 10^{-4}$ is already incompatible with the observed reads on 642 of 644 in-band (window, γ) cells: the family reads are tight at the multiplicity level.
- (ii) *Strength level*. The positivity-break threshold $s_{\min}(w, \gamma) = \min\{\delta : \exists k, R_A(k) + 2 \text{Re } T_k(\gamma + i\delta) < 0\}$ is mapped over all 67 windows and a declared γ grid (644 cells, geometric scan plus 18-step bisection):

γ	windows in band	blind	median $s_{\min}$	range	median $2\alpha s_{\min}$
5.0	67	0	0.205	0.153–0.464	2.15
8.0	67	0	0.150	0.108–0.385	1.60
14.1	67	2	0.156	0.068–0.452	1.43
20.0	67	1	0.151	0.061–0.445	1.60
30.0	67	2	0.182	0.064–0.485	1.74
50.0	67	0	0.219	0.090–0.466	2.04
80.0	67	1	0.215	0.058–0.476	2.36
130.0	65	2	0.256	0.073–0.497	2.38
210.0	53	1	0.257	0.083–0.450	2.63
340.0	38	0	0.215	0.112–0.479	2.40
550.0	19	0	0.233	0.161–0.372	2.49

(“Blind” cells, where no $\delta \leq 0.4999$ breaks the form, are excluded from median and range.) The dimensionless law $2\alpha s_{\min} \approx 1.4$ –2.6 echoes the Ihara detection law $K^* \cdot s \approx 2$ –3: in both laboratories a violation of exponential strength s needs window depth of order $(2\text{--}3)/s$.

Honest reach and honest caveats. The resolved band tops out at $\pi/D \leq 868$ over the family — nine to ten orders of magnitude below the classical verification height 3×10^{12} of Platt–Trudgian [10]. The value of the certificate is not range but *type*: it is a quadratic-form positivity detector, assembled prime-side, with no Z -function sign-change counting anywhere. Four caveats are declared, not waived: (a) the single-violator reading — conspiring negative contributions of several off-line zeros could mask a strong one (the same caveat holds for the Ihara ground truth); (b) the certificate excludes an *additional* quadruple on top of the observed comb, it does not re-derive the comb; (c) the strength floor: $\delta < s_{\min}$ is invisible; (d) the per-lag dictionary is verified at the $S(T)$ -blind level stated above.

5 The $C = 1$ quadrature theorem

The deployed uniform-constant law of the program states that on the complete-comb surface the model/real contraction obeys $\varepsilon \cdot h \leq 1$ with $\varepsilon = |q_r/q_m|$, measured maximum $\varepsilon h = 0.982$ at $h = 184$ (tertile medians 0.607/0.447/0.387 over the 67 windows). `c1_mechanism_probe.py` (12 checks, 0 failures, verdict C1-QUADRATURE-MECHANISM) derives the mechanism in four steps.

- (1) *Identity.* On a complete comb the window pairing is an exact Weil pair, so the numerator is the zero-side read of the lock profile,

$$q_r = 2 \sum_{\gamma > 0} h_v(\gamma) - [h(i/2) + h(-i/2)], \quad h_v(\gamma) = D \operatorname{sinc}^2\left(\frac{\gamma D}{2}\right) \sum_d W_d \cos(\gamma d D), \quad (8)$$

verified on all 67 windows (median relative residual 1.1×10^{-2} , worst 9.2×10^{-2} , inside the declared alias-sampling budgets; zero comb extended live to $n = 2500$ to cover the second Nyquist alias of every window). The predicted sign is positive on 67/67: the lock sign *is* zero-side positivity.

- (2) *The lock direction is the pole killer.* The discrete 2×2 pole block is numerically rank one (eigenvalue ratios $\sim 3 \times 10^{16}$) and its null direction coincides with the closed lock law to $\cos \angle = 1.000000$: along the lock the exponentially large pole read cancels and the zero tail carries q_r (residual pole share median 5.6%).
- (3) *The order h^{-1} is the DST normalization.* The rescaled read $q_r (N/2)/F_\alpha$ (with the closed window factor $F_\alpha = \omega_1^2(v_1 + 2v_2)^2$) is h -flat: measured log-log slope $+0.039$ (bar ± 0.15). The h^{-1} of the law is the ℓ^2 -mode normalization $4/N$, times an α -slowly-varying zero integral.
- (4) *The constant is a quadrature.* The zero-free Riemann–von Mangoldt density integral $q_{\text{dens}} = 2 \int_{\gamma_1}^{\infty} h_v(\gamma) dN_{\text{RvM}}(\gamma) - \text{poles}$ predicts the measured q_r with median ratio $q_{\text{dens}}/q_r = 1.005$ (IQR 0.82–1.28; preregistered band $[0.5, 2]$), and reproduces the falling tertile medians of εh : predicted 0.659/0.434/0.375 versus measured 0.607/0.447/0.387. The per-window spread around the prediction is the $\sin^2$ sampling of the actual low zeros — the only arithmetic left in the constant.

Three honest refinements sharpen the statement rather than soften it. *First*, on the clean surface ε decays like $h^{-1.28}$, steeper than h^{-1} ; the historical exponent -1.01 reconstructs exactly only when the two retired truncation-flip windows are added back — it was flip-contaminated. *Second*, the contract’s operator reading “residue operator of norm $\leq 1/h$ ” is false as an operator statement: the full-space G -norm of $R_h = A_{\text{model}} - A_{\text{real}}$ grows like $h^{+1.03}$ (measured $h \|R\|_G / \|A\|_G = 94/79/80$), and only along the pole-killed lock direction does $h |q_r/q_m| = 0.982/0.434/0.230 \leq 1$ hold — $C = 1$ is a *direction* statement. *Third*, the density average predicts the median, not the supremum: the measured headroom $\max \varepsilon h = 0.982 \leq 1$ is *not* proven by the average alone; the missing piece is the low-zero sampling budget $\sin^2 \leq 1$, named, not claimed. Since all arithmetic input sits at heights $\leq 2 \times$ the Nyquist edge $2\pi/D$ ($\leq 3471 \ll 3 \times 10^{12}$: verified-RH territory by [10]), the $C = 1$ lever is demystified as a computable discretization statement — the RH-hard content of the program is *not* hidden in it.

6 The Fejér density identity: the connecting lemma

Lemma 2 (Fejér density identity, verified). *For a window $[-a, a]$ let $s_{\text{tot}}(t; a)$ be the continuum total symbol per unit plane-wave mass (archimedean layer minus tent prime layer plus rank-two pole layer, assembled from primes and digamma only). Applying the explicit formula to the*

admissible pair $g_t(u) = \frac{1}{2a}(2a - |u|)_+ e^{-itu}$, $h_t(r) = 2\pi F_a(r - t)$, $F_a(s) = \sin^2(as)/(\pi a s^2)$, gives exactly

$$s_{\text{tot}}(\cdot; a) = 2\pi (F_a * dN), \quad dN = \sum_{\rho} \delta_{\gamma_{\rho}} : \quad (9)$$

the symbol is the Fejér-smoothed zero-counting density.

Verification (`fejerdensityboundprobe.py`, 14 checks, 0 failures, verdict FEJER-DENSITY-THEOREM): pointwise to 2.9×10^{-6} at 25 declared points on two windows, Hann t -averaged to 7.2×10^{-8} ; the first four peaks of the zero-free symbol sit on $\gamma_1, \dots, \gamma_4$ to 2×10^{-4} ; on $[10, 300]$ the peak-zero matching is 138/138 with median residual 0.0013; the dips are the zero gaps quantitatively (dip forensics residual 1.5×10^{-6} ; Fejér dip power measured 0.974 vs predicted 0.974).

Corollary 3 (unconditional smoothed-density bound). *Let $\rho_{a,\delta} = s_{\text{tot}} * \Phi_{\delta}$ with the unit-mass Fejér kernel Φ_{δ} of main-lobe half-width δ . Windowed Riemann–von Mangoldt counting with Trudgian’s bound $|S(t)| \leq 0.112 \log t + 0.278 \log \log t + 2.510$ [9], the Fejér box minorant and F_a mass transfer give, for $\delta \in \{4\pi, \pi a\}$ and every family window,*

$$\rho_{a,\delta}(t) \geq c_0 \log(2+t) - C_0 \quad (t \geq t_0(a, \delta)),$$

with the finite remainder $[10, t_0]$ machine-checked on the zero-free grids: all 10 (a, δ) certificates pass with margins $\geq +1.36$ (bar 0.05), $\min c_0 = 0.3058$, $C_0 \leq 2.50$, $t_0 \leq 3530$ — a $15.1\times$ gain over the raw measured hull slope 0.0202. The chain closes only above the certifiability floor $\delta > 4\pi \cdot 0.112 = 1.4074$: the literal width $1/a \leq 0.361$ and the plane-wave resolution $\pi/a \leq 1.133$ sit below the floor for every family window, so per-mode (pointwise) symbol control is not certifiable this way at any height — the honest structural explanation of the measured $1/\log$ Gårding drift.

The lemma is the hinge between the three results. The tent weight $(1 - u/2\alpha)$ of s_{tot} is exactly the closed-form DST mode weight (4) of the finite window; the finite alias comb T_k of Theorem 1 is the lattice reading of (9), its off-line term $\text{OFF}(x; \delta, a) \rightarrow 2F_a(x)$ as $\delta \rightarrow 0$; and the $C = 1$ quadrature integral (8) is the Riemann–von Mangoldt read of the same object along the lock profile. One identity, three appearances.

7 Scope: what is and is not claimed

Circularity ledger (five points, declared). (1) All window forms and mode reads are assembled from primes and digamma only; no zero enters any assembly. (2) Riemann zeros enter only the verification of the identities (5)–(6), through the Turing-certified cache. (3) The Epstein census is a measured output of the argument principle on lattice data, never an input table. (4) The certificate of Sec. 4 excludes an additional quadruple against the observed reads; it does not re-derive the observed comb. (5) The uniform statement is not approached from below by any ladder in this note; the only non-circular residual lever identified by the program is the $C = 1$ mechanism of Sec. 5, itself now typed as quadrature.

step	statement	status
W1	window form = Galerkin read of Suzuki's B_a	closed at exact/certified level (companion note [2]; suite modules v630–v644).
W2	form density; spectral density control	density identity (9) verified ; smoothed lower bound theorem-grade with $c_0 = 0.306$ (Cor. 3); the frame-level Gårding inequality $Q + C_0 G \succeq c_0 \sum_p w_p \varphi_p \varphi_p^\top$ holds at every M (packet probe; its packet floor is now <i>unconditional</i> by the zero-gap addendum below); the <i>pointwise</i> Gårding bound and the formal Mosco limit stay open (the packet ladder fails its stabilization gate; the $1/\log$ drift persists and is explained by the floor of Cor. 3).
W3	positivity of the window family as a zero statement	this note : structure theorem (Thm. 1) + two-laboratory validation + certificate with threshold map = “RH restricted to the resolved band $\gamma \leq \pi/D$ above the strength floor $s_{\min}$ ”. <i>Uniform</i> W3 over all windows/all a is the conjecture, equivalent to RH [3, 4, 5]; there is no ladder underneath that wall.
W4	positivity transfer to Weil positivity	<i>not claimed</i> , not touched here.

Addendum: the zero-gap theorem and the named pointwise target. After the numbers of this note were frozen, the zero-gap route (**zero_gap_theorem_probe.py**; nine checks plus one *declared* typed-residual failure, stated as such) supplied the missing unconditional input of the W2 row. An explicit unconditional zero-gap theorem is assembled from cited ingredients: for $t \geq 10$ every interval $(t, t + H_{\min}(t)]$ contains a zero ordinate, where $H_{\min}$ is the bisection root of the windowed Riemann–von Mangoldt S -difference minorant with S_{bound} the pointwise minimum of three unconditional bounds: Platt’s computed $|S(T)| \leq 2.5167$ for $T \leq 3.06 \times 10^{10}$ [11, 12] (quoted as eq. (1.2) in [12]), Trudgian’s bound [9] (floor $4\pi \cdot 0.112 = 1.40743$), and Bellotti–Wong [12, Cor. 1.5] (best asymptotic floor $4\pi \cdot 0.10076 = 1.26619$). The theorem is machine-verified on all 1999 Turing-certified comb gaps (0 violations on every chain separately; minimum air $2.05\times$, median $5.24\times$; $H_{\min}$ decreasing from 25.68 at $t = 10$ to 1.49 at $t = 10^{10}$), and adaptive packet bands $\delta_p = \kappa H_{\min}$ unconditionally hold $\geq \kappa$ ordinates each (0 misses at $\kappa = 2, 3$): the frame packet floor of the W2 row is now *unconditional* ($0 < B_p \leq V_p$ on all packets). The projection (pointwise) form does *not* close this way — the minimizer stays single-mode tight; re-partitioning by adaptive widths re-labels the obstruction, it does not remove it — and the last pointwise gap now carries a single named number: main-lobe capture of a guaranteed zero at the family anchor $a_0 = \log 16$ needs $H_{\min} < \pi/a_0 = 1.1331$, i.e. a leading argument constant $A_1 < 1/(4a_0) = 0.09017$ in $|S(t)| \leq A_1 \log t + \dots$. The best cited constant, $A_1 = 0.10076$ [12], misses by 11.7%, and the Platt branch bottoms out at $1.4183 > \pi/a_0$ at every height: the last W2 gap is a named target inequality at the frontier of explicit analytic number theory, independent of RH. Nothing in this note depends on it.

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 [15] as the modules **v668** (**ihara_ground_truth_probe.py** and **epstein_firewall_probe.py** jointly: the two laboratories of Sec. 3), **v669** (**fejer_density_bound_probe.py**: the Fejér density identity and the unconditional chain of Sec. 6), **v674** (**packet_garding_probe.py**: the frame-level packet Gårding form of the W2 row), **v676** (**c1_mechanism_probe.py**: the $C = 1$ quadrature of Sec. 5), **v677** (**w3_structure_theorem_probe.py**: the structure theorem and threshold map of Secs. 2–4), and **v678** (**zero_gap_theorem_probe.py**: the zero-gap addendum above). The declared-FAIL gates of the packet and zero-gap probes are promoted as typed-residual checks with inverted expectation, all numbers unchanged.

Every number above is the printed output of an executed, preregistered run (five probes, 83 checks, 0 failures on the quoted gates; the addendum quotes the zero-gap probe, whose

single declared typed-residual gate is stated as such); calibration repairs made between runs are declared inside the probes themselves. No positivity is claimed beyond the measured certificate, no marker of the verification program is moved by this note, and no statement about the Riemann Hypothesis is made.

References

- [1] M. Suzuki, *Weil’s quadratic form via the screw function*, arXiv:2606.09096 (2026).
- [2] S. Hamann, *The window quadratic form of the TFPT prime-front program is a Galerkin discretization of Suzuki’s localized Weil operator*, companion working note (August 2, 2026), submitted to arXiv in parallel with the present note; repository file `note_w1_suzuki_identification.pdf` [15].
- [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] E. Bombieri and J. C. Lagarias, *Complements to Li’s criterion for the Riemann hypothesis*, J. Number Theory 77 (1999), 274–287.
- [6] A. Cantoni and P. Butler, *Eigenvalues and eigenvectors of symmetric centrosymmetric matrices*, Linear Algebra Appl. 13 (1976), 275–288.
- [7] U. Grenander and G. Szegő, *Toeplitz forms and their applications*, Univ. of California Press, Berkeley (1958).
- [8] H. Iwaniec and E. Kowalski, *Analytic number theory*, AMS Colloq. Publ. 53, Providence (2004).
- [9] T. Trudgian, *An improved upper bound for the argument of the Riemann zeta-function on the critical line II*, J. Number Theory 134 (2014), 280–292.
- [10] D. J. Platt and T. S. Trudgian, *The Riemann hypothesis is true up to $3 \cdot 10^{12}$* , Bull. London Math. Soc. 53 (2021), 792–797.
- [11] D. J. Platt, *Isolating some non-trivial zeros of zeta*, Math. Comp. 86 (2017), 2449–2467; zero database at LMFDB.
- [12] C. Bellotti and P.-J. Wong, *Improved estimates for the argument and zero-counting function of the Riemann zeta-function*, Math. Comp. (2025), doi:10.1090/mcom/4133; arXiv:2412.15470.
- [13] H. Davenport and H. Heilbronn, *On the zeros of certain Dirichlet series*, J. London Math. Soc. 11 (1936), 181–185.
- [14] H. Bass, *The Ihara–Selberg zeta function of a tree lattice*, Internat. J. Math. 3 (1992), 717–797.
- [15] TFPT repository (<https://fixpoint-theory.com>): frozen window machinery `verification/v563_paper2_readouts.py`; the discovery probes of the title footnote, `packet_garding_probe.py` and `zero_gap_theorem_probe.py`, all in `experiments/tfpt-discovery/`; suite modules `v668`, `v669`, `v674`, `v676–v678` (Remark 4) (2026).