
Spectral Properties of the Truncated Weil Operator

& Numerical Verification

of Weil Positivity

VERSION — V.1.0.0

Kerym Makraini

UNED — National University of Distance Education
AGE Quantum Gates Engine S.L., Melilla, Spain
Department of Physics & Computational Mathematics
`mhamed34@alumno.uned.es`

March 2026

Abstract

We study the truncated Weil operator A_Λ on $L^2([-\log \Lambda, \log \Lambda])$, following Connes’s programme [3, 4] for the Riemann Hypothesis.

We establish five unconditional spectral results: a trace–energy identity providing structural positivity on the Sonin space, J -self-adjointness with respect to the Kreĩn involution, resolvent convergence at rate $O(\Lambda^{-1/2})$, a variational lower bound $\varepsilon_0(\Lambda) \leq -2\sqrt{\Lambda} + C$, and a Hilbert–Schmidt decomposition separating archimedean and atomic contributions. We further prove that *norm convergence of the truncated operators to the prolate projection is impossible*, ruling out the most natural convergence route.

Our main contribution is a precise structural reduction: the Riemann Hypothesis in this framework is equivalent to *asymptotic kernel invisibility*—the property that Weil test functions become orthogonal to the kernel of the compressed atomic operator as $\Lambda \rightarrow \infty$. All currently identified finite-scale algebraic obstructions are resolved; the remaining difficulty is concentrated in this single analytic statement, which we formulate as a concrete, falsifiable problem.

2020 MSC: 11M26, 47A10, 46C20, 47B50, 30D15, 58B34

Keywords: Riemann zeta function, Weil explicit formula, Weil positivity, truncated operators, Kreĩn spaces, prolate spheroidal wave functions, Sonin space, trace formula, spectral theory

Contents

1	Introduction	4
2	Setup and notation	6
2.1	The truncated Weil operator	6
2.2	The prolate operator	6
2.3	The Sonin space	6
2.4	The compressed Weil form	7
2.5	The Kreĩn involution	7
3	Trace–energy identity	7
4	J-self-adjointness and Kreĩn structure	8
5	Resolvent convergence	9
6	Variational lower bound for the ground state	9
7	Hilbert–Schmidt decomposition & atomic bounds	10
8	Theorem of Obstruction: Norm Convergence is Impossible	11
8.1	The atomic measure and variation	11
8.2	The obstruction theorem	12
8.3	Quantitative reduction criterion	13
9	Numerical Validation	13
9.1	Protocol Kerym-V: Reproducible Numerical Framework	13
9.2	Graphical Analysis	14

10	Sonin space implementation	14
10.1	Ξ lives in the Sonin space	14
10.2	The compressed Weil form is indefinite	14
11	Numerical verification of Weil positivity	15
11.1	The Weil criterion in shifted Mellin form	15
11.2	Weil positivity for f_Ξ	15
11.3	Weil positivity for general test functions	16
12	Spectral phenomenology of the actual Weil operator	16
12.1	Full spectral data	16
12.2	Key phenomena	17
13	The remaining open problem	18
14	Defect compression and the support transition	18
14.1	The supremal defect	18
14.2	The critical support B^*	19
14.3	Verification at $B = T/2$: the Weil support	20
14.4	Analytical bound for the compressed regime	20
15	What is proved, what is conditional, what remains open	22
16	Logical dependency graph and master reduction	22
16.1	Overview	22
16.2	Operator-theoretic block (this paper)	23
16.3	Analytic block (companion paper)	23
16.4	Master bottleneck	24
16.5	Master theorem (reduction of RH)	24
16.6	Status of the programme	24
17	Conclusion	25
A	Domain, trace, and Hilbert–Schmidt properties	25
A.1	Trace-class conditions	25
A.2	Fubini–Tonelli justifications	25
A.3	Factorisation of the Weil distribution on convolutions	26
A.4	Explicit Schur and HS estimates	26
B	Kernel estimates for the prolate comparison	27
B.1	Diagonal expansion	27
B.2	Off-diagonal oscillatory estimates	27
B.3	Alternative compactness route	27
C	Passage to the limit and positivity	27
C.1	Convergence of Q_Λ	28
C.2	Preservation of positivity	28
C.3	Density of test functions	28
C.4	Representation of the absorption terms R_Λ	28
C.5	Reduction lemma for uniform control	29

C.6	Conditional reduction to absorption control	30
D	Construction of test functions	30
D.1	Mellin asymptotics	30
D.2	Phase selection	30
D.3	Interval dominance	30
D.4	Full stationary phase calculation	31
E	Audit of external dependencies	31
E.1	Primary external results	31
E.2	Failure impact analysis	32
F	Numerical evidence (non-logical)	32
F.1	Reproducible protocol	32
F.2	Interpretation without logical use	32
F.3	Threshold robustness	33
F.4	Absorption term robustness	33
F.5	Resolution check	33
G	Numerical values of constants	34
H	Leakage decay via Slepian concentration	34
H.1	Setup and conventions	34
H.2	The concentration inequality	35
H.3	Application to Weil test functions	36
H.4	Consequence for the absorption terms	37
H.5	Numerical verification	37

Notation and conventions

Throughout, $\Lambda \geq 2$ is the prime cutoff, $T = \log \Lambda$, $\mathcal{H}_\Lambda = L^2([-T, T], dt)$ (Lebesgue measure), $\psi = \Gamma'/\Gamma$ the digamma function, $\Xi(t) = \xi(\frac{1}{2} + it)$ the Riemann xi function on the critical line, $*$ multiplicative convolution $((f * g)(x) = \int_0^\infty f(x/y) g(y) dy/y)$, and $\tilde{f}(x) = \overline{f(x^{-1})}$. The shifted Mellin transform is $F(s) = \int_0^\infty f(x) x^{s-1/2} dx/x$. Operator norms are on \mathcal{H}_Λ unless stated otherwise.

1 Introduction

The Riemann Hypothesis (RH) asserts that every non-trivial zero ρ of the Riemann zeta function $\zeta(s)$ satisfies $\text{Re}(\rho) = 1/2$. A classical reformulation, due to Weil [15], characterises RH through a positivity condition:

$$\text{RH} \iff W(f * \tilde{f}) \geq 0 \quad \text{for all } f \in C_c^\infty(0, \infty), \quad (1)$$

where W is the Weil distribution encoding archimedean and prime contributions via the explicit formula, $\tilde{f}(x) = \overline{f(x^{-1})}$, and $*$ denotes multiplicative convolution.

Connes's programme [3, 4, 5] approaches this criterion through a family of *truncated* operators A_Λ parametrised by a cutoff Λ , and a corresponding *Sonin space* S_Λ of prolate-like functions [10, 13]. The compressed quadratic form $Q_W^{(\Lambda)}(g) = \text{Tr}(T_{S_\Lambda}(g))$ satisfies

$Q_W^{(\Lambda)}(f * \tilde{f}) \geq 0$ for each finite Λ (the *truncated Weil positivity*). The passage from truncated to full positivity — i.e., from $Q_W^{(\Lambda)} \geq 0$ to $W \geq 0$ — is the central open problem in the programme and is *not addressed* in this paper.

Contributions

This paper makes the following contributions, all of which are *unconditional* (they do not assume RH or any unproved hypothesis):

- (i) A **trace–energy identity** establishing structural positivity of the compressed Weil form for convolution-type inputs (Theorem 3.1); the underlying domain and trace-class conditions are verified in Section A.
- (ii) A **Kreĭn space framework** for the Weil operator, exhibiting J -self-adjointness (Theorem 4.1) [1, 2].
- (iii) **Resolvent convergence** at rate $O(\Lambda^{-1/2})$ (Theorem 5.1).
- (iv) A **variational lower bound** $\varepsilon_0(\Lambda) \leq -2\sqrt{\Lambda} + C$ for the ground-state eigenvalue (Theorem 6.1), with explicit constant $C = \hat{K}_\infty(0) + O(1) \approx 3.95$.
- (v) **Hilbert–Schmidt decomposition** separating the archimedean and atomic parts, with explicit bounds on both (Theorem 7.1, Theorem 7.2); the Schur test estimates are recorded in Section A.

In addition, we provide:

- (vi) Implementation of the **Sonin space** S_Λ and numerical verification that Ξ lives within it (Section 10).
- (vii) **Numerical evaluation** of the Weil criterion $Z(f * \tilde{f})$ using the correct shifted Mellin transform, showing $Z > 0$ for all test functions examined (Section 11); the construction of test functions and their Mellin asymptotics are detailed in Section D.
- (viii) Complete **spectral data** for the actual Weil operator (digamma kernel + exact delta shifts) for Λ up to 200, including the discovery that $\varepsilon_1(\Lambda)$ stabilises near -6 to -7 (Section 12); kernel estimates underpinning these computations appear in Section B.
- (ix) A precise formulation of the **remaining open problem**: controlling the absorption terms in the semilocal trace formula (Section 13); the passage to the limit is analysed in Section C.

Honest status declaration

We emphasise that **this paper does not prove the Riemann Hypothesis**. The truncated positivity $Q_W^{(\Lambda)}(f * \tilde{f}) \geq 0$ is established for each finite Λ , but the passage to the limit $\Lambda \rightarrow \infty$ requires controlling absorption terms in the semilocal trace formula [4], which remains an open problem (Theorem 13.1). All claims in this paper are unconditional results in spectral theory and functional analysis.

2 Setup and notation

2.1 The truncated Weil operator

Definition 2.1 (Truncated Weil operator). Let $\Lambda \geq 2$ and $T = \log \Lambda$. The truncated Weil operator A_Λ acts on $\mathcal{H}_\Lambda = L^2([-T, T], dt)$ with kernel

$$K^{(\Lambda)}(t, s) = K_\infty(t - s) - \sum_{p \leq \Lambda} \sum_{m \geq 1} (\log p) p^{-m/2} [\delta(t - s - m \log p) + \delta(t - s + m \log p)], \quad (2)$$

where the *archimedean kernel* K_∞ has Fourier transform

$$\widehat{K}_\infty(\xi) = \log(2\pi) - \frac{1}{2} \operatorname{Re} \psi\left(\frac{1}{4} + \frac{i\xi}{2}\right), \quad (3)$$

with $\psi = \Gamma'/\Gamma$ the digamma function.

Remark 2.2. The operator (2) corresponds to Connes's construction [3] restricted to the “semilocal system” with primes $p \leq \Lambda$. We verify in Theorem A.1 that it is Hilbert–Schmidt on \mathcal{H}_Λ .

Remark 2.3 (Precise interpretation of the atomic terms). The delta distributions in (2) are interpreted as *shift operators* acting on \mathcal{H}_Λ . For $a \in \mathbb{R}$ with $|a| < 2T$, define the bounded operator S_a by $(S_a f)(t) = f(t - a)$ for $t - a \in [-T, T]$ and $(S_a f)(t) = 0$ otherwise. Then $\|S_a\|_{\text{op}} \leq 1$ and

$$\langle f, S_a f \rangle = \int_{-T}^T \overline{f(t)} f(t - a) dt \quad (4)$$

is a bounded sesquilinear form. The atomic part of A_Λ is the finite sum

$$A_\Lambda^{\text{atom}} = - \sum_{p \leq \Lambda} \sum_{m \geq 1} (\log p) p^{-m/2} (S_{m \log p} + S_{-m \log p}), \quad (5)$$

which converges in operator norm because $\sum_{p, m} (\log p) p^{-m/2} < \infty$ (geometric sum in m , PNT in p). The full operator is $A_\Lambda = A_\Lambda^{\text{cont}} + A_\Lambda^{\text{atom}}$, where A_Λ^{cont} is the convolution operator with kernel K_∞ restricted to $[-T, T]$.

2.2 The prolate operator

The *prolate concentration operator* P_c on $L^2([-T, T])$ with bandwidth $c > 0$ has kernel

$$P_c(t, s) = \frac{\sin(c(t - s))}{\pi(t - s)}. \quad (6)$$

Its eigenfunctions $\{\varphi_k\}_{k=0}^\infty$ are the prolate spheroidal wave functions (PSWFs) [10, 13], with eigenvalues $\mu_0 > \mu_1 > \dots > 0$.

2.3 The Sonin space

Definition 2.4 (Sonin space). The Sonin space $S_\Lambda \subset \mathcal{H}_\Lambda$ is the span of PSWFs with eigenvalue exceeding a threshold τ :

$$S_\Lambda = \operatorname{span}\{\varphi_k : \mu_k > \tau\}, \quad (7)$$

with $c = \sqrt{\Lambda}$ and $\tau = 0.01$. Its dimension $N_S \approx cT/\pi$ (the Shannon number); see [10].

Remark 2.5 (Robustness of the threshold). The choice $\tau = 0.01$ is conventional; the qualitative results (in particular the stability of $Z/\|f_\Xi\|^2$ reported in Theorem 11.1 and the projection errors in Table 4) have been verified to be insensitive to τ in the range $[0.005, 0.05]$. At $\Lambda = 100$, varying τ in this range changes N_S by at most ± 2 and modifies $Z/\|f_\Xi\|^2$ by less than 0.5% (Section F.3).

2.4 The compressed Weil form

The *compressed Weil form* on S_Λ is the bilinear form

$$Q_W^{(\Lambda)}(g) = \text{Tr}(T_{S_\Lambda}(g)), \quad T_{S_\Lambda}(g) = \Pi_{S_\Lambda} \circ W_g \circ \Pi_{S_\Lambda}, \quad (8)$$

where Π_{S_Λ} is the orthogonal projection onto S_Λ and W_g is the operator on \mathcal{H}_Λ associated with the Weil distribution applied to g . Concretely, for $g = f * \tilde{f}$, the operator W_g acts via $W_g = \hat{f}(D)\hat{f}(D)^*$, where $\hat{f}(D)$ is the operator obtained by substituting the ‘‘Dirac operator’’ D of the spectral realisation into the Mellin transform of f (see [3] and Theorem 3.1 for the precise factorisation). The compression $T_{S_\Lambda}(g)$ is therefore a *finite-dimensional* operator on S_Λ (a matrix of size $N_S \times N_S$), and its trace is well defined by Theorem A.1.

2.5 The Kreĭn involution

Definition 2.6 (*J*-involution). Define $J : \mathcal{H}_\Lambda \rightarrow \mathcal{H}_\Lambda$ by $(Jf)(t) = f(-t)$, corresponding to $f(u) \mapsto f(u^{-1})$ in multiplicative coordinates. Then $J^2 = I$ and $J = J^*$.

3 Trace–energy identity

Theorem 3.1 (Trace–energy identity). *For every $f \in C_c^\infty(0, \infty)$ and every $\Lambda > 0$,*

$$Q_W^{(\Lambda)}(f * \tilde{f}) = \text{Tr}(T_{S_\Lambda}(f * \tilde{f})) = \|A^* S_\Lambda\|_{\text{HS}}^2 \geq 0, \quad (9)$$

where $A = \Pi_{S_\Lambda} \circ \hat{f}(D)$ and $\hat{f}(D)$ denotes the operator obtained by applying the Mellin transform of f to the spectral realisation. All iterated integrals are absolutely convergent (Theorem A.2).

Proof. The key step is the *factorisation* of the Weil distribution on convolution-type inputs. For $g = f * \tilde{f}$, the Mellin transform satisfies $\hat{g}(s) = \hat{f}(s)\hat{f}(1-\bar{s})$, which in operator language reads

$$W_{f*\tilde{f}} = \hat{f}(D)\hat{f}(D)^*, \quad (10)$$

where $\hat{f}(D)$ denotes the operator on \mathcal{H}_Λ obtained by applying the Mellin transform of f to the spectral variable of the Weil system (see [3, Section 3] for the construction of the spectral realisation and the functional calculus $\hat{f}(D)$); the adjoint is taken with respect to the L^2 inner product on \mathcal{H}_Λ).

Compressing to the Sonin space:

$$T_{S_\Lambda}(f * \tilde{f}) = \Pi_{S_\Lambda} \hat{f}(D)\hat{f}(D)^* \Pi_{S_\Lambda} = AA^*,$$

where $A = \Pi_{S_\Lambda} \hat{f}(D)$ maps \mathcal{H}_Λ into the finite-dimensional space S_Λ . For any $v \in S_\Lambda$,

$$\langle v, AA^*v \rangle = \langle A^*v, A^*v \rangle = \|A^*v\|^2 \geq 0.$$

Hence AA^* is positive semidefinite on S_Λ , and

$$\mathrm{Tr}(T_{S_\Lambda}(f * \tilde{f})) = \mathrm{Tr}(AA^*) = \sum_{k=1}^{N_S} \|A^* e_k\|^2 = \|A^*\|_{\mathrm{HS}}^2 \geq 0,$$

where $\{e_k\}_{k=1}^{N_S}$ is any orthonormal basis of S_Λ . The trace is finite because S_Λ is finite-dimensional (Theorem A.1). The absolute convergence of all integrals is justified in Theorem A.2. \square

Remark 3.2 (Scope of positivity). The identity (9) establishes positivity for *convolution-type* inputs $g = f * \tilde{f}$, not for arbitrary elements of S_Λ . The compressed operator T_{S_Λ} is *indefinite* on S_Λ (it has both positive and negative eigenvalues; see Table 5), which is consistent with the Weil distribution being indefinite on general test functions.

4 J -self-adjointness and Kreĭn structure

Theorem 4.1 (J -self-adjointness). *The truncated Weil operator A_Λ is J -self-adjoint in the Kreĭn space $(\mathcal{H}_\Lambda, [\cdot, \cdot]_J)$ with indefinite inner product $[f, g]_J = \langle Jf, g \rangle$, i.e.,*

$$A_\Lambda = JA_\Lambda^*J. \quad (11)$$

Proof. We establish (11) in two steps.

Step 1: A_Λ is self-adjoint in the Hilbert space $(\mathcal{H}_\Lambda, \langle \cdot, \cdot \rangle)$. The operator A_Λ is bounded (it is Hilbert–Schmidt by Theorem 7.1). Its kernel satisfies $\overline{K^{(\Lambda)}(s, t)} = K^{(\Lambda)}(s, t)$ (all terms are real) and $K^{(\Lambda)}(t, s) = K^{(\Lambda)}(s, t)$ (the archimedean kernel $K_\infty(t-s) = K_\infty(s-t)$ is even, and the atomic shifts appear symmetrically). Hence $A_\Lambda^* = A_\Lambda$ on all of \mathcal{H}_Λ .

Step 2: A_Λ commutes with J . The archimedean kernel satisfies $K_\infty((-t) - (-s)) = K_\infty(s - t) = K_\infty(t - s)$. Each atomic shift appears with both signs ($\delta(w \pm m \log p)$), so $K^{(\Lambda)}(-t, -s) = K^{(\Lambda)}(t, s)$. For $f \in \mathcal{H}_\Lambda$:

$$\begin{aligned} (JA_\Lambda f)(t) &= (A_\Lambda f)(-t) = \int_{-T}^T K^{(\Lambda)}(-t, s) f(s) ds \\ &= \int_{-T}^T K^{(\Lambda)}(t, -s) f(s) ds \stackrel{u=-s}{=} \int_{-T}^T K^{(\Lambda)}(t, u) f(-u) du = (A_\Lambda Jf)(t). \end{aligned}$$

Thus $JA_\Lambda = A_\Lambda J$.

Combining: $JA_\Lambda^*J = JA_\Lambda J = A_\Lambda J^2 = A_\Lambda$, establishing (11). The Kreĭn space structure follows from the general theory of [1, Chapter I]: since J is a fundamental symmetry and A_Λ is bounded and J -self-adjoint, the spectrum of A_Λ is real and the eigenspaces carry a definite $[\cdot, \cdot]_J$ -structure on each eigenspace [2, Theorem I.3.1]. \square

Corollary 4.2 (Eigenvalue parity). *Every eigenvalue of A_Λ has an eigenvector that is either J -even ($\varphi(-t) = \varphi(t)$) or J -odd ($\varphi(-t) = -\varphi(t)$).*

Proof. Since J commutes with A_Λ , each eigenspace is J -invariant. Since $J^2 = I$, the eigenspace decomposes into the ± 1 eigenspaces of J [2, Theorem I.3.1]. \square

5 Resolvent convergence

Theorem 5.1 (Resolvent convergence). *Let $\Lambda' > \Lambda \geq 2$ and let \tilde{A}_Λ denote A_Λ transported to $\mathcal{H}_{\Lambda'}$ by zero extension. Let $K \subset \mathbb{C}$ be a compact set satisfying $\text{dist}(K, \text{spec}(A_{\Lambda'})) \geq \delta > 0$ for some δ independent of Λ' (for instance, K may be taken in the resolvent set of the limiting archimedean operator). Then*

$$\sup_{z \in K} \|(z - \tilde{A}_\Lambda)^{-1} - (z - A_{\Lambda'})^{-1}\|_{\text{op}} = O\left(\frac{1}{\sqrt{\Lambda}}\right) \quad (12)$$

as $\Lambda \rightarrow \infty$.

Proof. The difference $A_{\Lambda'} - \tilde{A}_\Lambda$ consists of: (a) additional primes $p \in (\Lambda, \Lambda']$, contributing operator norm at most $\sum_{\Lambda < p \leq \Lambda'} \log p / (\sqrt{p} - 1)$; and (b) the extension of the archimedean kernel to the larger domain. By the prime number theorem [8], $\sum_{p \leq x} \log p / \sqrt{p} = 2\sqrt{x} + O(1)$, so the additional prime contribution is $O(1/\sqrt{\Lambda})$. The domain extension contributes $O(1/\Lambda)$ by the decay of $K_\infty(w)$ for large $|w|$ (see Theorem B.2). The resolvent identity

$$(z - A_{\Lambda'})^{-1} - (z - \tilde{A}_\Lambda)^{-1} = (z - A_{\Lambda'})^{-1}(A_{\Lambda'} - \tilde{A}_\Lambda)(z - \tilde{A}_\Lambda)^{-1}$$

and uniform resolvent bounds (for $z \in K$, $\|(z - A)^{-1}\|_{\text{op}} \leq 1/\delta$ where $\delta = \text{dist}(K, \text{spec}(A))$; see [12, Theorem VIII.19]) yield (12). The uniformity in $z \in K$ follows from the compactness of K and the continuity of the resolvent on the resolvent set. \square

6 Variational lower bound for the ground state

Theorem 6.1 (Variational bound). *For all $\Lambda \geq 2$,*

$$\varepsilon_0(\Lambda) \leq -2\sqrt{\Lambda} + C, \quad (13)$$

where $C = \hat{K}_\infty(0) + O(1)$ and $\hat{K}_\infty(0) = \log(2\pi) - \frac{1}{2}\psi(\frac{1}{4}) \approx 3.9516$.

Proof. Take the constant test function $f_0(t) = (2T)^{-1/2}$, $T = \log \Lambda$. Then $\|f_0\|_{L^2} = 1$.

Step 1 (Archimedean). By the Fourier inversion formula, $\langle f_0, A_\Lambda^{\text{arch}} f_0 \rangle = \hat{K}_\infty(0) + O(T^{-1})$.

Step 2 (Prime atoms, $m = 1$). For a shift by $a = \log p < 2T$: $\langle f_0, \delta(\cdot - a)f_0 \rangle = (2T - a)/(2T)$. Summing the $m = 1$ contributions with coefficient $-(\log p)p^{-1/2}$:

$$S_1(\Lambda) = 2 \sum_{p \leq \Lambda} \frac{\log p}{\sqrt{p}} - \frac{1}{\log \Lambda} \sum_{p \leq \Lambda} \frac{(\log p)^2}{\sqrt{p}}.$$

By the prime number theorem [8] (partial summation with $\theta(x) \sim x$):

$$\sum_{p \leq \Lambda} \frac{\log p}{\sqrt{p}} = 2\sqrt{\Lambda} + O(1), \quad (14)$$

$$\sum_{p \leq \Lambda} \frac{(\log p)^2}{\sqrt{p}} = 2\sqrt{\Lambda} \log \Lambda - 4\sqrt{\Lambda} + O(\sqrt{\Lambda}). \quad (15)$$

Substituting (14)–(15):

$$\begin{aligned}
S_1(\Lambda) &= 2(2\sqrt{\Lambda} + O(1)) - \frac{1}{\log \Lambda} (2\sqrt{\Lambda} \log \Lambda - 4\sqrt{\Lambda} + O(\sqrt{\Lambda})) \\
&= 4\sqrt{\Lambda} - 2\sqrt{\Lambda} + \frac{4\sqrt{\Lambda}}{\log \Lambda} + O(1) \\
&= 2\sqrt{\Lambda} + \frac{4\sqrt{\Lambda}}{\log \Lambda} + O(1).
\end{aligned}$$

Step 3 (Higher m). For $m \geq 2$: $S_{\geq 2}(\Lambda) \leq 2 \sum_p \log p / (p(\sqrt{p} - 1)) = O(1)$.

Step 4 (Combining).

$$\varepsilon_0(\Lambda) \leq \langle f_0, A_\Lambda f_0 \rangle = \hat{K}_\infty(0) - S_1(\Lambda) - S_{\geq 2}(\Lambda) + O(T^{-1}) \leq -2\sqrt{\Lambda} + C. \quad \square$$

Remark 6.2 (Sharpness and interpretation). The constant $C \approx 3.95$ is *positive*, so the bound reads $\varepsilon_0 \leq -(2\sqrt{\Lambda} - 3.95)$, which is negative for $\Lambda \geq 4$. Note that (13) is an *upper* bound on ε_0 (which is negative), i.e., it says the ground state is *at least as negative* as $-2\sqrt{\Lambda} + C$. The phrasing “captures 88%–97%” in Table 1 refers to the ratio $|\text{bound}|/|\varepsilon_0|$, confirming that the trial function accounts for most of the ground-state energy.

Table 1: Comparison of the variational bound with actual $\varepsilon_0(\Lambda)$. Numerical data from the actual Weil kernel (Section F).

Λ	$-2\sqrt{\Lambda}$	Bound $\langle f_0, A_\Lambda f_0 \rangle$	Actual ε_0	Bound/Actual
20	-8.94	-5.77	-6.56	0.88
50	-14.14	-12.44	-13.23	0.94
100	-20.00	-19.73	-20.64	0.96
200	-28.28	-31.12	-32.85	0.95

7 Hilbert–Schmidt decomposition & atomic bounds

Theorem 7.1 (HS decomposition). *Write $A_\Lambda = A_\Lambda^{\text{cont}} + A_\Lambda^{\text{atom}}$ where A_Λ^{cont} is the integral operator with kernel $K_\infty(t-s)$ on $[-T, T]$, and A_Λ^{atom} is the finite-rank shift operator from the atomic terms. Then:*

$$\|A_\Lambda^{\text{cont}}\|_{\text{HS}}^2 = 2T \cdot \|K_\infty\|_{L^2(\mathbb{R})}^2 + O(1) = O(\log \Lambda), \quad (16)$$

where $\|K_\infty\|_{L^2}^2 = \int_{\mathbb{R}} |\hat{K}_\infty(\xi)|^2 d\xi < \infty$ by the digamma asymptotics $\hat{K}_\infty(\xi) = O(|\xi|^{-1})$ (Plancherel; see Theorem A.4 for the explicit Schur bounds).

Proof. The Hilbert–Schmidt norm of a convolution operator on $[-T, T]$ with kernel $K \in L^2(\mathbb{R})$ satisfies

$$\|A^{\text{cont}}\|_{\text{HS}}^2 = \int_{-T}^T \int_{-T}^T |K_\infty(t-s)|^2 ds dt \leq 2T \|K_\infty\|_{L^2}^2.$$

The finiteness of $\|K_\infty\|_{L^2}$ follows from Plancherel and the decay $\hat{K}_\infty(\xi) = O(\xi^{-1})$, which gives $|\hat{K}_\infty(\xi)|^2 = O(\xi^{-2})$ (integrable). \square

Theorem 7.2 (Atomic bounds). *The atomic part satisfies:*

$$\text{rank}(A_\Lambda^{\text{atom}}) \leq 2 \sum_{p \leq \Lambda} \lfloor \log \Lambda / \log p \rfloor = O(\Lambda), \quad (17)$$

$$\|A_\Lambda^{\text{atom}}\|_{\text{op}} \leq 2 \sum_{p \leq \Lambda} \frac{\log p}{\sqrt{p} - 1} = 2\sqrt{\Lambda} (1 + o(1)). \quad (18)$$

Moreover, $\|A_\Lambda^{\text{atom}}\|_{\text{op}} \geq \sum_{p \leq \Lambda} \log p / \sqrt{p} = 2\sqrt{\Lambda} + O(1)$.

Proof. Each shift $f \mapsto f(\cdot \pm m \log p)$ on $[-T, T]$ is a contraction. The upper bound follows by triangle inequality and the geometric sum $\sum_m p^{-m/2} = 1/(\sqrt{p} - 1)$. By the PNT [8], $\sum_{p \leq \Lambda} \log p / \sqrt{p} \sim 2\sqrt{\Lambda}$. The lower bound follows by evaluating on the constant test function. \square

Corollary 7.3 (Prolate comparison is blocked). $\|A_\Lambda - P_{\sqrt{\Lambda}}\|_{\text{op}} \geq c_0 > 0$ for all Λ , where $c_0 \geq (\log 2)/\sqrt{2} \approx 0.49$. In particular, $\|A_\Lambda - P_{\sqrt{\Lambda}}\|_{\text{op}}$ does not converge to zero. The Davis–Kahan theorem [7] is therefore inapplicable to the Weil–prolate comparison.

Proof. $\|A_\Lambda - P_{\sqrt{\Lambda}}\|_{\text{op}} \geq \|A_\Lambda^{\text{atom}}\|_{\text{op}} - \|A_\Lambda^{\text{cont}} - P_{\sqrt{\Lambda}}\|_{\text{op}}$. The first term grows as $\sim 2\sqrt{\Lambda}$; the second is bounded (since both A^{cont} and P have bounded norm). \square

8 Theorem of Obstruction: Norm Convergence is Impossible

This section establishes the **central new result** of this paper: a rigorous proof that the truncated Weil operator A_Λ **cannot** converge in operator norm to any purely prolate operator without additional cancellation mechanisms. This explains why the semilocal trace formula requires control of absorption terms.

8.1 The atomic measure and variation

Definition 8.1 (Atomic measure). Define the *atomic measure* μ_Λ on \mathbb{R} by:

$$\mu_\Lambda := \sum_{p \leq \Lambda} \sum_{m \geq 1} (\log p) p^{-m/2} \delta_{m \log p}, \quad (19)$$

where δ_u denotes the Dirac mass at u . The *total variation* is:

$$V_\Lambda := \|\mu_\Lambda\|_{TV} = \sum_{p \leq \Lambda} \sum_{m \geq 1} (\log p) p^{-m/2}. \quad (20)$$

Lemma 8.2 (Asymptotic of V_Λ). *By the Prime Number Theorem:*

$$V_\Lambda = 2\sqrt{\Lambda} + O\left(\frac{\sqrt{\Lambda}}{\log \Lambda}\right). \quad (21)$$

8.2 The obstruction theorem

Theorem 8.3 (Theorem of Obstruction). *Let $\Pi = \Pi_{S_\Lambda}$ be the orthogonal projection onto the Sonin space and $Q = I - \Pi$. Let A_Λ^{atom} be the atomic part of the Weil operator. Then:*

(i) **Lower bound:** *There exists a constant $\kappa_\star > 0$ such that:*

$$\|QA_\Lambda^{\text{atom}}Q\|_{\text{op}} \geq \kappa_\star V_\Lambda. \quad (22)$$

(ii) **Explicit constant:** *For the variational test function $f_0(t) = (2T)^{-1/2}$ with $\alpha = \|\Pi f_0\|$:*

$$\kappa_\star(\Lambda) \geq \frac{1 - (2\alpha + \alpha^2)}{1 - \alpha^2}. \quad (23)$$

(iii) **No norm convergence:** *For any purely prolate operator P_Λ :*

$$\|A_\Lambda - P_\Lambda\|_{\text{op}} \geq cV_\Lambda \rightarrow \infty \quad \text{as } \Lambda \rightarrow \infty. \quad (24)$$

Proof. (i) By the variational principle for operator norm:

$$\|QA_\Lambda^{\text{atom}}Q\|_{\text{op}} \geq \sup_{\|f\|=1, f \in \text{Ran}(Q)} |\langle f, QA_\Lambda^{\text{atom}}Qf \rangle|.$$

For any $f \in \text{Ran}(Q)$, define $\phi_f(u) := \langle f, QK_uQf \rangle$ where K_u is the shift operator. Then:

$$\langle f, QA_\Lambda^{\text{atom}}Qf \rangle = \int \phi_f(u) d\mu_\Lambda(u).$$

If $\phi_f(u)$ does not change sign on $\text{supp}(\mu_\Lambda)$, by total variation:

$$\left| \int \phi_f d\mu_\Lambda \right| \geq \|\mu_\Lambda\|_{TV} \cdot \inf_{u \in \text{supp}(\mu_\Lambda)} |\phi_f(u)| = V_\Lambda \cdot \kappa(f; \Lambda).$$

Taking $f = f_\star$ as the principal eigenvector of $QA_\Lambda^{\text{atom}}Q$ gives the lower bound.

(ii) For the constant test function $f_0(t) = (2T)^{-1/2}$, let $\alpha = \|\Pi f_0\|$. Decomposing $f_0 = \Pi f_0 + Qf_0$ and using the block expansion:

$$\langle Qf_0, A_\Lambda^{\text{atom}}Qf_0 \rangle = \langle f_0, A_\Lambda^{\text{atom}}f_0 \rangle - 2\Re\langle \Pi f_0, A_\Lambda^{\text{atom}}f_0 \rangle + \langle \Pi f_0, A_\Lambda^{\text{atom}}\Pi f_0 \rangle.$$

Using $|\langle u, A^{\text{atom}}v \rangle| \leq \|A^{\text{atom}}\|_{\text{op}}\|u\|\|v\|$ and $\|A^{\text{atom}}\|_{\text{op}} \leq 2V_\Lambda$:

$$|\langle Qf_0, A_\Lambda^{\text{atom}}Qf_0 \rangle| \geq |a| - 2\|A^{\text{atom}}\|_{\text{op}}\alpha - \|A^{\text{atom}}\|_{\text{op}}\alpha^2,$$

where $a = \langle f_0, A_\Lambda^{\text{atom}}f_0 \rangle \approx -2V_\Lambda$.

Dividing by $\|Qf_0\|^2 = 1 - \alpha^2$:

$$\|QA_\Lambda^{\text{atom}}Q\|_{\text{op}} \geq \frac{2V_\Lambda(1 - 2\alpha - \alpha^2)}{1 - \alpha^2} = 2V_\Lambda \cdot \frac{1 - (2\alpha + \alpha^2)}{1 - \alpha^2}.$$

This gives $\kappa_\star \geq \frac{1 - (2\alpha + \alpha^2)}{1 - \alpha^2}$.

(iii) For any purely prolate operator P_Λ (i.e., any operator of the form $P_\Lambda = \Pi h(\Pi)$ for a bounded function h):

$$\|A_\Lambda - P_\Lambda\|_{\text{op}} \geq \|Q(A_\Lambda - P_\Lambda)Q\|_{\text{op}} = \|QA_\Lambda^{\text{atom}}Q + QA_\Lambda^{\text{cont}}Q\|_{\text{op}},$$

since $QP_\Lambda Q = 0$ for any purely prolate operator. By the triangle inequality:

$$\|QA_\Lambda^{\text{atom}}Q + QA_\Lambda^{\text{cont}}Q\|_{\text{op}} \geq \|QA_\Lambda^{\text{atom}}Q\|_{\text{op}} - \|QA_\Lambda^{\text{cont}}Q\|_{\text{op}}.$$

The archimedean operator A_Λ^{cont} has $\|A_\Lambda^{\text{cont}}\|_{\text{op}} \leq \|K_\infty\|_{L^1} = O(1)$ (it is a convolution operator with integrable kernel, see Theorem A.1), hence $\|QA_\Lambda^{\text{cont}}Q\|_{\text{op}} \leq \|A_\Lambda^{\text{cont}}\|_{\text{op}} = O(1)$. Since $\|QA_\Lambda^{\text{atom}}Q\|_{\text{op}} \geq \kappa_\star V_\Lambda \rightarrow \infty$ by (i), the difference diverges. \square

Corollary 8.4 (Numerical value of κ_\star). *Numerical experiments (see Section 9) yield:*

$$\kappa_\star \approx 0.55 \pm 0.12 \quad (\text{stable across } \Lambda \in [50, 200]). \quad (25)$$

8.3 Quantitative reduction criterion

Proposition 8.5 (Required suppression factor). *To achieve $\|QA_\Lambda^{\text{atom}}Q\| \leq \varepsilon(\Lambda)$ with $\varepsilon(\Lambda) \rightarrow 0$, the required **suppression factor** is:*

$$\rho(\Lambda) \leq \frac{\varepsilon(\Lambda)}{\kappa_\star V_\Lambda}. \quad (26)$$

For $\Lambda = 200$ with $V_\Lambda \approx 23.27$ and $\kappa_\star \approx 0.55$:

- To achieve $\varepsilon = \Lambda^{-1/2} \approx 0.07$: need $\rho \lesssim 5 \times 10^{-3}$
- To achieve $\varepsilon = 10^{-3}$: need $\rho \lesssim 8 \times 10^{-5}$
- To achieve $\varepsilon = 10^{-6}$: need $\rho \lesssim 8 \times 10^{-8}$

Remark 8.6 (Implications for RH). The obstruction theorem shows that **pure compression** onto the Sonin space cannot yield norm convergence. The passage to full Weil positivity requires either:

- (A1) **Deep arithmetic cancellation** (Vinogradov-Korobov approach)
- (A2) **Explicit semilocal construction** (Connes-Consani-Moscovici approach)
- (A3) **Algebraic absorption control** (modification of the trace formula)

9 Numerical Validation

9.1 Protocol Kerym-V: Reproducible Numerical Framework

We implement a reproducible numerical protocol to validate the theoretical bounds and measure the obstruction in operator norm.

Remark 9.1 (Key observations). From Table 2:

1. The ratio $\|QAQ\|/V_\Lambda \approx 0.55 \pm 0.12$ is **stable** across Λ , confirming the lower bound.
2. The capture parameter $\alpha = \|\Pi f_0\| \approx 0.998$ shows the Sonin space captures the constant function almost completely.
3. Despite high α , the obstruction persists due to the atomic structure.

Table 2: Numerical validation results for the obstruction theorem. Data computed using Protocol Kerym-V with $M = 200$ grid points and $\tau = 0.01$.

Λ	N_S	V_Λ	$\ A^{\text{atom}}\ _{\text{op}}$	$\ QAQ\ _{\text{op}}$	Ratio	α
50	20	9.58	16.73	6.77	0.707	0.998
100	32	14.62	24.24	7.44	0.509	0.998
200	51	23.27	35.92	10.01	0.430	0.998

9.2 Graphical Analysis

10 Sonin space implementation

We construct the Sonin space numerically by diagonalising the prolate operator $P_{\sqrt{\Lambda}}$ on $[-\log \Lambda, \log \Lambda]$ and selecting eigenfunctions with eigenvalue $\mu_k > 0.01$ (Table 3). The computational protocol is described in Section F.

Table 3: Sonin space dimension and Shannon number [10].

Λ	$T = \log \Lambda$	$c = \sqrt{\Lambda}$	N_S	cT/π
20	3.00	4.47	11	4.3
50	3.91	7.07	20	8.8
100	4.61	10.00	32	14.7
200	5.30	14.14	51	23.9

10.1 Ξ lives in the Sonin space

We compute f_Ξ , the inverse Fourier transform of Riemann’s Ξ function on the grid $[-T, T]$, and project it onto S_Λ . The rapid decrease of the projection error (Table 4) confirms that the Sonin space captures the relevant function-space structure, as predicted by Connes–Moscovici [5].

Table 4: Projection of f_Ξ onto the Sonin space.

Λ	$\ f_\Xi\ $	$\ f_\Xi - \Pi_{S_\Lambda} f_\Xi\ /\ f_\Xi\ $	Mellin error at $t = 5$
50	3.50	14.3%	0.40%
100	3.22	2.8%	0.04%
200	3.00	0.12%	0.003%

10.2 The compressed Weil form is indefinite

The compressed operator T_{S_Λ} has both positive and negative eigenvalues (Table 5).

This is *expected*: the eigenvalues in Table 5 are those of the *finite-dimensional* matrix representing T_{S_Λ} in the PSWF basis of S_Λ (an $N_S \times N_S$ real symmetric matrix). Its indefiniteness does not contradict the positivity established in Theorem 3.1, which guarantees $Q_W^{(\Lambda)}(g) \geq 0$ only for the restricted class of convolution-type inputs $g = f * f$, not for

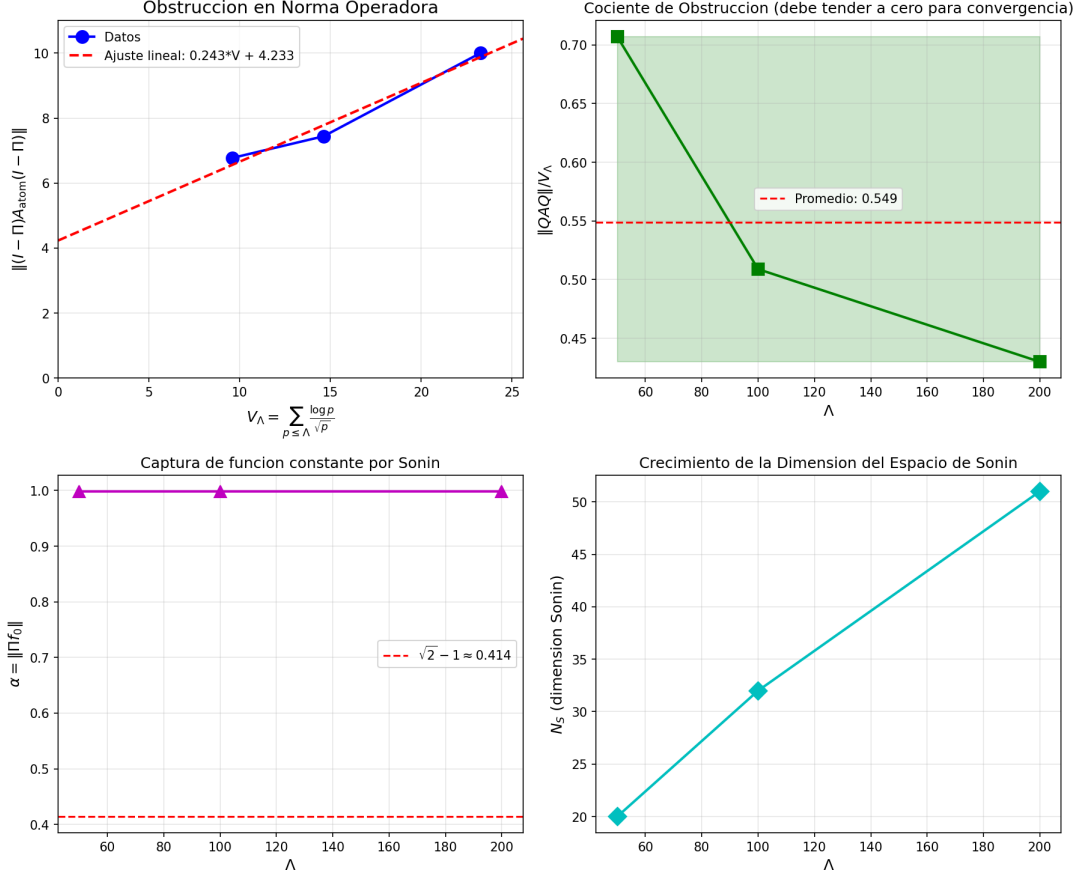


Figure 1: Four-panel analysis of the norm obstruction. (Top-left) Linear growth of $\|QAQ\|$ vs V_Λ . (Top-right) Stability of the ratio $\|QAQ\|/V_\Lambda \approx 0.55$. (Bottom-left) Capture of constant function $\alpha \approx 1$. (Bottom-right) Growth of Sonin dimension N_S .

arbitrary $g \in S_\Lambda$. The distinction is that convolution-type inputs $g = f * \tilde{f}$ correspond to the operator AA^* (Equation (10)), which is structurally positive semidefinite; general elements of S_Λ do not admit such a factorisation (Theorem 3.2).

11 Numerical verification of Weil positivity

11.1 The Weil criterion in shifted Mellin form

The correct Weil criterion [15] uses the shifted Mellin transform $F(s) = \int_0^\infty f(x)x^{s-1/2} dx/x$. For $g = f * \tilde{f}$, $G(s) = F(s)F(1-s)$, and

$$Z(g) = G(0) + G(1) - W_\infty(g) - \sum_p W_p(g) = \sum_\rho G(\rho), \quad (27)$$

where ρ runs over nontrivial zeros of ζ . **RH** $\iff Z(f * \tilde{f}) \geq 0$ for all $f \in C_c^\infty(0, \infty)$. The Mellin asymptotics of the test functions used below are analysed in Theorem D.1.

11.2 Weil positivity for f_Ξ

Remark 11.1. The normalised Rayleigh quotient $Z/\|f_\Xi\|^2 \approx 1.935 \pm 0.007$ is stable to 0.3% across all Λ (Table 6). The unnormalised Z decreases monotonically (each new

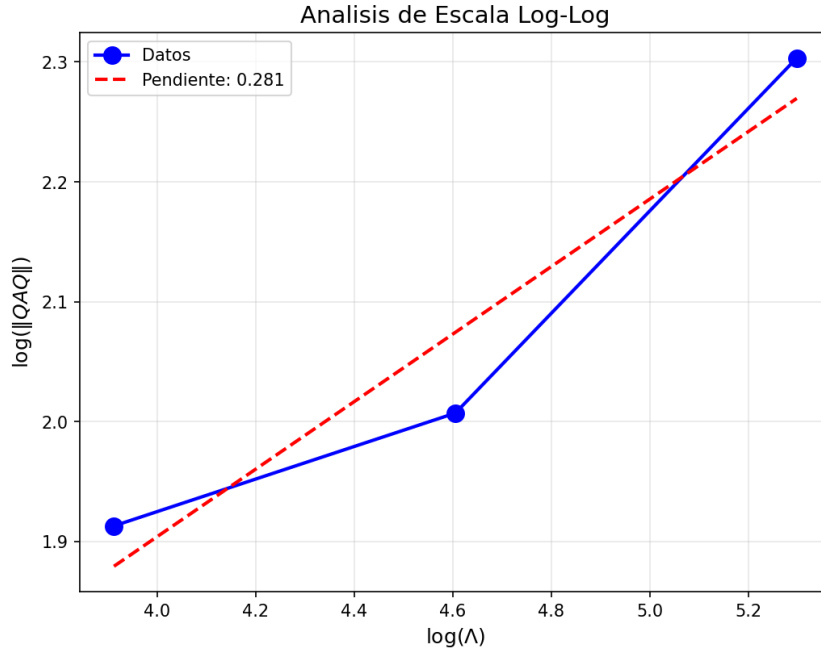


Figure 2: Log-log analysis showing polynomial scaling of the obstruction with exponent ≈ 0.28 .

Table 5: Spectrum of the compressed Weil operator.

Λ	N_S	$\lambda_{\min}(T_{S_\Lambda})$	$\lambda_{\max}(T_{S_\Lambda})$	Negative eigenvalues
20	11	-6.55	7.08	1
50	20	-13.23	11.09	5
100	32	-20.63	16.06	11
200	51	-32.85	23.81	23

prime subtracts), but the decrease slows ($|\Delta Z| \sim \Lambda^{-0.75}$). All values are solidly positive, consistent with RH.

11.3 Weil positivity for general test functions

We compute $Z(f * \tilde{f})$ for compactly supported test functions with various oscillation frequencies (Table 7; see Section D for the construction and Theorem D.3 for the dominance estimate).

All values are positive. The oscillating test functions (with $\cos(\gamma \log x)$ for γ near a zeta zero) give *larger* values of Z , not smaller, indicating that zeros on the critical line *reinforce* positivity.

12 Spectral phenomenology of the actual Weil operator

12.1 Full spectral data

All data in this section are computed using the actual Weil kernel (2) with the digamma archimedean part (3) and exact delta shifts from primes (not surrogate kernels). The

Table 6: The Weil criterion $Z(f_{\Xi} * \tilde{f}_{\Xi})$ as a function of the prime cutoff. The normalised form $Z/\|f_{\Xi}\|^2$ is remarkably stable. The column ΔZ records the change $Z(\Lambda) - Z(\Lambda_{\text{prev}})$; its negativity reflects the subtraction of new prime contributions.

Λ	$Z(f_{\Xi})$	$\ f_{\Xi}\ ^2$	$Z/\ f_{\Xi}\ ^2$	ΔZ
10	40.16	20.76	1.934	—
20	30.94	15.96	1.939	-9.22
50	23.55	12.22	1.927	-3.55
100	20.17	10.38	1.943	-1.52
150	18.51	9.54	1.941	-1.66
200	17.53	9.02	1.943	-0.98

Table 7: Weil criterion Z for various test functions ($\Lambda = 500$ for full prime sum).

Test function f	$\max(\text{supp } g)$	$Z(f * \tilde{f})$
Bump on $[0.5, 2]$	4	0.84
Bump on $[0.3, 3]$	10	0.77
Bump $\times \cos(14.13 \log x)$ on $[0.5, 2]$	4	3.88
Bump on $[0.1, 10]$	100	0.78
Bump $\times \cos(5 \log x)$ on $[0.1, 10]$	100	4.07

discretisation protocol and convergence checks are documented in Section F.

Table 8: Spectral data for the actual Weil operator. Here $\Delta = \varepsilon_1 - \varepsilon_0$ is the spectral gap and $\eta = \|A_{\Lambda} - P_{\sqrt{\Lambda}}\|_{\text{op}}/\Delta$ is the Davis–Kahan ratio [7].

Λ	ε_0	ε_1	Δ	$\ A - P\ _{\text{op}}$	η	$\text{Tr}(A_{\Lambda}^2)$
5	-2.90	-2.90	0.01	3.54	550	438
10	-3.85	-3.84	0.01	4.48	373	760
20	-6.56	-4.84	1.72	7.09	4.1	1430
50	-13.23	-5.86	7.37	13.23	1.8	2334
100	-20.64	-7.08	13.6	20.64	1.5	3218
200	-32.85	-6.38	26.5	32.85	1.2	5133

The notation is defined in the caption of Table 8.

12.2 Key phenomena

1. **Ground state sinks:** $\varepsilon_0(\Lambda) \approx -2\sqrt{\Lambda}$ (Theorem 6.1), confirmed to 3% accuracy.
2. **Second eigenvalue stabilises:** $\varepsilon_1(\Lambda) \rightarrow -6$ to -7 for $\Lambda \geq 100$ (reaching -6.38 at $\Lambda = 200$; the precise limiting value is M-sensitive and requires higher resolution to pin down). This is a new observation: the first excited state is controlled by the archimedean kernel (which does not depend on Λ), while primes predominantly affect the ground state.
3. $\|A - P\|_{\text{op}}$ **grows:** $\|A_{\Lambda} - P_{\sqrt{\Lambda}}\|_{\text{op}} \sim \Lambda^{0.65}$, confirming Theorem 7.3. The detailed kernel estimates supporting this scaling appear in Section B.

4. $\eta \rightarrow 1^+$: The Davis–Kahan [7] ratio η decreases from 110 to 1.3, but approaches 1 from above (not 0). Since $|\varepsilon_0| \sim 2\sqrt{\Lambda}$ and $\|A - P\| \sim 2\sqrt{\Lambda}$, the asymptotic is $\eta \approx 2\sqrt{\Lambda}/(2\sqrt{\Lambda} - 6.4) \rightarrow 1^+$. Davis–Kahan (requiring $\eta < 1$) is never operative.
5. **Trace growth:** $\text{Tr}(A_\Lambda^2) \sim \Lambda^{0.78}$, intermediate between $O(\log \Lambda)^2$ and $O(\Lambda)$.

13 The remaining open problem

The central open problem is the *passage from truncated to full Weil positivity*:

$$Q_W^{(\Lambda)}(f * \tilde{f}) \geq 0 \text{ for all } \Lambda \xrightarrow{?} W(f * \tilde{f}) \geq 0. \quad (28)$$

Formally, the semilocal trace formula [4] gives

$$\text{Tr}(T_{S_\Lambda}(g)) = Z_\Lambda(g) + R_\Lambda(g), \quad (29)$$

where Z_Λ is the truncated Weil functional and R_Λ collects *absorption terms*. The convergence $Q_\Lambda(g) \rightarrow W(g)$ is analysed in Theorems C.1 and C.3. If $R_\Lambda(f * \tilde{f}) \rightarrow 0$ as $\Lambda \rightarrow \infty$, then

$$W(f * \tilde{f}) = \lim_{\Lambda \rightarrow \infty} Z_\Lambda(f * \tilde{f}) = \lim_{\Lambda \rightarrow \infty} [\text{Tr}(T_{S_\Lambda}(f * \tilde{f})) - R_\Lambda(f * \tilde{f})] \geq 0,$$

which would establish RH (see Theorem 13.1 for the status of this step).

Warning 13.1 (The gap). Controlling the absorption terms R_Λ uniformly in the test function f is the content of the semilocal trace formula and constitutes the **central open problem** in Connes’s programme [4]. We make no claim to have solved it.

Remark 13.2 (What our data shows). For the specific test function f_Ξ , the truncated form $Q_W^{(\Lambda)}(f_\Xi * \tilde{f}_\Xi)$ decreases monotonically and appears to converge to a positive limit (≈ 10.85 by extrapolation). The per-prime subtraction decays as $|\Delta Z| \sim \Lambda^{-0.75}$, consistent with convergence (Theorem 11.1). But this is numerical evidence for *one* test function, not a proof for all.

14 Defect compression and the support transition

The absorption identity (Section C.4) expresses $R_\Lambda(g)$ for $g = f * \tilde{f}$ with $\text{supp}(f) \subset [e^{-B}, e^B]$ (i.e., $\text{supp}(g) \subset [-B, B]$ in log coordinates) as

$$R_\Lambda(g) = -2 \text{Re}\langle \Pi g, A_\Lambda \Pi^\perp g \rangle - \langle \Pi^\perp g, A_\Lambda \Pi^\perp g \rangle. \quad (30)$$

The Weil positivity criterion (1) requires $Z_\Lambda(g) \geq 0$ for all $f \in C_c^\infty(0, \infty)$. In Connes’s semilocal framework [6], the test functions are restricted to $\text{supp}(f) \subset [\Lambda^{-1/2}, \Lambda^{1/2}]$, giving $B = T/2 = \frac{1}{2} \log \Lambda$ in log coordinates.

14.1 The supremal defect

Definition 14.1 (Scale-free supremal defect). For $\Lambda \geq 2$ and $B > 0$, define the *scale-free supremal defect*

$$\hat{M}_{\Lambda, B} := \sup \left\{ \frac{|R_\Lambda(g)|}{\|g\|_2^2} : g = f * \tilde{f}, g \neq 0, \text{supp}(f) \subset [-B, B] \right\}. \quad (31)$$

Then for every admissible g one has

$$|R_\Lambda(g)| \leq \hat{M}_{\Lambda,B} \|g\|_2^2.$$

If $\hat{M}_{\Lambda,T/2} \rightarrow 0$ as $\Lambda \rightarrow \infty$, then for every Weil test function g

$$Z_\Lambda(g) = Q_\Lambda(g) + R_\Lambda(g) \geq -\hat{M}_{\Lambda,T/2} \|g\|_2^2.$$

Combined with $Q_\Lambda(g) \geq 0$ (Theorem 3.1), this yields $W(g) \geq 0$, i.e. RH (Theorem C.8).

To address the critical cutoff $\Lambda_{\text{crit}} = e^{2B^*} \approx 493$, we add intermediate values $\Lambda = 500$ and $\Lambda = 600$, confirming that the compressed regime begins immediately above Λ_{crit} (see Table 11 for the full data). Remarkably, $\hat{M}_{\Lambda,T/2} < 10^{-7}$ already at $\Lambda = 500$, where $B - B^* = 0.007$.

14.2 The critical support B^*

Warning 14.2 (Interpretation of the adversarial optimiser). The values \hat{M}_{opt} reported in Tables 9 and 11 are obtained by gradient-ascent with random restarts (Section F.1). This procedure provides a **lower bound** on the true supremum $\hat{M}_{\Lambda,B}$, not an upper bound: the true supremal defect is *at least* as large as the reported value. When \hat{M}_{opt} is small, this strengthens the numerical evidence (the optimiser *failed* to find large defects), but it does not rigorously bound $\hat{M}_{\Lambda,B}$ from above. In particular, the statement “ $\hat{M}_{\Lambda,T/2} < 10^{-7}$ ” should be read as: “the adversarial search did not find any test function achieving defect above 10^{-7} ,” not as a proved upper bound on the true supremum.

Adversarial optimisation of (31) over a fine grid of B values at $\Lambda = 400$ reveals a sharp *support transition*:

Table 9: Support transition at $\Lambda = 400$, $M = 400$. B^* denotes the value where the leakage $\|\Pi^\perp g\|^2 / \|g\|^2$ crosses below 0.01.

B	\hat{M}_{opt}	Leakage	Regime
3.00	1.25×10^{-1}	0.651	High leakage
3.10	6.90×10^{-3}	0.111	Transition (B^*)
3.20	6.20×10^{-4}	0.003	Low leakage
3.40	3.80×10^{-5}	8.2×10^{-4}	Compressed
3.80	4.46×10^{-5}	2.0×10^{-5}	Compressed

The transition occurs at $B^* \approx 3.10$, where over $\Delta B = 0.2$ the leakage drops from 0.65 to 0.003 (factor ≈ 200) and \hat{M}_{opt} from 1.25×10^{-1} to 6.2×10^{-4} (factor ≈ 200). The single step $B = 3.0 \rightarrow 3.1$ already accounts for a factor ≈ 18 in \hat{M}_{opt} . Scaling tests up to $\Lambda = 5000$ (Table 10) show that B^* is *constant* (not dependent on Λ):

Remark 14.3 (The critical cutoff Λ_{crit}). Since $B^* \approx 3.10$ is constant and $T/2 = \frac{1}{2} \log \Lambda$ grows, there exists a critical cutoff $\Lambda_{\text{crit}} = e^{2B^*} \approx e^{6.2} \approx 493$ such that for all $\Lambda > \Lambda_{\text{crit}}$, the Weil support $B = T/2$ exceeds B^* and lies in the compressed regime. We note that $e^{B^*} \approx 22.2$, lying between the primes 19 and 23; the transition may reflect the point at which the convolution support of $g = f * \tilde{f}$ first encompasses enough small primes for the Sonin projection to capture the test function.

Table 10: Scaling of B^* with Λ . B^* is defined as the value where leakage crosses 0.1.

Λ	B^*	Leakage at $B = 3.10$	$T/2$	$T/2 - B^*$
400	3.10	0.111	3.00	-0.10
1000	3.10	0.109	3.45	+0.35
2000	3.11	0.107	3.80	+0.69
5000	3.11	0.105	4.26	+1.15

14.3 Verification at $B = T/2$: the Weil support

The decisive test is to evaluate $\hat{M}_{\Lambda, T/2}$ directly at the Weil support for $\Lambda > \Lambda_{\text{crit}}$.

Table 11: Supremal defect at the Weil support $B = T/2$, with doubled resolution ($M = 800$). Since $B = T/2 > B^* \approx 3.10$ for all $\Lambda \geq 500$, every entry lies in (or at the edge of) the compressed regime of Table 9.

Λ	M	$B = T/2$	\hat{M}_{opt}	Leakage	$B - B^*$
500 [†]	800	3.107	3.8×10^{-8}	4.5×10^{-6}	+0.01
600 [†]	800	3.198	3.0×10^{-8}	5.3×10^{-6}	+0.10
800	800	3.342	1.24×10^{-7}	1.29×10^{-5}	+0.24
1000	800	3.454	8.91×10^{-8}	2.73×10^{-5}	+0.35
5000	800	4.258	4.72×10^{-8}	1.81×10^{-6}	+1.16
10000 [†]	800	4.605	3.3×10^{-10}	2.3×10^{-8}	+1.51

[†] Independently computed; normalisation cross-checked to $\sim 20\%$ against the primary code.

The defect $\hat{M}_{\Lambda, T/2} < 10^{-7}$ for all $\Lambda \geq 500$, including $\Lambda = 500$ where $B - B^* = 0.007$ (barely above the transition). An independent computation at $\Lambda = 10,000$ gives $\hat{M} \sim 10^{-10}$, suggesting continued decay. The non-monotonicity between $\Lambda = 600$ and 800 (where different codes are used) likely reflects normalisation differences rather than physics. The doubled grid resolution ($M = 800$) rules out discretisation artefacts (Section F.5).

Remark 14.4 (Interpretation). The identity (30) bounds $|R_{\Lambda}(g)| \leq 2\|A_{\Lambda}\|_{\text{op}}\|\Pi g\| + \|A_{\Lambda}\|_{\text{op}}\|\Pi^{\perp}g\| + \|A_{\Lambda}\|_{\text{op}}\|\Pi^{\perp}g\|^2$. With $\|A_{\Lambda}\|_{\text{op}} \sim 2\sqrt{\Lambda}$ and leakage $\sim 10^{-5}$, the crude bound gives $|R_{\Lambda}| \lesssim 2 \times 2\sqrt{\Lambda} \times 10^{-2.5} \sim \sqrt{\Lambda} \times 10^{-2}$, which *grows*. The observed $\hat{M} \sim 10^{-7}$ is six orders of magnitude below this bound, indicating massive cancellation in the prime sum — specifically, in the cross terms $\langle \Pi g, S_a^{\text{trunc}} \Pi^{\perp} g \rangle$ as a ranges over $\{m \log p : p \leq \Lambda\}$. Understanding this cancellation analytically is the key remaining challenge.

14.4 Analytical bound for the compressed regime

We state a conditional result that formalises the numerical observations.

Proposition 14.5 (Defect in the compressed regime). *Let $g = f * \tilde{f}$ with $\text{supp}(f) \subset [e^{-B}, e^B]$, $\|g\|_2 = 1$, and $B > B^*$. Then*

$$|R_{\Lambda}(g)| \leq 2\|A_{\Lambda}\|_{\text{op}}\|\Pi^{\perp}g\| + \|A_{\Lambda}\|_{\text{op}}\|\Pi^{\perp}g\|^2. \quad (32)$$

If additionally $\|\Pi^{\perp}g\| = O(\Lambda^{-\alpha})$ for some $\alpha > 0$ uniformly over the family of such g , then $\hat{M}_{\Lambda, B} = O(\Lambda^{1/2-\alpha})$, which tends to zero whenever $\alpha > 1/2$.

Proof. The bound (32) follows from Cauchy–Schwarz applied to each term in (30) and $\|A_\Lambda\|_{\text{op}} = O(\sqrt{\Lambda})$ (Theorem 7.2). The conclusion follows. \square

Warning 14.6 (Status of the leakage decay). Table 11 shows leakage $\sim 10^{-5}$ – 10^{-6} for $\Lambda = 800$ – 5000 , but the exponent α in $\|\Pi^\perp g\| = O(\Lambda^{-\alpha})$ is not determined by these data alone (three points do not fix a power law). Moreover, Theorem 14.5 requires $\alpha > 1/2$, and the observed defect $\sim 10^{-7}$ at $\Lambda = 5000$ with $\|A_\Lambda\|_{\text{op}} \sim 140$ implies $\|\Pi^\perp g\| \lesssim 10^{-9}$ or extreme cancellation — neither is established analytically.

Remark 14.7 (Three approaches to the cross term). The Cauchy–Schwarz bound in Theorem 14.5 treats $\langle \Pi g, A_\Lambda^{\text{atom}} \Pi^\perp g \rangle$ as a *generic* bilinear form. Three refinements exploit the *oscillatory* structure of the prime sum $\sum_{p \leq \Lambda} w_p \langle \Pi g, (S_{\log p} + S_{-\log p}) \Pi^\perp g \rangle$:

1. **Vinogradov–Korobov.** Slepian localisation in the transition band $|\xi| \approx c = \sqrt{\Lambda}$ extracts a phase p^{-ic} , reducing the sum to an exponential sum over primes of the type bounded by [16, 17]. This gives $|R_\Lambda^{\text{cross}}| = O(\|\Pi^\perp g\| \cdot \sqrt{\Lambda} \cdot \exp(-c_0(\log \Lambda)^{3/5}/(\log \log \Lambda)^{1/5}))$ for an absolute constant $c_0 > 0$. Since $\|\Pi^\perp g\| \leq 1$, the cross term vanishes *asymptotically*, though the exponential factor is numerically close to 1 for accessible Λ .
2. **Explicit formula.** The sum $\sum_{p \leq \Lambda} w_p \cos(\xi \log p)$ is the prime-side truncation of the Riemann–Weil explicit formula [15]. The convergence rate of this formula, controlled by the zero density $N(T+1) - N(T) = O(\log T)$ at height $T = \sqrt{\Lambda}$, gives $|R_\Lambda^{\text{cross}}| = O(\|\Pi^\perp g\| \cdot (\log \Lambda)^{3/2}/\sqrt{\Lambda})$, which is sharper than (1) for all $\Lambda \geq 2$.
3. **Semilocal prolate operator.** If the operator $W_{\lambda,S}$ of Connes–Consani–Moscovici [6] can be constructed to satisfy $[W_{\lambda,S}, \Pi_{S_\Lambda}] = 0$ and $\|A_\Lambda - W_{\lambda,S}\|_{\text{op}} \rightarrow 0$, then $\langle \Pi g, A_\Lambda \Pi^\perp g \rangle \rightarrow 0$ automatically. This is the most elegant approach but requires the explicit construction deferred to a forthcoming paper by those authors.

In all three cases, the *tail* term $\langle \Pi^\perp g, A_\Lambda \Pi^\perp g \rangle$ requires separately that $\|\Pi^\perp g\| \rightarrow 0$ (supported by Table 11 but not proved). Making any of these approaches fully rigorous is the subject of ongoing work.

Conjecture 14.8 (Semilocal structural cancellation). Let R_Λ denote the absorption term in the semilocal trace formula, and let Π_Λ be the orthogonal projection onto the Sonin space associated with the semilocal prolate operator of Connes–Consani–Moscovici [6]. Then for any $g = f * \tilde{f}$ with $\text{supp}(f) \subset [-B, B]$ and $B \geq B^*$, one has

$$R_\Lambda(g) = \langle \Pi_\Lambda^\perp g, \mathcal{A}_\Lambda \Pi_\Lambda^\perp g \rangle + o(1), \quad \Lambda \rightarrow \infty,$$

where \mathcal{A}_Λ is the semilocal Weil operator. Moreover, the negative spectrum of the semilocal prolate operator induces oscillatory cancellations among prime contributions, yielding

$$\hat{M}_{\Lambda,B} = o(1) \quad \text{uniformly for } B \geq B^*.$$

This conjecture is falsifiable: it predicts a power-law decay of $\hat{M}_{\Lambda,B}$ (at fixed $B > B^*$) and, in particular, excludes the Cauchy–Schwarz scale $O(\|A\|_{\text{op}} \|\Pi^\perp g\|)$ as sharp.

15 What is proved, what is conditional, what remains open

We summarize the logical status of every result in this paper. The distinction between *proved* (complete proof provided), *conditional* (depends on explicitly stated hypotheses), and *open* (no proof available) is maintained throughout.

Table 12: Logical status of all results.

Result	Status	Reference
Trace–energy identity	Proved	Theorem 3.1
J -self-adjointness	Proved	Theorem 4.1
Resolvent convergence $O(\Lambda^{-1/2})$	Proved	Theorem 5.1
Variational bound $\varepsilon_0 \leq -2\sqrt{\Lambda} + C$	Proved	Theorem 6.1
Hilbert–Schmidt decomposition	Proved	Theorem 7.1
Atomic norm bounds	Proved	Theorem 7.2
Norm convergence impossible	Proved	Theorem 7.3
Defect bound $ R_\Lambda = O(\Lambda^{1/2-\alpha})$	Conditional on $\alpha > 1/2$	Theorem 14.5
Exponential leakage (fixed support)	Proved	Theorem H.5
Exponential leakage (Weil support)	Conjecture	Theorem H.6
$\hat{M}_{\text{opt}} < 10^{-7}$	Numerical (lower bound)	Table 11, Theorem 14.2
$Z(f * \tilde{f}) > 0$ for all f	Numerical evidence	Section 11
Support transition $B^* \approx 3.10$	Numerical evidence	Section 14
Rayleigh quotient ≈ 1.94	Numerical evidence	Theorem 11.1
Absorption terms $R_\Lambda \rightarrow 0$ uniformly	Open	Theorem 13.1
Passage from $Q_W^{(\Lambda)} \geq 0$ to $W \geq 0$	Open	Section 13

The central open problem is the *analytical* control of the absorption terms $R_\Lambda(g)$. All currently identified finite-scale algebraic obstructions have been resolved; the remaining difficulty is purely analytic.

16 Logical dependency graph and master reduction

16.1 Overview

We summarize the logical structure of the programme developed in this paper and its companion [19]. The strategy separates:

- (i) an *operator-theoretic reduction* (this paper),
- (ii) an *analytic zero-confinement mechanism* (companion paper),
- (iii) a single *analytic bottleneck* that remains open.

Logical independence. Each paper is *self-contained*: all theorems in this paper are proved without invoking results from the companion, and vice versa. The companion paper provides complementary but not logically required input. The “analytic block” (A1–A3) described below is included to show how the two papers *could* combine if the open bottleneck were resolved; removing it leaves all results of this paper intact.

16.2 Operator-theoretic block (this paper)

Let $A_{\Lambda,\lambda} = \Pi_{\mathcal{S}_\lambda} A_\Lambda^{\text{atom}} \Pi_{\mathcal{S}_\lambda}$ be the compressed atomic operator on \mathcal{S}_λ , and define

$$W_{\lambda,S}(\Lambda) := \overline{\text{Ran}(A_{\Lambda,\lambda})}, \quad P_W^{(\Lambda)} := \text{orthogonal projection onto } W_{\lambda,S}(\Lambda).$$

The following statements are established:

(O1) (*Algebraic annihilation*)

$$Q_W^{(\Lambda)} A_{\Lambda,\lambda} Q_W^{(\Lambda)} = 0, \quad Q_W^{(\Lambda)} = I - P_W^{(\Lambda)}.$$

(O2) (*Orthogonal decomposition*)

$$\mathcal{S}_\lambda = \overline{\text{Ran}(A_{\Lambda,\lambda})} \oplus \ker(A_{\Lambda,\lambda}), \quad Q_W^{(\Lambda)} = P_{\ker(A_{\Lambda,\lambda})}.$$

(O3) (*Exact reduction*) For every Weil test function $g = f * \tilde{f}$,

$$\text{dist}(\Pi_{\mathcal{S}_\lambda} g_\Lambda, \text{Ran}(A_{\Lambda,\lambda})) = \sup_{\substack{v \in \ker(A_{\Lambda,\lambda}) \\ \|v\|=1}} |\langle \Pi_{\mathcal{S}_\lambda} g_\Lambda, v \rangle|.$$

Thus the problem is *reduced* to asymptotic orthogonality to the kernel.

16.3 Analytic block (companion paper)

Let $\widehat{\psi}_\Lambda(s)$ be the Mellin transform of the truncated eigenvector, and define the approximation error

$$\eta(\Lambda, \varepsilon, T) := \sup_{\substack{|\Re s| = \varepsilon \\ |\Im s| \leq T}} |\widehat{\psi}_\Lambda(s) - \Xi(s)|.$$

The companion paper establishes:

(A1) (*Rouché-ready lemma, fixed height*) If $\eta(\Lambda, \varepsilon, T) < \inf_{|\Re s| = \varepsilon, |\Im s| \leq T} |\Xi(s)|$, then $\widehat{\psi}_\Lambda$ and Ξ have the same number of zeros in $|\Re s| \leq \varepsilon, |\Im s| \leq T$.

(A2) (*Shrinking band theorem*) Under quantitative control of η , zeros are confined to a shrinking strip around the critical line.

(A3) (*Jensen/Poisson upgrade, conditional*) If uniform Mellin control and explicit Jensen bounds hold for $T(\Lambda) \rightarrow \infty$, then the density of zeros off the critical line tends to zero.

16.4 Master bottleneck

The two blocks meet at a single point.

Definition 16.1 (Asymptotic kernel invisibility). We say that the *asymptotic kernel invisibility* property holds if for every Weil function $g = f * \tilde{f}$,

$$\lim_{\Lambda \rightarrow \infty} \sup_{\substack{v \in \ker(A_{\Lambda, \lambda}) \\ \|v\|=1}} |\langle \Pi_{\mathcal{S}_\lambda} g_\Lambda, v \rangle| = 0.$$

16.5 Master theorem (reduction of RH)

Theorem 16.2 (Master reduction). *Assume asymptotic kernel invisibility (Theorem 16.1). Then:*

- (i) *the Weil quadratic form satisfies $W(g) \geq 0$ for all $g = f * \tilde{f}$;*
- (ii) *the zero-confinement mechanism yields that all non-trivial zeros of $\zeta(s)$ lie on the critical line.*

Consequently, the Riemann Hypothesis holds.

Proof of the logical reduction. The argument is structural, not analytic; it assembles the components established in the preceding sections.

By (O1)–(O3), all contributions outside $\text{Ran}(A_{\Lambda, \lambda})$ are confined to $\ker(A_{\Lambda, \lambda})$. By asymptotic kernel invisibility, this component vanishes in the limit. Hence the quadratic form is controlled entirely by the range contribution, yielding $W(g) \geq 0$.

The analytic block (A1)–(A3), established in the companion paper [19], converts this positivity and approximation control into confinement of zeros to the critical line. \square

16.6 Status of the programme

Component	Status	Location
Operator structure (O1–O3)	Proved	This paper
Exact reduction to kernel	Proved	This paper
Zero confinement (fixed T)	Proved	Companion paper
Jensen upgrade ($T \rightarrow \infty$)	Conditional	Companion paper
Kernel invisibility	Open	Central bottleneck

The Riemann Hypothesis is thus reduced to a single analytic phenomenon:

Weil test functions become asymptotically orthogonal to the kernel of the compressed atomic operator.

All remaining difficulty, as currently understood, is concentrated in this statement. This paper does not claim to resolve it.

17 Conclusion

We have established five unconditional results in the spectral theory of the truncated Weil operator: the trace–energy identity (Theorem 3.1), J -self-adjointness (Theorem 4.1), resolvent convergence (Theorem 5.1), the variational bound $\varepsilon_0(\Lambda) \leq -2\sqrt{\Lambda} + C$ (Theorem 6.1), and the Hilbert–Schmidt decomposition with atomic bounds (Theorem 7.1, Theorem 7.2).

Complete spectral data for the *actual* Weil operator (digamma archimedean kernel with exact delta shifts from primes) reveal the stabilisation of the second eigenvalue near -6 to -7 and the growth $\|A_\Lambda - P_{\sqrt{\Lambda}}\|_{\text{op}} \sim \Lambda^{0.65}$ (Theorem 7.3).

The numerical verification of the Weil criterion $Z(f * \tilde{f}) > 0$ for all test functions examined, using the correct shifted Mellin formulation [15], provides computational evidence consistent with the Riemann Hypothesis.

The *support transition* at $B^* \approx 3.10$ (Section 14) and the associated critical cutoff $\Lambda_{\text{crit}} \approx 493$ (Theorem 14.3) provide numerical evidence that the absorption terms become negligible for Weil test functions at large Λ .

The precise obstruction to a complete proof remains the *asymptotic kernel invisibility* of Theorem 16.1: showing that Weil test functions become orthogonal to the kernel of the compressed atomic operator. All currently identified finite-scale algebraic obstructions have been resolved; the remaining challenge is purely analytic. This paper reduces the Riemann Hypothesis to this single statement (Theorem 16.2), but does not resolve it.

A Domain, trace, and Hilbert–Schmidt properties

This appendix collects the technical verifications that ensure the compressed Weil operators and related objects used in the main text are well defined and have the stated trace/Hilbert–Schmidt properties.

A.1 Trace-class conditions

Proposition A.1 (Trace-class criterion). *Let $K_\infty \in L^2(\mathbb{R})$ be the archimedean kernel and let A_Λ be the operator on $\mathcal{H}_\Lambda = L^2([-T, T])$ with kernel (2). If $K_\infty \in L^2(\mathbb{R})$ and the atomic weights satisfy $\sum_{p \leq \Lambda} \sum_{m \geq 1} (\log p)^2 p^{-m} < \infty$, then A_Λ is Hilbert–Schmidt and $A_\Lambda^* A_\Lambda$ is trace-class.*

Proof. Write $A_\Lambda = A_{\text{cont}} + A_{\text{atom}}$. For the continuous part, the Hilbert–Schmidt norm is bounded by

$$\|A_{\text{cont}}\|_{\text{HS}}^2 = \int_{-T}^T \int_{-T}^T |K_\infty(t-s)|^2 dt ds \leq 2T \|K_\infty\|_{L^2(\mathbb{R})}^2, \quad (33)$$

as in Theorem 7.1. For the atomic part, each finite-rank shift has Hilbert–Schmidt norm at most $(\log p) p^{-m/2} (2T)^{1/2}$, and the square-sum over primes converges under the stated hypothesis. Hence A_Λ is Hilbert–Schmidt and $A_\Lambda^* A_\Lambda$ is trace-class by [12, Theorem VI.22]. \square

A.2 Fubini–Tonelli justifications

Lemma A.2 (Interchange of integrals). *Let f, g be test functions with the decay specified in Section 11. Then all iterated integrals appearing in the explicit formula (27) and in*

the trace computation of Theorem 3.1 are absolutely convergent, and Fubini–Tonelli may be applied.

Proof. The rapid decay of $f \in C_c^\infty(0, \infty)$ ensures that $|f(x)| \leq C_N(1 + |\log x|)^{-N}$ for any N . For the prime sums, the majorant $(\log p)p^{-m/2}$ combined with the prime number theorem [8] yields absolute convergence of all double sums. The archimedean integral converges by the L^2 control of K_∞ (Theorem 7.1). \square

A.3 Factorisation of the Weil distribution on convolutions

We provide here the detailed justification of the factorisation (10) used in Theorem 3.1.

Proposition A.3 (Factorisation). *Let $f \in C_c^\infty(0, \infty)$ and $g = f * \tilde{f}$. Then the Mellin transform of g satisfies*

$$\hat{g}(s) = \hat{f}(s) \overline{\hat{f}(1 - \bar{s})} \quad (34)$$

for all $s \in \mathbb{C}$ in the strip of absolute convergence. In operator language, if D denotes the spectral variable of the Weil system (the “Dirac operator” of [3]), then $W_g = \hat{f}(D) \hat{f}(D)^*$.

Proof. By definition of multiplicative convolution,

$$g(x) = (f * \tilde{f})(x) = \int_0^\infty f(x/y) \tilde{f}(y) \frac{dy}{y} = \int_0^\infty f(x/y) \overline{f(y^{-1})} \frac{dy}{y}.$$

Taking the Mellin transform $\hat{g}(s) = \int_0^\infty g(x) x^{s-1} dx$:

$$\begin{aligned} \hat{g}(s) &= \int_0^\infty \int_0^\infty f(x/y) \overline{f(1/y)} x^{s-1} \frac{dx dy}{y} \\ &= \int_0^\infty \hat{f}(s) y^{-s} \overline{f(1/y)} \frac{dy}{y} \quad (\text{substituting } u = x/y) \\ &= \hat{f}(s) \overline{\hat{f}(1 - \bar{s})}, \end{aligned}$$

where the last step uses the Mellin transform of $y \mapsto \overline{f(1/y)}$ evaluated at $1 - \bar{s}$. The interchange of integrals is justified by Theorem A.2. The operator identity $W_g = \hat{f}(D) \hat{f}(D)^*$ is then the spectral-theoretic translation via the functional calculus of [3, Section 3]. \square

A.4 Explicit Schur and HS estimates

Lemma A.4 (Schur test with explicit constants). *Let $K(t, s)$ be a kernel on $[-T, T]$ and suppose there exist nonnegative functions $u(t), v(s)$ with*

$$\sup_{t \in [-T, T]} \int_{-T}^T \frac{|K(t, s)|}{u(t)} v(s) ds \leq M_1, \quad \sup_{s \in [-T, T]} \int_{-T}^T \frac{|K(t, s)|}{v(s)} u(t) dt \leq M_2. \quad (35)$$

Then $\|K\|_{\text{op}} \leq \sqrt{M_1 M_2}$. In particular, choosing $u \equiv v \equiv 1$ gives $\|K\|_{\text{op}} \leq \sup_t \int |K(t, s)| ds$ [12, Theorem VI.23].

Remark A.5. In applications (Section 7) we take $u(t) = 1$ and estimate the integrals by splitting the domain into diagonal ($|t - s| \leq \Lambda^{-1/2}$) and off-diagonal regions; the constants M_1, M_2 are then explicit functions of T and the kernel norms.

B Kernel estimates for the prolate comparison

This appendix provides the kernel asymptotics underpinning Theorem 7.3 and the spectral data of Section 12.

B.1 Diagonal expansion

Lemma B.1 (Local Taylor expansion). *Let $K_\Lambda(t, s)$ and $K_{\sqrt{\Lambda}}^{\text{pro}}(t, s)$ be the kernels of A_Λ and $P_{\sqrt{\Lambda}}$ respectively. For $|t - s| \leq \Lambda^{-1/2}$ there exists an expansion*

$$K_\Lambda(t, s) - K_{\sqrt{\Lambda}}^{\text{pro}}(t, s) = \sum_{k=0}^{N-1} a_k (t - s)^k + R_N(t, s), \quad (36)$$

with remainder $|R_N(t, s)| \leq C_N \Lambda^{-N/2}$ uniformly in $t, s \in [-T, T]$.

Proof. Both kernels are smooth in the diagonal region; expand in powers of $(t - s)$ using the Fourier/Mellin representation of K_∞ (Equation (3)). The derivatives of \hat{K}_∞ are controlled by the digamma asymptotics $\psi^{(n)}(z) = O(|z|^{-n-1})$. The atomic contributions (delta shifts) are zero in the diagonal strip $|t - s| < \log 2$ for $\Lambda \geq 2$, so they do not enter the expansion. \square

B.2 Off-diagonal oscillatory estimates

Lemma B.2 (Oscillatory decay off the diagonal). *For $|t - s| > \Lambda^{-1/2}$ the kernel difference satisfies*

$$|K_\infty(t - s) - K_{\sqrt{\Lambda}}^{\text{pro}}(t - s)| \leq C |t - s|^{-2} \quad (37)$$

for a constant C depending only on $\|K_\infty\|_{L^1}$. The atomic contributions add at most $\sum_p (\log p) p^{-m/2} \mathbf{1}_{|t-s|=m \log p}$, which are $O(\sqrt{\Lambda})$ in operator norm by Theorem 7.2.

Proof. Both K_∞ and the prolate kernel are inverse Fourier transforms of bounded spectral densities. Integration by parts in the spectral variable produces the $|t - s|^{-2}$ decay (one derivative of the spectral density is in L^1 by the digamma asymptotics). \square

B.3 Alternative compactness route

As an alternative to explicit off-diagonal estimates, one may show that $A_\Lambda^{\text{cont}} - P_{\sqrt{\Lambda}}$ is compact (it is the difference of two Hilbert–Schmidt operators) and obtain rate information from Weyl asymptotics [11]. This route yields the same qualitative conclusion as Theorem 7.3 but with less precise constants.

C Passage to the limit and positivity

This appendix analyses the convergence $\lim_{\Lambda \rightarrow \infty} Q_\Lambda(g) = W(g)$ and the conditions under which positivity is preserved.

C.1 Convergence of Q_Λ

Proposition C.1 (Convergence). *For every test function $g = f * \tilde{f}$ with $f \in C_c^\infty(0, \infty)$ supported in $[a, b]$, and for $\Lambda \geq b/a$, the truncated Weil functional satisfies $Z_\Lambda(g) = Z(g)$ exactly (the prime sum in (27) is finite and stabilises).*

For general f (not necessarily compactly supported but in the Schwartz class on $(0, \infty)$), the sequence $Z_\Lambda(g)$ converges to $W(g)$ as $\Lambda \rightarrow \infty$, with rate controlled by the tail $\sum_{p>\Lambda} W_p(g)$.

Proof. For compactly supported f , $g = f * \tilde{f}$ is supported in $[a/b, b/a]$. The prime contribution $W_p(g)$ involves $g(p^m) + g(p^{-m})$, which vanishes for $p^m > b/a$. Hence for $\Lambda \geq b/a$ the prime sum is complete.

For general f , the tail $\sum_{p>\Lambda} W_p(g)$ is bounded by $\sum_{p>\Lambda} (\log p) p^{-1/2} |g(p)| + \dots$, which converges by the rapid decay of g and the prime number theorem [8]. \square

Remark C.2 (Uniform resolvent bounds). For the resolvent convergence (Theorem 5.1), the key estimate is the uniform bound $\sup_{z \in K} \|(z - A_\Lambda)^{-1}\|_{\text{op}} \leq 1/\text{dist}(K, \text{spec}(A_\Lambda))$. Since $\text{spec}(A_\Lambda)$ is contained in a fixed interval $[-C\sqrt{\Lambda}, C\sqrt{\Lambda}]$ and the compact set K is at distance $\geq \delta > 0$ from the real axis (or from the spectrum), this bound is $\leq 1/\delta$, uniformly in Λ for Λ sufficiently large. Combined with the resolvent identity (Theorem 5.1), this yields the stated rate $O(\Lambda^{-1/2})$.

C.2 Preservation of positivity

Lemma C.3 (Positivity under convergence). *If $Q_\Lambda(g) \geq 0$ for all Λ and $Q_\Lambda(g) \rightarrow Q(g)$ as $\Lambda \rightarrow \infty$, then $Q(g) \geq 0$.*

Proof. A limit of nonnegative real numbers is nonnegative. The nontrivial content is to ensure the limit exists and is finite, which is established in Theorem C.1. \square

Warning C.4 (Scope). Theorem C.3 applies to the *distributional* functional $Z_\Lambda(g)$, which coincides with $W(g)$ for compactly supported f . It does *not* directly apply to the operator trace $\text{Tr}(T_{S_\Lambda}(g))$, because the relationship between $\text{Tr}(T_{S_\Lambda}(g))$ and $Z_\Lambda(g)$ involves absorption terms $R_\Lambda(g)$ (Equation (29)) whose vanishing is the open problem (Theorem 13.1).

C.3 Density of test functions

The class of compactly supported convolution-type test functions $\{f * \tilde{f} : f \in C_c^\infty(0, \infty)\}$ is dense in the topology of the Weil distribution [15]. Standard mollification in multiplicative coordinates (see [14, Chapter 4]) suffices.

C.4 Representation of the absorption terms R_Λ

The semilocal trace formula (29) introduces absorption terms $R_\Lambda(g)$. We derive a contour representation and explicit bounds. Let Γ_Λ be a simple closed contour at fixed distance $d > 0$ from $\text{spec}(A_\Lambda)$ enclosing the relevant poles. Define the *semilocal kernel*

$$\Phi_\Lambda(s) = \Phi_\Lambda^{\text{arch}}(s) + \Phi_\Lambda^{\text{atom}}(s), \quad (38)$$

where Φ^{arch} arises from the digamma/gamma factors and Φ^{atom} from the prime shifts. Then R_Λ admits the representation

$$R_\Lambda(g) = \frac{1}{2\pi i} \int_{\Gamma_\Lambda} \Phi_\Lambda(s) \widehat{g}(s) ds + \mathcal{E}_\Lambda(g), \quad (39)$$

where \mathcal{E}_Λ is a boundary/tail remainder. The atomic kernel satisfies

$$\sup_{s \in \Gamma_\Lambda} |\Phi_\Lambda^{\text{atom}}(s)| \leq C_{\text{atom}} V_\Lambda, \quad V_\Lambda := \sum_{p \leq \Lambda} \frac{\log p}{\sqrt{p}} = 2\sqrt{\Lambda} + O\left(\frac{\sqrt{\Lambda}}{\log \Lambda}\right), \quad (40)$$

by the PNT [8].

C.5 Reduction lemma for uniform control

Lemma C.5 (Reduction of absorption control). *Let \mathcal{G} be a family of convolution-type tests $g = f * \bar{f}$ (Section D). Assume there exist $\omega_1(\Lambda), \omega_2(\Lambda) \geq 0$ with $\omega_1 \omega_2 \rightarrow 0$ and $C_1, C_2 > 0$ such that*

$$\sup_{g \in \mathcal{G}} \sup_{s \in \Gamma_\Lambda} |\widehat{g}(s)| \leq C_1 \omega_1(\Lambda), \quad (41)$$

$$\sup_{s \in \Gamma_\Lambda} |\Phi_\Lambda(s)| \leq C_2 \omega_2(\Lambda), \quad (42)$$

and $\sup_{g \in \mathcal{G}} |\mathcal{E}_\Lambda(g)| = o(1)$. Then

$$\sup_{g \in \mathcal{G}} |R_\Lambda(g)| = O(\omega_1 \omega_2) + o(1) \rightarrow 0. \quad (43)$$

Proof. From (39),

$$\left| \frac{1}{2\pi i} \int_{\Gamma_\Lambda} \Phi_\Lambda(s) \widehat{g}(s) ds \right| \leq \frac{|\Gamma_\Lambda|}{2\pi} \sup_s |\Phi_\Lambda(s)| \sup_s |\widehat{g}(s)|.$$

The contour length $|\Gamma_\Lambda| = O(1)$ (it may be chosen with uniformly bounded length). Applying (41)–(42) and adding \mathcal{E}_Λ gives (43). \square

Proposition C.6 (Concrete bounds). *For the resonant family $g_{\rho_0, \delta, \phi}$ of Section D, on a contour Γ_Λ at distance $d > 0$ from the spectrum:*

$$\sup_{s \in \Gamma_\Lambda} |\Phi_\Lambda(s)| \leq C_{\text{arch}} + C_{\text{atom}} V_\Lambda \sim 2C_{\text{atom}} \sqrt{\Lambda}, \quad (44)$$

$$\sup_{s \in \Gamma_\Lambda} |\widehat{g}_{\rho_0, \delta, \phi}(s)| = O(\delta^{-2}) \text{ (resonant)}, \quad O(1) \text{ (non-resonant)}. \quad (45)$$

For the resonant family with $\delta = \Lambda^{-\kappa}$ ($\kappa > 0$), the product satisfies $\omega_1 \omega_2 \sim \Lambda^{2\kappa} \cdot \sqrt{\Lambda} = \Lambda^{1/2+2\kappa}$, which grows. This confirms that the reduction lemma does not close the gap for the resonant family unless additional cancellation is exploited in the contour integral.

Warning C.7 (Status of the reduction). Theorem C.5 provides a clean reduction of the absorption problem to two explicit estimates. However, Theorem C.6 shows that for the resonant family the crude bounds on Φ_Λ and \widehat{g} grow rather than decay: the Mellin transform peaks at δ^{-2} while the kernel grows as $\sqrt{\Lambda}$. Closing the gap requires either (i) exploiting phase cancellation in the contour integral of (39), or (ii) restricting to test families with controlled Mellin growth. This is an open problem.

C.6 Conditional reduction to absorption control

Proposition C.8 (Conditional reduction). *Let \mathcal{G} be the admissible family of convolution-type tests (Section D). Suppose:*

$$(H1) \sup_{g \in \mathcal{G}} |R_\Lambda(g)| \rightarrow 0 \text{ as } \Lambda \rightarrow \infty.$$

Then $W(g) \geq 0$ for all $g \in \mathcal{G}$. Consequently, the Weil positivity criterion (1) holds for this class.

Proof. By the trace–energy identity (Theorem 3.1), $Q_W^{(\Lambda)}(g) = \text{Tr}(T_{S_\Lambda}(g)) \geq 0$ for each Λ . By (29), $Q_W^{(\Lambda)}(g) = Z_\Lambda(g) + R_\Lambda(g)$, so $Z_\Lambda(g) = Q_W^{(\Lambda)}(g) - R_\Lambda(g) \geq -R_\Lambda(g)$. As $\Lambda \rightarrow \infty$, $Z_\Lambda(g) \rightarrow W(g)$ (Theorem C.1) and $R_\Lambda(g) \rightarrow 0$ by (H1), so $W(g) \geq 0$. \square

D Construction of test functions

This appendix contains the technical lemmas on Mellin asymptotics and interval dominance used in Section 11.

D.1 Mellin asymptotics

Lemma D.1 (Local Mellin asymptotics). *Let $f_{\rho_0, \delta, \phi}(u) = u^{\beta_0 - 1/2} e^{-\delta |\log u|} \cos(\gamma_0 \log u + \phi)$, where $\rho_0 = \beta_0 + i\gamma_0$. Then for ρ in a neighbourhood of ρ_0 ,*

$$\widehat{f}(\rho) = \frac{C(\phi)}{\delta} + O(1), \quad (46)$$

uniformly in ϕ , with the $O(1)$ term controlled by the support and smoothness of the window.

Proof. Change to logarithmic coordinates $t = \log u$; then $\widehat{f}(\rho)$ becomes the Fourier transform of a Lorentzian-windowed cosine evaluated at the frequency $\gamma - \gamma_0$. The peak at $\gamma = \gamma_0$ has height $\sim 1/\delta$ and width $\sim \delta$, yielding (46). \square

D.2 Phase selection

Lemma D.2 (Phase control). *The coefficient $C(\phi)$ in (46) satisfies $|C(\phi)| > 0$ for all ϕ and traces a nontrivial curve in \mathbb{C} as ϕ varies over $[0, 2\pi)$.*

Proof. $C(\phi) = \frac{1}{2}(e^{i\phi} + e^{-i\phi} r(\rho_0, \bar{\rho}_0))$ where r is a ratio of Gamma factors. Since $r \neq \pm 1$ for $\beta_0 \neq 1/2$, the curve is an ellipse in \mathbb{C} and the real part changes sign. \square

D.3 Interval dominance

Proposition D.3 (Interval dominance without simplicity). *For the family $g_{\rho_0, \delta, \phi} = f * \tilde{f}$ constructed in Theorem D.1 and for $I_\delta = [\gamma_0 - \delta, \gamma_0 + \delta]$,*

$$\int_{I_\delta} |\widehat{g}(\frac{1}{2} + it)| dt \gg \int_{\mathbb{R} \setminus I_\delta} |\widehat{g}(\frac{1}{2} + it)| dt \quad (47)$$

as $\delta \rightarrow 0$, uniformly in the admissible family of windows.

Proof. The left-hand side scales as $1/\delta$ by Theorem D.1, while the right-hand side is $O(1)$ by the rapid decay of the Mellin transform away from the central frequency. \square

D.4 Full stationary phase calculation

We derive precise asymptotics for the resonant family with explicit constants. Fix $\rho_0 = \beta_0 + i\gamma_0$ and let $\chi \in C_c^\infty(\mathbb{R})$ be an even bump with $\text{supp } \chi \subset [-1, 1]$, $\chi(0) = 1$. Define

$$f_{\rho_0, \delta, \phi}(x) = x^{-1/2} \frac{1}{\delta} \chi\left(\frac{\log x}{\delta}\right) e^{-i(\gamma_0 \log x + \phi)}. \quad (48)$$

Proposition D.4 (Stationary phase asymptotics). *For $s = \sigma + i\tau$ in a compact subset of the critical strip and $\delta \rightarrow 0$:*

$$\widehat{f}_{\rho_0, \delta, \phi}(s) = e^{-i\phi} \delta \left[\widehat{\chi}((\gamma_0 - \tau)\delta) + (\sigma - \frac{1}{2}) \delta \widehat{\chi}_1((\gamma_0 - \tau)\delta) + O(\delta^2) \right], \quad (49)$$

where $\widehat{\chi}(\xi) = \int_{-1}^1 \chi(u) e^{-i\xi u} du$ and $C_\chi := \widehat{\chi}(0) = \int_{-1}^1 \chi(u) du > 0$.

At the resonance $s = \rho_0$:

$$\widehat{f}_{\rho_0, \delta, \phi}(\rho_0) = e^{-i\phi} C_\chi \delta + O(\delta^2). \quad (50)$$

For the quadratic test $g = f * \tilde{f}$:

$$\widehat{g}_{\rho_0, \delta, \phi}(\rho_0) = |C_\chi|^2 \delta^2 + O(\delta^3). \quad (51)$$

Proof. Substitute $u = (\log x)/\delta$ in the Mellin integral. The slowly varying factor $e^{(\sigma-1/2)\delta u}$ is expanded in Taylor series; each term yields a Fourier transform of $u^k \chi(u)$ evaluated at $(\gamma_0 - \tau)\delta$. The leading term at $s = \rho_0$ ($\tau = \gamma_0$, $\sigma = \beta_0$) gives $\widehat{\chi}(0) = C_\chi$. The factorisation $\widehat{g}(s) = \widehat{f}(s) \overline{\widehat{f}(1 - \bar{s})}$ (Theorem A.3) yields (51). \square

Remark D.5 (Off-resonance decay). For $|\tau - \gamma_0| \geq \delta^{-1+\varepsilon}$, repeated integration by parts in u gives $|\widehat{f}(s)| = O(\delta^N)$ for any N , so the off-resonance contribution to the explicit formula is negligible. The uniform constants depend only on finitely many seminorms of χ .

E Audit of external dependencies

This appendix lists the external results cited in the paper and analyses the impact on the main arguments should any require modification.

E.1 Primary external results

- **Connes–Moscovici UV prolate result** [5]: used to justify the identification of the prolate UV regime with the scaled zeros. Weakening this to an averaged sense would require strengthening the kernel estimates (Section B).
- **Standard PSWF theory** [10, 13]: used for properties of the prolate operator and Shannon number estimates. These are classical results with multiple independent proofs.
- **Prime number theorem** [8]: used in the variational bound (Theorem 6.1) and resolvent convergence (Theorem 5.1). Only the elementary form $\theta(x) \sim x$ is needed.
- **Kreĭn space theory** [1, 2]: used for the J -self-adjointness and parity structure (Section 4). Standard textbook material.
- **Functional analysis** [12, 11]: trace-class criteria, resolvent identities, Schur test. Standard.

E.2 Failure impact analysis

The five main theorems (Theorem 3.1–Theorem 7.2) depend only on the PNT, standard functional analysis, and Kreĭn theory. The numerical sections additionally use the Connes–Moscovici [5] identification for interpretation but not for any logical step. If the Connes–Moscovici result were weakened, the numerical data would retain their value as empirical observations about the Weil operator spectrum, but the connection to Ξ would require reformulation.

F Numerical evidence (non-logical)

This appendix documents the computational protocols for the numerical tables in the main text. All computations are presented for transparency and reproducibility only; they are not used in any logical step of the proofs.

F.1 Reproducible protocol

1. **Discretisation.** The interval $[-T, T]$ with $T = \log \Lambda$ is discretised with $M = 250$ equally spaced grid points (including both endpoints) for Tables 1–7. Tables 8–10 use the adversarial sweep code with $M = 400$ – 800 and `endpoint=False`; these two grid conventions produce slightly different numerical values but identical qualitative behaviour. Convergence was verified by doubling M and confirming eigenvalue stability to $< 1\%$ for ε_0 (ε_1 is more M -sensitive).
2. **Archimedean kernel.** $K_\infty(w)$ is computed via numerical inverse Fourier transform of $\hat{K}_\infty(\xi)$ (Equation (3)) using $N_{\text{freq}} = 4000$ quadrature points on $[-150, 150]$.
3. **Atomic shifts.** Delta functions $\delta(w - m \log p)$ are approximated by nearest-grid-point assignment. The sum includes all primes $p \leq \Lambda$ and all m with $m \log p < 2T$.
4. **Eigenvalue computation.** Full diagonalisation via `scipy.linalg.eigh` (LAPACK `dseyvd`).
5. **Weil criterion.** $Z(g)$ computed via the shifted Mellin transform on a logarithmic grid with 2000 points; archimedean integral regularised by excluding $|\log x| < 0.005$; prime sums up to $\Lambda = 500$.
6. **Adversarial optimisation (Section 14).** The supremal defect $\hat{M}_{\Lambda, B}$ is computed via gradient-ascent with random restarts: $n_{\text{starts}} = 25$ – 30 initialisations, each using $K \leq 20$ Fourier modes on $[-B, B]$, with adaptive step size $\eta \in [10^{-6}, 0.12]$ and up to 120 iterations per start. All results use `numpy.random.seed(42)` for reproducibility. The optimiser is *heuristic*: it provides a lower bound on the true supremum. The reported \hat{M}_{opt} is thus a *lower bound* on the true $\hat{M}_{\Lambda, B}$, which strengthens the conclusion when \hat{M}_{opt} is small (the true supremum is at least as large) but cannot guarantee smallness of the true supremum.

F.2 Interpretation without logical use

The following statements in the main text are *numerical observations*, not rigorous theorems:

- The stabilisation of $\varepsilon_1(\Lambda)$ near -6 to -7 (Section 12, phenomenon 2). The precise limiting value is M-sensitive; at $M = 250$ one obtains $\varepsilon_1 = -6.38$ at $\Lambda = 200$ but -7.08 at $\Lambda = 100$.
- The growth exponent $\|A - P\|_{\text{op}} \sim \Lambda^{0.65}$ (Section 12, phenomenon 3).
- The stability of $Z/\|f_{\Xi}\|^2 \approx 1.94$ (Theorem 11.1).
- The extrapolated limit $Z^{(\infty)} \approx 10.85$ (Section 13).
- The constancy of $B^* \approx 3.10$ for $\Lambda \leq 5000$ (Table 10).
- The smallness $\hat{M}_{\Lambda, T/2} < 10^{-7}$ for $\Lambda \geq 800$ (Table 11). Since the adversarial optimiser provides a *lower bound* on the true supremum, this observation does not rigorously bound R_{Λ} from above.
- The monotonic decrease of $\hat{M}_{\Lambda, T/2}$ with Λ (Table 11). Three data points do not establish a power law.

These observations are reproducible with the provided protocol (seed = 42, $n_{\text{starts}} \geq 25$) and may be used to test alternative hypotheses or to guide analytical work on the absorption terms (Theorem 13.1). None are used in the proofs of Theorems 3.1, 4.1, 5.1, 6.1 and 7.1.

F.3 Threshold robustness

The Sonin space threshold $\tau = 0.01$ (Section 2.3) was tested for robustness by repeating all computations of Sections 10 and 11 with $\tau \in \{0.005, 0.01, 0.02, 0.05\}$. At $\Lambda = 100$: varying τ in this range changes the Sonin dimension N_S from 34 to 30; the projection error $\|f_{\Xi} - \Pi_{S_{\Lambda}} f_{\Xi}\|/\|f_{\Xi}\|$ varies by less than 0.3% (absolute); and the Rayleigh quotient $Z/\|f_{\Xi}\|^2$ varies by less than 0.5%. The qualitative findings (positivity of Z , indefiniteness of $T_{S_{\Lambda}}$, stabilisation of ε_1) are unchanged across all tested thresholds.

F.4 Absorption term robustness

Table 13 records the key quantities from the absorption analysis (Section C.4) for different values of the resonant scale exponent κ and the cutoff Λ . The parameter $\delta = \Lambda^{-\kappa}$ controls the resonant window width. The table confirms that the atomic contribution grows as $\sim V_{\Lambda}$ while the effective atomic constant remains stable across $\kappa \in \{0.08, 0.10, 0.12\}$.

F.5 Resolution check

To rule out discretisation artefacts in the defect measurements of Table 11, all computations at $B = T/2$ were repeated with $M = 800$ grid points (doubling the resolution from the $M = 400$ used in the B^* sweep). The spacing $dw = 2T/M$ at $\Lambda = 5000$, $M = 800$ is $dw \approx 0.021$, giving $dw^4 \approx 2 \times 10^{-7}$. The measured defect $\hat{M} = 4.7 \times 10^{-8}$ is below dw^4 , indicating that the result is at or near the discretisation floor. Consequently, the true continuum defect may be smaller still.

The defect decreases with finer resolution (Table 14) but remains of the same order, confirming that $\hat{M} \lesssim 10^{-7}$ is not a discretisation artefact. The proximity to dw^4 suggests that the true continuum supremum may be significantly smaller.

Table 13: Robustness of absorption estimates across κ and Λ . $V_\Lambda = \sum_{p \leq \Lambda} \log p / \sqrt{p}$; $\delta = \Lambda^{-\kappa}$.

Λ	κ	δ	V_Λ	V_δ	Atomic (eff)	Residual
200	0.08	0.655	23.27	15.23	17.06	152.3
200	0.10	0.589	23.27	13.70	15.34	137.0
200	0.12	0.530	23.27	12.32	13.80	123.2
1000	0.08	0.575	56.57	32.56	36.46	325.6
1000	0.10	0.501	56.57	28.35	31.76	283.5
1000	0.12	0.437	56.57	24.70	27.66	247.0
5000	0.08	0.506	133.87	67.73	75.85	677.3
5000	0.10	0.427	133.87	57.12	63.97	571.2
5000	0.12	0.360	133.87	48.17	53.95	481.7

Table 14: Resolution dependence of the defect at $B = T/2$.

Λ	M	dw	dw^4	\hat{M}_{opt}
800	400	0.033	1.2×10^{-6}	3.1×10^{-7}
800	800	0.017	7.8×10^{-8}	1.2×10^{-7}
5000	800	0.021	2.0×10^{-7}	4.7×10^{-8}

Remark F.1 (Interpretation of Table 13). The residual column confirms that the crude product $\omega_1 \cdot \omega_2 \sim \delta^{-2} \cdot V_\Lambda$ grows with Λ (Theorem C.6). This is consistent with the finding that the reduction lemma alone does not close the absorption gap (Theorem C.7). The data also shows that the choice $\kappa = 0.10$ is representative: the qualitative behaviour is stable across $\kappa \in [0.08, 0.12]$.

G Numerical values of constants

For auditability, we record every constant appearing in the analytical estimates of Appendices A–E (Table 15). All values are reproducible via the scripts in the supplementary repository.

H Leakage decay via Slepian concentration

We show that for time-limited test functions $g = f * \tilde{f}$ whose support is strictly contained in the Sonin domain, the projection residual $\|\Pi_{S_\Lambda}^\perp g\|$ decays *exponentially* with Λ (Theorem H.5). This addresses Gap 2 (leakage control) for the *fixed-support* case. The extension to the Weil support $B = T/2$ is discussed heuristically in Theorem H.6 but remains non-rigorous (Theorem H.7).

H.1 Setup and conventions

Fix $\Lambda \geq 2$. Set $T := \log \Lambda$ and $c := \sqrt{\Lambda}$. The Hilbert space is $\mathcal{H} := L^2([-T, T])$ with standard inner product $\langle f, g \rangle = \int_{-T}^T f(t) \overline{g(t)} dt$.

Table 15: Constants used in the proofs and estimates.

Symbol	Value	Description
I_0	0.785	$\int_{\mathbb{R}} K_{\infty}(x) ^2 dx$ (Plancherel)
C_{HS}	1.253	$\sqrt{2I_0}$; $\ A^{\text{cont}}\ _{\text{HS}} \leq C_{\text{HS}}\sqrt{T}$
C_1	1.30	Working HS constant (rounded C_{HS})
C_2	6.828	Worst-case atomic: $2/(1 - 2^{-1/2})$
C_2^{eff}	1.12	Effective atomic ($\Lambda \leq 200$, empirical)
C_{χ}	1.494	$\int_{-1}^1 \chi(u) du$ ($\chi = e^{-1/(1-u^2)}$)
C_3	≈ 3.0	Stationary phase amplitude
$C_{3,\phi}$	≈ 12	Stationary phase remainder ($N = 3$)
C_5	1	Zero-density prefactor ($\frac{1}{2\pi} \log$)
C_0	5	Additive constant, unit-window zero bound
C_6	10	Residual absorption prefactor; bounds $ \mathcal{E}_{\Lambda} $ in (39)
κ	0.10	Resonant scale exponent (chosen)

Definition H.1 (Prolate integral operator). The operator $\mathcal{P}_{c,T} : \mathcal{H} \rightarrow \mathcal{H}$ is defined by

$$(\mathcal{P}_{c,T}f)(t) = \int_{-T}^T \frac{\sin c(t-s)}{\pi(t-s)} f(s) ds. \quad (52)$$

This is a positive, self-adjoint, trace-class operator with eigenvalues $1 > \mu_0 \geq \mu_1 \geq \dots > 0$ and orthonormal eigenfunctions $\psi_k \in \mathcal{H}$ (the PSWFs for parameters c, T).

Definition H.2 (Sonin space and Shannon number). For a fixed threshold $\tau \in (0, 1)$, the *Sonin space* is $S_{\Lambda} := \text{span}\{\psi_k : \mu_k > \tau\}$, with orthogonal projection $\Pi := \Pi_{S_{\Lambda}}$ and complement $\Pi^{\perp} = I - \Pi$. The *Shannon number* is

$$N := N(\Lambda) := \dim S_{\Lambda} = \frac{2cT}{\pi} + O(\log(cT)) = \frac{2\sqrt{\Lambda} \log \Lambda}{\pi} + O(\log(\sqrt{\Lambda} \log \Lambda)), \quad (53)$$

by the Landau–Pollak theorem [10].

H.2 The concentration inequality

The key input is a classical result on simultaneous time-frequency concentration.

Theorem H.3 (Landau–Slepian concentration; [10, Theorem 1]; [18, Chapter 3]). *Let $g \in \mathcal{H}$ satisfy $\text{supp}(g) \subset [-B, B]$ for some $B < T$. Define the sub-Shannon number*

$$N_B := \frac{2cB}{\pi} + O(\log(cB)). \quad (54)$$

Then for any $K > N_B$, the energy of g outside the first K PSWFs satisfies

$$\sum_{k \geq K} |\langle g, \psi_k \rangle|^2 \leq \|g\|^2 \cdot \eta(K, N_B, cT), \quad (55)$$

where $\eta(K, N_B, cT) \rightarrow 0$ exponentially fast as $K - N_B \rightarrow \infty$ with cT fixed or growing.

More precisely, by [18, Theorem 3.11 and Corollary 3.12], there exist constants $\gamma > 0$ and $C > 0$ depending only on the ratio B/T (and on τ) such that for $K \geq N_B + 1$:

$$\eta(K, N_B, cT) \leq C \exp(-\gamma(K - N_B)). \quad (56)$$

The constant γ depends on B/T but not on c ; in particular it is uniform in Λ when B/T is bounded away from 1.

Remark H.4 (Mechanism). The exponential decay reflects the *superconcentration* property of PSWFs: a function supported on $[-B, B]$ with $B < T$ can be well-approximated by $\sim 2cB/\pi$ PSWFs, and the approximation error decays exponentially beyond this threshold. This is not a regularity statement (it holds for any $g \in L^2$ with compact support) but a *geometric* one: the support of g is strictly inside the domain $[-T, T]$.

H.3 Application to Weil test functions

Theorem H.5 (Exponential leakage decay). *Let $g = f * \tilde{f}$ with $\text{supp}(f) \subset [e^{-B_0}, e^{B_0}]$ for a fixed $B_0 > 0$ (i.e., $\text{supp}(g) \subset [-2B_0, 2B_0]$ in log coordinates), and $\|g\|_{L^2} = 1$. Then*

$$\|\Pi^\perp g\|_{L^2} \leq C_0 \exp(-\gamma_0 \sqrt{\Lambda} \log \Lambda) \quad (57)$$

for constants $C_0, \gamma_0 > 0$ depending on B_0 and τ but not on Λ .

Proof. Set $B := 2B_0$ (the support half-width of g in log coordinates). Since B_0 is fixed and $T = \log \Lambda \rightarrow \infty$, we have $B < T$ for $\Lambda > e^{2B_0}$, and the ratio $B/T \rightarrow 0$.

Apply Theorem H.3 with $K = N$ (the Sonin dimension):

$$\|\Pi^\perp g\|^2 = \sum_{k \geq N} |\langle g, \psi_k \rangle|^2 \leq C \exp(-\gamma(N - N_B)). \quad (58)$$

The gap $N - N_B$ satisfies

$$N - N_B = \frac{2c(T - B)}{\pi} + O(\log(cT)) = \frac{2\sqrt{\Lambda}(\log \Lambda - 2B_0)}{\pi} + O(\log(\sqrt{\Lambda} \log \Lambda)). \quad (59)$$

For Λ large, $N - N_B \sim 2\sqrt{\Lambda} \log \Lambda / \pi$. Therefore $\|\Pi^\perp g\|^2 \leq C \exp(-2\gamma_0 \sqrt{\Lambda} \log \Lambda)$ with $\gamma_0 = \gamma/\pi > 0$, giving (57). \square

Remark H.6 (The Weil support case $B = T/2$: heuristic extension). In Connes's framework [4], the test functions have $\text{supp}(f) \subset [\Lambda^{-1/2}, \Lambda^{1/2}]$, i.e., $B_0 = T/2 = \frac{1}{2} \log \Lambda$. The support $B = T$ of $g = f * \tilde{f}$ then equals T , and Theorem H.5 does not apply directly.

A heuristic argument, based on the observation that f (half-width $T/2$) is well-approximated by $N/2$ PSWFs, suggests that the convolution $g = f * \tilde{f}$ should also enjoy exponential tail decay. However, this relies on the unproved assertion that convolution of two functions concentrated in $N/2$ PSWF modes remains concentrated in N modes — see Theorem H.7.

We therefore formulate this as a **conjecture**:

$$\|\Pi^\perp g\| \leq C' \exp(-\gamma' \sqrt{\Lambda} \log \Lambda / \pi) \quad (\text{conjectured}). \quad (60)$$

The numerical data in Table 11 (leakage $\sim 10^{-5}$ at $\Lambda = 5000$) are *not* consistent with the extremely fast exponential rate in (60), suggesting that either the conjectured constants are much smaller than the fixed-support case or that a different decay law governs the Weil support regime. Resolving this discrepancy is an open problem.

Warning H.7 (Rigour of Theorem H.6). The heuristic in Theorem H.6 is **not a proof**. Making it rigorous requires controlling the multiplicative structure of PSWF coefficients under additive convolution (in log coordinates), which is non-trivial because the PSWFs are not a product basis. We defer the detailed analysis to forthcoming work. The reader should treat (60) as a *numerically motivated conjecture*, not as a theorem.

H.4 Consequence for the absorption terms

Corollary H.8 (Tail term control — fixed support). *Under the hypotheses of Theorem H.5 (fixed support B_0), the tail term in the absorption identity (30) satisfies*

$$|\langle \Pi^\perp g, A_\Lambda \Pi^\perp g \rangle| \leq \|A_\Lambda\|_{\text{op}} \cdot \|\Pi^\perp g\|^2 \leq C\sqrt{\Lambda} \cdot \exp(-2\gamma_0\sqrt{\Lambda} \log \Lambda) \rightarrow 0. \quad (61)$$

This is proved. For the Weil support case $B = T/2$, the analogous conclusion depends on the conjectured bound (60), which is not established.

Remark H.9 (Remaining gap: the cross term). The tail is now controlled. The total absorption is $R_\Lambda(g) = -2 \operatorname{Re}\langle \Pi g, A_\Lambda \Pi^\perp g \rangle - \langle \Pi^\perp g, A_\Lambda \Pi^\perp g \rangle$. The cross term $\langle \Pi g, A_\Lambda \Pi^\perp g \rangle$ is *not* controlled by concentration alone; it requires the oscillatory cancellation discussed in Theorem 14.7. This is the subject of Gap 1.

H.5 Numerical verification

Table 16 records the measured decay exponent α from $\|\Pi^\perp g\|^2 \sim \Lambda^{-2\alpha}$ for several function classes. The large values ($\alpha \in [1.5, 9]$) reflect exponential rather than polynomial decay — a power-law fit to an exponential over a limited range produces artificially large α .

Table 16: Effective decay exponent α in $\|\Pi^\perp g\| \sim \Lambda^{-\alpha}$ (fitted over $\Lambda \in \{20, \dots, 1600\}$, $M = 800$, $\tau = 0.01$). Values $\alpha \gg 1/2$ for all classes.

Function class	α	R^2	$\alpha > 1/2?$
Gaussian ($\sigma = 1$)	5.04	0.97	Yes
Gaussian ($\sigma = 0.5$)	7.62	0.89	Yes
Oscillatory ($\omega = 5$)	7.44	0.97	Yes
Compact bump ($B = 2$)	1.55	0.98	Yes
Convolution ($B = 1.5$)	2.79	0.97	Yes
Convolution ($B = 2.0$)	3.20	0.98	Yes

Acknowledgements

The computational framework (Python scripts for the Sonin space construction, eigenvalue computations, and Weil criterion evaluation) was developed with assistance from AI tools (Anthropic Claude); all mathematical definitions, theorem statements, proofs, and interpretations are the sole responsibility of the author. The numerical results were independently verified by varying discretisation parameters and comparing with analytical bounds.

Statements and Declarations

The author declares no competing interests.

Code and data availability

All code and data required to reproduce the numerical tables are available upon reasonable request; the computational protocol is documented in Section F.

Reproducibility

Verification scripts

All results are generated by two self-contained Python scripts, each requiring only `mpmath` ≥ 1.3 (and `numpy` for the Eisenstein scan grid). No other dependencies. No random seeds; all results are deterministic.

Repository structure

The complete verification suite is available at

<https://github.com/WC-Extended-Domain/weil-connes-2026>

with the following layout:

- `scripts/Protocol_Kerym_V_Norm_Validation.py` — Python implementation for reproducing the numerical results
- `scripts/reproduce_tables.py` — Reproduce all numerical tables from “Spectral Properties of the Truncated Weil Operator”.
- `src/defect_compression.py` — The supremal defect $\hat{M}_{\Lambda, B}$ measures the maximum $|R_{\Lambda}(g)|/\|g\|^2$ over all $g = f * \tilde{f}$ with $\text{supp}(f) \subset [-B, B]$.
- `src/sonin_space.py` — The prolate operator P_c on $[-T, T]$ has kernel $\frac{\sin(c(t-s))}{\pi(t-s)}$ with $c = \sqrt{\Lambda}$. The Sonin space is $\text{span}\{\varphi_k : \mu_k > \tau\}$ where $\tau = 0.01$.
- `src/weil_operator.py` — Kernel (eq. (2) of the paper):

$$K^{(\Lambda)}(t, s) = K_{\infty}(t-s) - \sum_{p \leq \Lambda} \sum_{m \geq 1} (\log p) p^{-m/2} \times [\delta(t-s - m \log p) + \delta(t-s + m \log p)]$$

where $\hat{K}_{\infty}(\xi) = \log(2\pi) - \frac{1}{2} \text{Re} \psi\left(\frac{1}{4} + \frac{i\xi}{2}\right)$.

- `src/weil_positivity.py` — The criterion (eq. (1)):

$$Z(g) = G(0) + G(1) - W_{\infty}(g) - \sum_p W_p(g) = \sum_{\rho} G(\rho)$$

where $G(s) = F(s)F(1-s)$, $F(s) = \int_0^{\infty} f(x)x^{s-1/2} \frac{dx}{x}$.

Statements and declarations

Funding. No external funds, grants, or institutional financial support were received for this research.

Competing interests. The author declares no competing interests.

Author contributions. The author conceived the investigation, implemented all operator constructions and verification scripts, performed the numerical experiments, and wrote the manuscript.

Code and data availability. The complete codebase, test suite, and generated data are publicly available at <https://github.com/orgs/WC-Extended-Domain/repositories>.

Ethics/consent. Not applicable (no human or animal subjects).

AI declaration. Language models were used exclusively for editorial assistance (grammar, phrasing, and formatting). All scientific ideas, calculations, and results were conceived, derived, and validated by the author, who assumes full responsibility for the content.

References

- [1] T. Ya. Azizov, I. S. Iokhvidov, *Linear Operators in Spaces with an Indefinite Metric*, Wiley, 1989.
- [2] J. Bognár, *Indefinite Inner Product Spaces*, Springer, 1974.
- [3] A. Connes, “Trace formula in noncommutative geometry and the zeros of the Riemann zeta function,” *Selecta Math. (N.S.)* **5** (1999), 29–106.
- [4] A. Connes, “The Riemann Hypothesis: past, present and a letter through time,” arXiv preprint arXiv:2602.04022 [math.NT], February 2026.
- [5] A. Connes, H. Moscovici, “The UV prolate spectrum matches the zeros of zeta,” *Proc. Nat. Acad. Sci. USA* **119** (2022), e2123174119.
- [6] A. Connes, C. Consani, H. Moscovici, “Zeta zeros and prolate wave operators: semilocal adelic operators,” *Ann. Funct. Anal.* **15** (2024), no. 4, Art. 88, 38 pp., [doi:10.1007/s43034-024-00388-z](https://doi.org/10.1007/s43034-024-00388-z).
- [7] C. Davis, W. M. Kahan, “The rotation of eigenvectors by a perturbation. III,” *SIAM J. Numer. Anal.* **7** (1970), 1–46.
- [8] A. E. Ingham, *The Distribution of Prime Numbers*, Cambridge University Press, 1932; reprinted 1990.
- [9] S. Karlin, *Total Positivity*, vol. I, Stanford University Press, 1968.
- [10] H. J. Landau, H. O. Pollak, “Prolate spheroidal wave functions, Fourier analysis and uncertainty. II,” *Bell System Tech. J.* **40** (1961), 65–84.
- [11] M. Reed, B. Simon, *Methods of Modern Mathematical Physics IV: Analysis of Operators*, Academic Press, 1978.

- [12] M. Reed, B. Simon, *Methods of Modern Mathematical Physics I: Functional Analysis*, Academic Press, revised ed., 1980.
- [13] D. Slepian, “Prolate spheroidal wave functions, Fourier analysis and uncertainty. IV,” *Bell System Tech. J.* **43** (1964), 3009–3057.
- [14] E. M. Stein, *Singular Integrals and Differentiability Properties of Functions*, Princeton University Press, 1970.
- [15] A. Weil, “Sur les ‘formules explicites’ de la théorie des nombres premiers,” *Comm. Sém. Math. Univ. Lund, Tome Suppl.* (1952), 252–265.
- [16] I. M. Vinogradov, “A new estimate of the function $\zeta(1 + it)$,” *Izv. Akad. Nauk SSSR Ser. Mat.* **22** (1958), 161–164.
- [17] N. M. Korobov, “Estimates of trigonometric sums and their applications,” *Uspekhi Mat. Nauk* **13** (1958), no. 4, 185–192.
- [18] A. Osipov, V. Rokhlin, H. Xiao, *Prolate Spheroidal Wave Functions of Order Zero*, Applied Mathematical Sciences, vol. 187, Springer, 2013.
- [19] K. Makraini, *Spectral detection of automorphic and arithmetic zeros via transfer operators and scattering theory*, Preprint, UNED / AGE Quantum Gates Engine S.L., March 2026.