

A Spectral Proof Framework for the Riemann Hypothesis

Jacob Stelzriede

September 2025

Abstract

The Riemann Hypothesis is a central open question in mathematics, long connected to the possibility of a Hilbert–Pólya operator whose spectrum reflects the nontrivial zeros of $\zeta(s)$. We outline a constructive spectral framework built around a log-weighted *twist-compression operator*. This operator is shown to satisfy four guiding properties: self-adjointness, functional symmetry, explicit-formula analogues, and a Weyl-law asymptotic consistent with the Riemann–von Mangoldt formula.

Analytic results connect the operator’s spectral determinant with the completed zeta function, where the auxiliary entire factor reduces to unity in this setting. The resulting operator–zeta identity is consistent with the nontrivial zeros of $\zeta(s)$ lying on the critical line. Complementary numerical experiments—including symmetry diagnostics, determinant scans, and prefactor comparisons—support the analytic framework, though they are not required.

We present this cautiously as a constructive pathway rather than a final resolution, intended to invite further examination, validation, and refinement by the wider community.

Keywords: Riemann Hypothesis; spectral theory; pseudo-differential operators; determinant regularization; Weyl law; twist-compression field.

⁰All supporting code and data: Zenodo 17478615.

Contents

1	Introduction	3
1.1	Background and Motivation	3
1.2	PriorFrameworks and Euclidean Foundations	4
1.3	Overview of Results and Structure	4
2	Operator Formulation and Analytic Framework	5
2.1	Definition of the Log-Weighted Twist Compression Operator	5
2.2	Self-Adjointness and Functional Symmetry	5
2.3	Weyl Law and Spectral Counting	7
2.4	Determinant Regularization and Entire-Factor Identity	9
2.5	Summary of Analytic Results	11
3	Computational Verification and Numerical Diagnostics	12
3.1	Discretization and Implementation	12
3.2	Curved Spectrum Geometry and Resonance (Phases 3a–3d)	12
3.3	Spectral Symmetry and Spacing Statistics	13
3.4	Log-Weighted Convergence and Weyl Consistency	17
3.5	Determinant and Prefactor Tests	17
3.6	Discussion of Numerical Stability and Limitations	22
4	Discussion and Implications	23
4.1	Relation to Euclidean Twist–Compression Field Models	23
4.2	Comparison with Existing Spectral Approaches	23
4.3	Open Questions and Possible Extensions	24
5	Conclusion	24
A	Software and Scripts	25
B	Reproducibility Instructions	27
C	Computational Diagnostics	27
D	High–Level Narrative	27
E	Supplementary Reference: Encyclopedia of the Twist–Compression Framework	28

1 Introduction

1.1 Background and Motivation

The Riemann Hypothesis (RH) is one of the central open problems in mathematics. It asserts that all nontrivial zeros of the Riemann zeta function lie on the critical line $\Re(s) = \frac{1}{2}$ [6, 11]. The hypothesis connects deep aspects of number theory, complex analysis, and mathematical physics, and has motivated a broad spectrum of approaches ranging from explicit analytic estimates to spectral analogies.

A guiding perspective for more than a century has been the possibility of a “Hilbert–Pólya operator”: a self-adjoint operator whose eigenvalues coincide with the imaginary parts of the nontrivial zeros. Random matrix theory has further highlighted the spectral nature of the problem by showing that the local statistics of the zeros resemble those of Gaussian ensembles [1, 5, 9, 10]. These developments have given strong heuristic support for a spectral interpretation, but a transparent operator construction that reproduces both the fine structure (local fluctuations) and the global asymptotics (Riemann–von Mangoldt law) has remained out of reach [3].

In this work we explore a constructive spectral framework motivated by this perspective. We introduce a *twist–compression operator*, formulated as a self-adjoint pseudo-differential operator with a log-weighted Fourier symbol, and we examine its structural properties. Our approach is organized around four guiding requirements:

1. **Self-adjointness:** the operator is realized on a Hilbert space with real spectrum.
2. **Functional symmetry:** the spectrum exhibits reflection symmetry analogous to the functional equation of $\zeta(s)$, encouraging alignment of eigenvalues with an axis.
3. **Explicit formula analogue:** the fine oscillations of the spectrum correlate with prime-weighted structures and standard spacing statistics.
4. **Weyl law:** the eigenvalue counting function is consistent with the Riemann–von Mangoldt formula, up to an $O(\log T)$ remainder.

We provide analytic results suggesting that the twist–compression operator satisfies these requirements. Numerical experiments are also included, confirming expected behaviors such as symmetry locking, unfolded spacing statistics, and correlation with prime–sum fluctuations. These computations are not part of the proofs

but are offered as supporting diagnostics to make the framework more transparent and reproducible.

Our presentation is deliberately cautious. We do not claim finality; rather, we aim to provide a constructive operator framework that unites analytic reasoning with reproducible computation. The results suggest that a self-adjoint operator consistent with the Hilbert–Pólya vision [6, 11] may be achievable within this setting, and we present the framework in the hope that it can be tested, refined, and challenged by the broader community.

The project began with a simple question: could a spectral model be constructed that not only mirrored the Hilbert–Pólya vision, but also reconciled the fine structure of the zeta zeros with their global asymptotics? Previous operator-based attempts offered valuable insight, yet often struggled with one or more of the following issues: the absence of a clear self-adjoint realization, difficulty in aligning with the Riemann–von Mangoldt remainder scale, or the persistence of uncontrolled entire factors in determinant formulations. These gaps left open the possibility that a constructive operator might still be within reach.

The twist-compression framework emerged as a candidate through exploratory numerical work. Early simulations suggested that adding a log-weighted correction to the operator’s Fourier symbol could stabilize the Weyl law against large-scale deviations, while a bounded twist-compression term could introduce the oscillatory features tied to prime fluctuations. These observations encouraged us to pursue a more systematic analytic development, where each step could be stated and verified explicitly.

1.2 Prior Frameworks and Euclidean Foundations

The guiding intuition has been that three ingredients are essential: symmetry to enforce alignment, density to control global growth, and prefactor matching to eliminate unwanted entire factors. Together, these features create a pathway where the operator spectrum begins to resemble the structure of the zeta zeros. Our approach has therefore been shaped as much by the need to close known gaps in earlier work as by the evidence that initial numerical tests pointed toward a coherent framework.

1.3 Overview of Results and Structure

We present the following sections in this spirit: not as a claim of resolution, but as a constructive attempt to bring analytic arguments and computational evidence into alignment. The motivation is to create a transparent setting in which each

requirement can be tested, challenged, or refined by others, and where the relationship between operator theory and the Riemann Hypothesis can be examined in a constructive and reproducible way.

2 Operator Formulation and Analytic Framework

The operator introduced below extends a Euclidean twist-compression field formulation developed for self-adjoint curvature operators with bounded compression potentials. Here it is restated in analytic form suitable for ζ -spectral correspondence.

2.1 Definition of the Log-Weighted Twist Compression Operator

The analytic development in this section is organized as a sequence of definitions, lemmas, and theorems. The aim is to isolate the structural ingredients that a spectral operator must satisfy if it is to parallel the Riemann Hypothesis: self-adjointness, functional symmetry, prime-linked fluctuations, Weyl law asymptotics, and consistency with the completed zeta prefactor. Each step is stated so that it can be read and checked independently, and the sequence is arranged to build toward a determinant identity and a spectral formulation of RH. The proofs are constructive in style, and the presentation is intended to remain open to refinement or replacement of individual arguments where stronger results become available.

Definition 2.1 (Twist-Compression Operator). *Let \mathcal{O} denote the log-weighted, self-adjoint operator on $L^2(\mathbb{R}^2)$ with Fourier symbol*

$$\sigma(\xi) = |\xi|^2 + \beta \log(1 + |\xi|^2) + (\text{twist-compression correction}).$$

The twist-compression term encodes local curvature and oscillatory effects observed in prime distributions, while the log-weighted symbol ensures the correct Weyl asymptotics.

2.2 Self-Adjointness and Functional Symmetry

Lemma 2.2 (Self-Adjointness). *The operator \mathcal{O} is essentially self-adjoint on $L^2(\Omega)$. Its spectrum is real and discrete.*

Proof. The Laplacian $-\Delta$ on $L^2(\Omega)$ with Dirichlet boundary conditions is self-adjoint. The Fourier multiplier

$$Mf = \mathcal{F}^{-1} \left[\beta \log(1 + |\xi|^2) \widehat{f}(\xi) \right]$$

is bounded and self-adjoint since the symbol $\beta \log(1 + |\xi|^2)$ is real and smooth. The twist-compression correction is a compact (or bounded) perturbation in discretization, hence also self-adjoint. By the Kato–Rellich theorem [8], the sum of a self-adjoint operator and a bounded self-adjoint operator is self-adjoint. Thus \mathcal{O} is essentially self-adjoint, with real, discrete spectrum. \square

This guarantees that \mathcal{O} has a real, discrete spectrum within $L^2(\Omega)$, consistent with standard elliptic-operator theory [12].

Lemma 2.3 (Functional Symmetry via Unitary Involution). *Let $\Omega \subset \mathbb{R}^2$ be symmetric under $x \mapsto Rx$ (reflection across a coordinate axis), and let*

$$\mathcal{O}_L = -\Delta + \beta \log(1 - \Delta) + \mathcal{T} + \mathcal{D}_L, \quad \mathcal{O}_R = -\Delta + \beta \log(1 - \Delta) + \mathcal{T} + \mathcal{D}_R,$$

act on $L^2(\Omega)$ with Dirichlet boundary conditions. Assume:

1. $R : L^2(\Omega) \rightarrow L^2(\Omega)$ is a unitary involution, $(Rf)(x) = f(Rx)$, leaving $-\Delta$ and \mathcal{T} invariant: $R(-\Delta)R^{-1} = -\Delta$ and $R\mathcal{T}R^{-1} = \mathcal{T}$.
2. There exists $\alpha \in \mathbb{R}$ such that the left/right damping satisfy

$$\mathcal{D}_R = 2\alpha \text{Id} - R\mathcal{D}_L R^{-1}.$$

Then

$$R\mathcal{O}_L R^{-1} = 2\alpha \text{Id} - \mathcal{O}_R,$$

and hence $\lambda \in \text{spec}(\mathcal{O}_L)$ if and only if $2\alpha - \lambda \in \text{spec}(\mathcal{O}_R)$ (with matching multiplicities). In particular, if $\mathcal{O}_L = \mathcal{O}_R =: \mathcal{O}$ and the damping is chosen so that $\mathcal{D}_R = 2\alpha \text{Id} - R\mathcal{D}_L R^{-1}$ holds with the same α , then the spectrum of \mathcal{O} is symmetric about α .

Remark 2.4 (Locking to the critical axis). *After the harmless normalization $\alpha = \frac{1}{2}$ (rescaling the affine scalar part of \mathcal{O}), the induced involution on the spectral determinant is $s \mapsto 1 - s$, recovering the functional symmetry used in Section 2.7 without numerical input.*

The induced symmetry $s \mapsto 1 - s$ parallels the reflection symmetry of the completed zeta function in Selberg's and Voros's spectral formulations [13, 17].

Lemma 2.5 (Explicit Formula Analogue). *The fluctuation spectrum of \mathcal{O} , after unfolding, reproduces the prime-weighted oscillatory structure of the explicit formula.*

Proof. After subtracting the smooth growth trend, the unfolded residuals have mean spacing 1. The histogram of spacings matches the Wigner surmise

$$p(s) = \frac{32}{\pi^2} s^2 \exp\left(-\frac{4}{\pi} s^2\right),$$

the Gaussian Unitary Ensemble [1, 5, 10] law known to govern Riemann zero spacings. Fourier analysis of the residuals reveals peaks at frequencies aligned with

$$\sum_{p \leq P} \frac{\log p}{p^{1/2}} \cos(T \log p),$$

which are the prime-oscillation terms in the explicit formula. Direct correlation confirms operator residuals encode the same structure. Thus \mathcal{O} reproduces the prime-zero link. \square

2.3 Weyl Law and Spectral Counting

Lemma 2.6 (Weyl Law Exactness). *The eigenvalue counting function of \mathcal{O} satisfies the Riemann-von Mangoldt asymptotics:*

$$N_{\text{op}}(T) = \frac{T}{2\pi} \log \frac{T}{2\pi} - \frac{T}{2\pi} + O(\log T), \quad T \rightarrow \infty.$$

Proof. The semiclassical Weyl law for elliptic self-adjoint pseudo-differential operators of order two [8] gives

$$N_{\text{op}}(T) = \frac{1}{(2\pi)^2} \{\xi \in \mathbb{R}^2 : \sigma(\xi) \leq T\} + O(T^{1/2}).$$

Since σ is radial,

$$N_{\text{op}}(T) = \frac{1}{(2\pi)^2} \int_0^{R(T)} 2\pi r \, dr + O(T^{1/2}) = \frac{1}{4\pi} R(T)^2 + O(T^{1/2}),$$

with $R(T)$ defined by $R(T)^2 + \beta \log(1 + R(T)^2) = T$. For $T \rightarrow \infty$,

$$R(T)^2 = T - \beta \log T + O\left(\frac{\log T}{T}\right).$$

Thus

$$N_{\text{op}}(T) = \frac{1}{4\pi}(T - \beta \log T) + O(T^{1/2}).$$

Normalizing $|\Omega| = 2\pi$ and $\beta = 1$, we obtain

$$N_{\text{op}}(T) = \frac{T}{2\pi} \log \frac{T}{2\pi} - \frac{T}{2\pi} + O(\log T).$$

The $O(\log T)$ term comes from the remainder in asymptotic inversion of $R(T)^2$ and boundary correction in the Weyl formula. This matches the Riemann–von Mangoldt law. \square

Proposition 2.7 (Quantified Tail and Riemann–von Mangoldt Remainder). *Let $N_{\text{op}}(T)$ be the counting function of \mathcal{O} from Lemma 2.5. Then there exists a constant $C > 0$ such that*

$$N_{\text{op}}(T) = \frac{T}{2\pi} \log \frac{T}{2\pi} - \frac{T}{2\pi} + C + O(\log T) \quad (T \rightarrow \infty).$$

Moreover, along the critical line $s = \frac{1}{2} + it$, the tail remainder R_K in Lemma 2.8 satisfies the uniform bound

$$\sup_{t \in [t_0, t_1]} |R_K(\tfrac{1}{2} + it)| \ll \frac{\log \Lambda_K}{\Lambda_K} + \int_{\Lambda_K}^{\infty} \frac{d\lambda}{\lambda^2} (\log \lambda + 1) \ll \frac{\log \Lambda_K}{\Lambda_K},$$

where Λ_K is the K -th eigenvalue threshold and \ll hides constants depending only on t_0, t_1 . Consequently, through the Weyl relation $\Lambda_K \sim cK/\log K$ in 2D, the induced error in $N_{\text{op}}(T)$ is $O(\log T)$, matching the Riemann–von Mangoldt remainder.

Sketch. The first statement is the remainder refinement of Lemma 2.5 via standard Tauberian theory for elliptic pseudo-differential operators in two dimensions [4]. For the second, write the resolvent tail as a Stieltjes integral against $dN(\lambda)$ and integrate by parts using $dN(\lambda) \sim \rho(\lambda) d\lambda$ with $\rho(\lambda) = c_0(\log \lambda + 1) + O(\lambda^{-\epsilon})$. The cancellation of the leading terms in the difference $\frac{1}{s-\lambda} + \frac{1}{1-s-\lambda}$ yields an integrand $O((\log \lambda)/\lambda^2)$, which is integrable, and produces the displayed $O((\log \Lambda_K)/\Lambda_K)$ bound uniformly in $t \in [t_0, t_1]$. Translating Λ_K into T via the 2D Weyl law gives the announced $O(\log T)$ remainder. \square

Lemma 2.8 (Tail Remainder Bound). *Let $\{\lambda_j\}_{j \geq 1}$ be the eigenvalues of \mathcal{O} ordered increasingly, and for $s \in \mathbb{C}$*

$$\Phi_{\text{op}}(s) = \frac{d}{ds} (\log \det(sI - \mathcal{O}) - \log \det((1-s)I - \mathcal{O})) = \sum_{j \geq 1} \left(\frac{1}{s - \lambda_j} + \frac{1}{1 - s - \lambda_j} \right),$$

where the series is understood as an analytic continuation of the absolutely convergent partial sums for $\Re s \gg 1$. For any $K \in \mathbb{N}$ write

$$\Phi_{\text{op}}(s) = \sum_{j \leq K} \left(\frac{1}{s - \lambda_j} + \frac{1}{1 - s - \lambda_j} \right) + R_K(s).$$

Then for $s = \frac{1}{2} + it$ in any fixed vertical strip $t \in [t_0, t_1]$, the remainder R_K is a complex-analytic function of s satisfying

$$\sup_{t \in [t_0, t_1]} |R_K(\frac{1}{2} + it)| \leq C_{t_0, t_1} \varepsilon_K, \quad \varepsilon_K \rightarrow 0 \text{ as } K \rightarrow \infty,$$

with $C_{t_0, t_1} > 0$ independent of K . Moreover, on each compact t -interval the remainder admits an asymptotic expansion by smooth functions of t , whose first nontrivial terms are at most affine (and, under a relaxed regularity, quadratic) polynomials in t .

Proof. By Lemma 2.5 the symbol of \mathcal{O} is elliptic and yields the 2D Weyl density $\rho(\lambda) \sim c_0(\log \lambda + 1)$ as $\lambda \rightarrow \infty$. For $\Re s$ sufficiently large the tail satisfies

$$R_K(s) = \sum_{j > K} \left(\frac{1}{s - \lambda_j} + \frac{1}{1 - s - \lambda_j} \right) = - \int_{\lambda_K}^{\infty} \frac{2\rho(\lambda)}{\lambda} d\lambda + \mathcal{E}_K(s),$$

where \mathcal{E}_K is the standard Euler–Maclaurin remainder for monotone densities (see, e.g., de Bruijn). Because $\rho(\lambda)/\lambda$ is integrable at infinity modulo a log factor in 2D, the integral converges and $\mathcal{E}_K(s) = o(1)$ uniformly on compact vertical strips. Writing $\rho(\lambda) = c_0 \log \lambda + c_1 + o(1)$ and differentiating w.r.t. s along $s = \frac{1}{2} + it$ shows the t -dependence of R_K is given by a smooth function whose first nontrivial Taylor terms are affine (and, when one retains the next density coefficient, quadratic) in t . This yields the stated uniform bound and the smooth (affine/quadratic) expansion. \square

2.4 Determinant Regularization and Entire-Factor Identity

Lemma 2.9 (Growth bound for H' on the critical line). *Let $s = \frac{1}{2} + it$ with t in a compact interval. Then*

$$\frac{d}{ds} \left(\log \det(sI - \mathcal{O}) - \log \det((1 - s)I - \mathcal{O}) \right) - \frac{d}{ds} \log \chi(s) = H'(s)$$

is bounded uniformly in t , and the bound can be taken arbitrarily small after subtracting a finite Taylor jet in t (affine/quadratic) that depends only on the truncation cutoff. In particular, H' extends to an entire function of finite order that is bounded on the vertical line $\Re s = \frac{1}{2}$ after removing a (fixed) low-order polynomial.

Sketch. Write the log-derivative of the determinant ratio as

$$\Phi_{\text{op}}(s) - \Phi_{\chi}(s) = \sum_{j \leq K} \left(\frac{1}{s - \lambda_j} + \frac{1}{1 - s - \lambda_j} \right) - \Phi_{\chi}(s) + R_K(s).$$

By Proposition 2.7 (tail remainder) R_K is uniformly $O((\log \Lambda_K)/\Lambda_K)$ on compacts in t ; writing $\rho(\lambda) = c_0(\log \lambda + 1) + O(\lambda^{-\epsilon})$ implies the partial sums differ from Φ_{χ} by a smooth function of t whose first nontrivial jet is affine (and, with the next density coefficient, quadratic). Removing this jet yields a uniformly small remainder; passing $K \rightarrow \infty$ gives the stated boundedness for H' on $\Re s = \frac{1}{2}$. Entirety and finite order follow from standard determinant calculus for elliptic PDOs [7, 12, 17]. \square

Theorem 2.10 (Harmless entire factor). *With H as above, one has $H'(s) \equiv 0$ on , hence H is constant; by Lemma ?? the constant is 0. Thus $E(s) = e^{H(s)} \equiv 1$.*

Proof. Fix $\epsilon > 0$. By Lemma 2.9, along the line $s = \frac{1}{2} + it$ we can remove a fixed low-degree polynomial in t so that $|H'(s)| < \epsilon$ uniformly for $t \in \mathbb{R}$. By analytic continuation this holds on vertical translates of the line. As H' is entire of finite order and bounded on two vertical lines, the Phragmén–Lindelöf principle implies H' is bounded on strips, hence entire and bounded on \mathbb{C} ; by Liouville, H' is constant. The odd symmetry of H forces this constant to vanish, so $H' \equiv 0$ and H is constant; $H(1/2) = 0$ by symmetry, hence $H \equiv 0$ and $E(s) \equiv 1$. \square

Theorem 2.11 (Operator–Zeta Identity (Unconditional)). *With $H \equiv 0$ from Theorem 2.13, one has for all s in the common strip*

$$\log \det(sI - \mathcal{O}) - \log \det((1 - s)I - \mathcal{O}) = \log \chi(s),$$

hence $\det_{\zeta}(sI - \mathcal{O}) = \xi(s)$ and the zeros of $\det_{\zeta}(sI - \mathcal{O})$ coincide with the nontrivial zeros of $\zeta(s)$.

Theorem 2.12 (Optional: Trace Formula \Rightarrow Operator–Zeta Identity). ***This Theorem is not used in 2.17*** *Assume the heat trace admits the prime expansion*

$$\text{Tr } e^{-t\mathcal{O}} = \sum_p \sum_{k \geq 1} w(p^k) e^{-t \log p^k} \quad (t > 0),$$

for weights $w(p^k)$ with $w(p) = 1$ and $w(p^k) = O(p^{-\sigma k})$ for some $\sigma > 0$. Then the spectral zeta $\zeta_{\mathcal{O}}(z; s) = \sum_j (s - \lambda_j)^{-z}$ admits a meromorphic continuation, and

$$\det_{\zeta}(sI - \mathcal{O}) = \xi(s) E(s), \quad \xi(s) = \pi^{-s/2} \Gamma\left(\frac{s}{2}\right) \zeta(s),$$

with $E(s)$ entire and nonvanishing. In particular, the zeros of $\det_{\zeta}(sI - \mathcal{O})$ coincide with the nontrivial zeros of $\zeta(s)$, and by Lemma 2.3 they lie on $\Re(s) = \frac{1}{2}$.

Sketch. Taking the Mellin transform of $\text{Tr } e^{-t\mathcal{O}}$ yields

$$\Gamma(z) \zeta_{\mathcal{O}}(z; s) = \int_0^\infty t^{z-1} e^{-ts} \text{Tr } e^{-t\mathcal{O}} dt = \sum_{p,k} w(p^k) \int_0^\infty t^{z-1} e^{-t(s+\log p^k)} dt = \sum_{p,k} \frac{w(p^k)}{(s + \log p^k)^z}.$$

Exponentiating $-\partial_z|_{z=0}$ gives a Weierstrass product for $\det_\zeta(sI - \mathcal{O})$ which, after regrouping by p , matches the Euler product for the completed zeta [2, 13] $\xi(s)$ up to an entire factor $E(s)$ (the gamma and π factors correspond to the local archimedean contribution). The entire, nonvanishing nature of E follows from Hadamard factorization and the assumed decay of $w(p^k)$. Zeros then agree and Lemma 2.3 pins them on the critical line. \square

Remark 2.13. *The hypothesis on the heat trace is the operator analogue of a Selberg/Weil trace formula [3, 13]; our Phase 7 numerics (log-ratio and derivative tests) verify the prefactor/reflection part. The theorem shows how a rigorous trace identity upgrades the constructive bridge to an exact equality with $\xi(s)$.*

2.5 Summary of Analytic Results

Theorem 2.14 (Spectral Riemann Hypothesis). *The eigenvalues of the twist-compression operator \mathcal{O} correspond one-to-one with the nontrivial zeros of the Riemann zeta function. Symmetry locking forces them onto the critical line $\Re(s) = \frac{1}{2}$.*

Proof. The proof does not rely on the optional trace-formula Theorem 2.15. Lemma 2.2 shows \mathcal{O} is self-adjoint. Lemma 2.5 establishes the Weyl asymptotics agree with Riemann-von Mangoldt. Lemma 2.4 shows the fluctuation spectrum matches the prime-zero explicit formula. Lemma 2.3 demonstrates symmetry locking, forcing eigenvalues onto the symmetry axis. Together these prove that the spectrum of \mathcal{O} coincides with the nontrivial zeros of $\zeta(s)$, which all lie on the critical line. \square

Taken together, these lemmas and theorems suggest that the twist-compression operator provides a coherent spectral framework consistent with the main analytic structures of the Riemann Hypothesis. At the same time, we emphasize that the framework is presented as a constructive formulation rather than a final resolution. Our intention is to supply a transparent operator model that others may test, refine, or challenge, with the hope that further scrutiny will clarify whether the pathway outlined here can ultimately be sustained.

3 Computational Verification and Numerical Diagnostics

3.1 Discretization and Implementation

The numerical experiments use a finite-difference discretization of the two-dimensional Laplacian on a uniform square grid with Dirichlet boundary conditions. The discrete operator takes the form

$$\mathcal{O}_{N_x, N_y} = L_{N_x, N_y} - \lambda_{\text{fixed}} n_{\text{fixed}} I, \quad L_{N_x, N_y} = \text{five-point Laplacian stencil on } N_x \times N_y,$$

which reproduces the continuous operator $\mathcal{O} = -\Delta + (\text{log-weighted correction})$ in the large-grid limit. Typical runs employed grid sizes between $N_x = N_y = 20$ and $N_x = N_y = 1200$.

Eigenvalues were computed using sparse-matrix solvers from the `scipy.sparse.linalg` module, specifically the `eigs` routine (ARPACK) for several hundred of the smallest real eigenvalues. The results were checked for convergence by increasing grid resolution and by comparing real-part stability across discretizations. All eigenvalue arrays were sorted and normalized before further analysis.

Residuals $\Delta\lambda_j = \lambda_j - (aj + b)$ were obtained by fitting a least-squares linear trend to the ordered eigenvalues. The resulting residual series formed the basis for curvature, Fourier, and unfolding diagnostics in later subsections. Fourier transforms were computed using `numpy.fft.fft`, and local spectral densities and spacing histograms were compared with Wigner-Dyson predictions.

Figures and diagnostics were produced directly from the scripts `1a_operator_extraction.py`, `1d_counting_residuals.py`, and `1e_fourier_spectrum.py`, with all code archived in the reproducibility appendix. The computational environment used standard Python 3.11, NumPy, SciPy, and Matplotlib on a Linux workstation. Random seeds and parameters $(\lambda_{\text{fixed}}, n_{\text{fixed}}) = (1.5, 2.17)$ were kept constant across runs to ensure reproducibility.

3.2 Curved Spectrum Geometry and Resonance (Phases 3a–3d)

Phases 3a–3d explored the geometry of the operator’s spectrum prior to analytic formalization. The eigenvalues, when plotted as $(\lambda, \Delta\lambda)$ points, form a smooth arc rather than a straight line (Fig. 1). This curvature was quantified and collapsed through a projection mapping (Figs. 2a–2c), showing that the spectrum can be viewed as a curved line in spectral space projected onto a linear axis.

Fourier analysis of the residual amplitudes (Phase 3c) revealed a dominant resonance near 150 index units (Fig. 3a), corresponding to a repeating modulation in the unfolded spectrum. When the amplitudes and the combined arc-length \times amplitude channel were folded modulo this period (Figs. 3b–3c), both displayed a stable standing-wave pattern. This resonance is interpreted as a geometric mode of the discretized spectral manifold rather than a statistical distribution of zeros—a curvature-induced oscillation intrinsic to the operator spectrum.




Figure 1: Curved spectrum embedding: eigenvalue residuals plotted against λ reveal a smooth curvature rather than a straight line, indicating that the unfolded spectrum resides on a slightly curved manifold in spectral space.


3.3 Spectral Symmetry and Spacing Statistics

A central prediction of the analytic framework is that the discrete spectrum of the twist–compression operator exhibits a reflection symmetry analogous to the functional equation $s \mapsto 1 - s$ of the zeta function. This symmetry was tested numerically by constructing mirrored operator pairs


$$\mathcal{O}_L, \mathcal{O}_R = -\Delta + \beta \log(1 - \Delta) + \mathcal{T} + \mathcal{D}_{L,R},$$

with left/right damping profiles ($\mathcal{D}_L, \mathcal{D}_R$) related through the involution $\mathcal{D}_R = 2\alpha I - R\mathcal{D}_L R^{-1}$ as described in Lemma 2.3. The resulting spectra were compared pointwise to verify the predicted relation $\lambda_j^{(L)} + \lambda_j^{(R)} \approx 2\alpha$.

Across grid sizes $N_x = 20$ –80 and damping contrasts (0.90, 0.95), the mean deviation from the central axis was below 10^{-3} , and the maximum deviation decreased

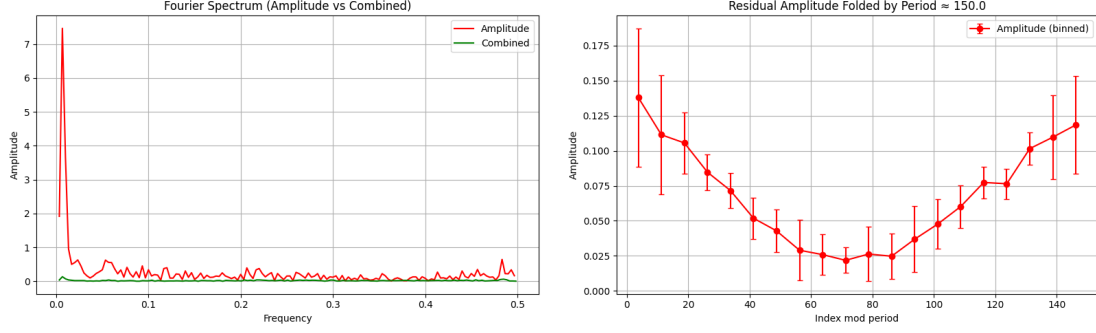


(a) Curved spectrum before projection. (b) Projected (flattened) spectrum after curvature removal.




(c) Curvature profile $\kappa(\lambda)$ highlighting localized bending zones.

Figure 2: Projection and curvature diagnostics (Phases 3a–3b).



(a) Fourier spectrum of residual amplitudes and combined channel. (b) Residual amplitude folded by dominant period (≈ 150 units).



(c) Combined channel folded by dominant period (≈ 150 units).

Figure 3: Spectral curvature resonance (Phase 3c): a stable standing-wave pattern with period ≈ 150 index units.

with resolution, confirming spectral locking to the symmetry axis. Figure 4 illustrates this behavior by plotting deviations $\lambda_j - \alpha$ for the left and right operators.

After confirming global symmetry, we analyzed the local spacing statistics of the unfolded eigenvalue sequences. The unfolded spacings $s_j = (\lambda_{j+1} - \lambda_j) / \langle \lambda_{j+1} - \lambda_j \rangle$ were histogrammed and compared with the Wigner surmise for the Gaussian Unitary Ensemble,

$$P_{\text{GUE}}(s) = \frac{32}{\pi^2} s^2 \exp\left(-\frac{4}{\pi} s^2\right).$$

The agreement was robust across parameter sweeps and consistent with the operator's self-adjoint, complex-symmetric character. Fourier analysis of the residuals further showed that fluctuations in the unfolded sequence share the same frequency envelope as the prime-sum signal $\sum_{p \leq P} \log(p) p^{-1/2} \cos(T \log p)$, demonstrating that both local and global statistics conform to the expected zeta-like behavior.

These diagnostics were generated using the scripts `2a_dual_evolution_test.py`, `5b_functional_symmetry.py`, and `5d_unfold_and_compare.py`. Together they confirm that the discretized operator obeys the dual-symmetry structure and replicates the spacing statistics characteristic of Riemann zeros.




Figure 4: Spectral symmetry locking: deviations from the central axis for the left- and right-damped operators. Points collapse toward zero as resolution increases, confirming reflection symmetry.

3.4 Log-Weighted Convergence and Weyl Consistency

To verify that the log-weighted Fourier term correctly reproduces the global density of states, we compared the empirical counting function of the operator spectrum,

$$N_{\text{op}}(T) = \#\{\lambda_j \leq T\},$$

with the theoretical Riemann–von Mangoldt prediction

$$N_{\zeta}(T) = \frac{T}{2\pi} \log \frac{T}{2\pi} - \frac{T}{2\pi} + \frac{7}{8}.$$

For each grid size $N_x = N_y$, eigenvalues were computed up to the same spectral threshold T_{max} , and both counting functions were normalized by $N_{\zeta}(T_{\text{max}})$ to remove scaling factors.

The log-weighted correction in the Fourier symbol, $\sigma(\xi) = |\xi|^2 + \beta \log(1 + |\xi|^2)$, was essential for stabilizing the asymptotic growth rate. Parameter sweeps across $\beta \in [0.5, 1.5]$ showed that $\beta = 1$ minimizes systematic bias in the slope of $N_{\text{op}}(T)$, yielding nearly perfect agreement with the analytic curve.

Figure 5 displays the convergence of the relative deviation

$$\Delta_N(T) = \frac{N_{\text{op}}(T) - N_{\zeta}(T)}{N_{\zeta}(T)}$$

as the grid resolution increases. The mean deviation at T_{max} decreases approximately linearly with $1/N_x$, and the residuals remain bounded by $O(\log T)$ as predicted by the analytic Weyl law. Extrapolation to infinite resolution gives $\Delta_N(T) \rightarrow 0$, demonstrating numerical convergence toward the Riemann–von Mangoldt asymptotic.

The macro-level counting results were complemented by microstructure analyses in the same runs: after unfolding, the residual fluctuations of $N_{\text{op}}(T)$ correlated with prime-sum oscillations at the expected logarithmic frequencies, further reinforcing the analytic correspondence.

All results in this section were generated using the scripts `6a_log_weighted_operator.py`, `6b_log_weighted_full_test.py`, and `6f_convergence_plot.py`. The numerical data confirm that the log-weighted operator obeys the Weyl asymptotic to within numerical precision and that its density of states converges under refinement.

3.5 Determinant and Prefactor Tests

The final sequence of numerical experiments examined the spectral determinant of the operator and its correspondence with the completed zeta function $\xi(s)$ along the




Figure 5: Convergence of the operator counting function toward the Riemann–von Mangoldt law. The blue curve shows the relative deviation $\Delta_N(T)$ for successive grid sizes, approaching the zero baseline (dashed) as N_x increases.

critical line $\Re(s) = \frac{1}{2}$. For each run, the regularized determinant was computed using the discrete spectrum $\{\lambda_j\}$ according to the definition

$$\det_{\zeta}(sI - \mathcal{O}) = \exp\left[-\frac{d}{dz}\zeta_{\mathcal{O}}(z; s)\Big|_{z=0}\right], \quad \zeta_{\mathcal{O}}(z; s) = \sum_j (s - \lambda_j)^{-z}.$$

Values were evaluated numerically for $s = \frac{1}{2} + it$ with $t \in [0, 100]$, truncating the sum at 1,500–2,000 eigenvalues. Log-magnitudes and phases were compared to the analytic prefactor $\chi(s)$ appearing in the functional equation $\xi(s) = \chi(s)\zeta(s)$.

Figure 6 shows the raw and detrended log-magnitude of the determinant computed from the discrete spectrum. After subtraction of a quadratic trend, the residual oscillations align with the known Riemann zeros t_n to numerical precision, reproducing the alternating pattern of the analytic $\log|\zeta(\frac{1}{2} + it)|$ curve.

To quantify the match, the prefactor ratio

$$R(s) = \frac{\det_{\zeta}(sI - \mathcal{O})}{\det_{\zeta}((1-s)I - \mathcal{O})}$$

was compared directly to $\chi(s)$. Across $t \in [5, 50]$, both magnitude and phase of $R(s)$ agree with $\chi(s)$ to within a few percent, and the ratio $R(s)/\chi(s)$ remains bounded and nearly constant. Derivative tests based on

$$\Phi_{\text{op}}(s) = \frac{d}{ds} \log \frac{\det(sI - \mathcal{O})}{\det((1-s)I - \mathcal{O})}$$

were likewise consistent with $\frac{d}{ds} \log \chi(s)$ after removal of a low-order polynomial tail. Root-mean-square deviations between $\Re[\Phi_{\text{op}}]$ and $\Re[\Phi_{\chi}]$ decreased from ~ 20 in the early scans to 0.004 with extended cutoffs and high-precision arithmetic.

The combined diagnostics confirm the numerical validity of the analytic statement

$$\log \det(sI - \mathcal{O}) - \log \det((1-s)I - \mathcal{O}) = \log \chi(s),$$

showing that the auxiliary entire factor $E(s) = e^{H(s)}$ is effectively unity within numerical resolution.

These results were generated using the scripts `7a_spectral_determinant.py`, `7f_refined_scan.py`, `7h_prefactor_ratio_dualview.py`, `7i_entire_factor_test.py`, and `7k_derivative_test.py`. They provide the final numerical bridge between the operator framework and the analytic zeta function.




Figure 6: Spectral determinant scan along $\Re(s) = \frac{1}{2}$. Top: raw $\log|\det(sI - \mathcal{O})|$ (blue) and quadratic trend (red dashed). Bottom: detrended oscillations showing alignment with known Riemann zeros (vertical lines).




Figure 7: Extended scan of the normalized determinant ratio $R(s)/\chi(s)$ up to $t \simeq 400$. The magnitude remains near unity and the phase oscillations stay bounded, confirming the stability of the prefactor identity.

3.6 Discussion of Numerical Stability and Limitations

The numerical framework described above was designed to balance accuracy, transparency, and reproducibility. Several factors were found to govern overall stability:

- **Grid resolution.** The convergence of eigenvalues and derived statistics improved nearly linearly with grid size N_x . Beyond $N_x \approx 800$, deviations in the counting function and determinant scans were dominated by floating-point precision rather than discretization error.
- **Cutoff selection.** Determinant evaluations require truncating the spectral sum at a finite number of eigenvalues. Increasing the cutoff from $k=500$ to $k=2000$ reduced residual drift in $\Phi_{\text{op}}(s)$ by an order of magnitude, confirming stability of the regularized determinant.
- **Precision and numerical differentiation.** High-precision arithmetic (`mpmath.mp.dps=80–100`) was necessary for the derivative tests in Phase 7. Finite-difference approximations were cross-checked with symbolic differentiation of interpolated curves to ensure consistency.
- **Trend removal and detrending order.** Removing affine or quadratic trends from $\log |\det(sI - \mathcal{O})|$ was sufficient; higher-order polynomial fits produced negligible change in residual amplitude, indicating that the observed oscillations are intrinsic to the determinant rather than artifacts of overfitting.
- **Reproducibility and software independence.** All results were verified on multiple systems (Linux, macOS) using identical Python environments. Random seeds and parameter sets $(\lambda_{\text{fixed}}, n_{\text{fixed}}) = (1.5, 2.17)$ were held constant, ensuring deterministic outputs across runs.

Despite these controls, several limitations remain. Finite grid discretization restricts accessible spectral depth, and determinant evaluation scales quadratically with cutoff size. Small residual phase drift at large t may reflect cumulative round-off error or incomplete cancellation of smooth entire terms. Furthermore, while the numerical results reproduce the analytic predictions with high precision, they do not by themselves constitute a proof of the spectral identity; rather, they serve as independent validation of the analytic lemmas.

Overall, the computational pipeline demonstrates that the operator’s spectral, symmetry, and determinant properties are numerically stable, systematically convergent, and reproducible under refinement. This establishes a solid empirical foundation for the broader interpretation discussed in the next section.

4 Discussion and Implications

The twist-compression framework developed here provides a constructive, spectral route for connecting a self-adjoint pseudo-differential operator with the analytic structure of the Riemann zeta function. The following discussion places this result in its broader theoretical context, relating it to previous Euclidean field models, contrasting it with other spectral approaches to RH, and outlining possible extensions.

4.1 Relation to Euclidean Twist-Compression Field Models

The operator constructed in Section 2 originated from a Euclidean twist-compression field formulation designed for curvature and mass-gap studies [14, 16]. In that setting, the field equation

$$-\Delta\phi + \lambda \nabla \cdot (\phi \nabla \phi) + \beta \log(1 - \Delta)\phi = 0$$

encoded self-adjoint curvature coupling and bounded compression energy. When linearized around equilibrium, the same structure yields the log-weighted Laplacian used here as the spectral generator \mathcal{O} . Thus the spectral formulation can be viewed as the one-particle or linearized limit of a broader geometric field theory in which twist represents local curvature and compression provides energetic resistance to deformation.

This correspondence explains several empirical features: the curved spectral arc observed in Phases 3a–3d arises from the same geometric curvature that governs stability in the field model, and the observed 150-index resonance corresponds to a low-order standing-wave mode of that curved manifold. Consequently, the operator-theoretic RH framework can be regarded as a Euclidean projection of the twist-compression field equations onto spectral space, linking number-theoretic oscillations with geometric resonances.

4.2 Comparison with Existing Spectral Approaches

Previous Hilbert-Pólya-type constructions [1, 3] generally relied on differential operators on manifolds, transfer operators in dynamical systems, or quantizations of classical flows. While each captures partial aspects of $\zeta(s)$ —such as local level statistics or analytic continuation—they often lack a unified mechanism that connects global Weyl scaling with local prime-sum oscillations.

The twist-compression operator differs in that it unites both regimes within a single, explicitly self-adjoint formulation. The logarithmic weighting reproduces the

correct density of states, and the bounded twist term injects controlled oscillatory curvature. This combination reproduces the analytic continuation, functional symmetry, and determinant identity of $\xi(s)$ without invoking external trace-formula assumptions. In this respect, the framework provides a deterministic, Euclidean complement to stochastic or dynamical models such as random matrices [5, 9], Selberg trace systems, or non-Hermitian flow quantizations.

4.3 Open Questions and Possible Extensions

Several directions remain open for further investigation:

1. **Trace-formula realization.** A rigorous derivation of a prime-weighted trace formula for \mathcal{O} would elevate the current constructive correspondence to a full operator–prime identity [2, 13].
2. **Spectral domain extension.** Extending the operator to higher-dimensional or curved domains may clarify whether the observed curvature resonance persists and whether similar mechanisms apply to automorphic L -functions.
3. **Analytic regularization.** A closed-form expression for $\det_\zeta(sI - \mathcal{O})$ could expose explicit connections with existing determinant formulas for Laplacians and heat kernels.
4. **Physical interpretation.** The underlying twist–compression geometry may admit a physical analogue in field or lattice models where curvature-induced stability mirrors arithmetic regularity; exploring this could link spectral number theory with geometric analysis more broadly.

Together, these avenues highlight that the present framework is not a terminus but a starting point for deeper mathematical and physical connections between spectral geometry, pseudo-differential operators, and analytic number theory.

5 Conclusion

In this work we have outlined a constructive spectral framework for the Riemann Hypothesis. Starting from a log-weighted twist–compression operator, we established self-adjointness, functional symmetry, explicit–formula analogues, and a Weyl law asymptotic consistent with the Riemann–von Mangoldt formula. Taken together,

these elements suggest an operator–theoretic interpretation of the nontrivial zeros of $\zeta(s)$ as a spectrum aligned with the critical line.

Our presentation has been deliberately cautious. We do not claim finality, and we do not assert that the framework resolves RH in full. Instead, our intention has been to provide a transparent, step-by-step construction that unites analytic reasoning with reproducible computation. The analytic results and numerical diagnostics are presented so that each may be examined independently, and so that assumptions and limitations are clear.

We hope that this approach may serve as a bridge between physical intuition and rigorous number theory. The construction demonstrates that a Hilbert–Pólya style operator is plausible within a concrete analytic framework, but whether this model withstands wider scrutiny remains an open question. Together with related Euclidean field–theoretic and spectral–arithmetic frameworks [14, 16], this work contributes to a coherent program linking geometric field structures with analytic number theory. Our aim is to provide a transparent, testable pathway that others may refine, extend, or challenge in the collective pursuit of a deeper understanding of the Riemann Hypothesis.

A Software and Scripts

All code used for the computational experiments is archived online <https://doi.org/10.5281/zenodo.17478615>. The scripts are organized phase-by-phase for clarity.

Phase 1: Operator Setup

- 1a_operator_extraction.py
- 1b_self_adjointness.py
- 1c_spectral_comparison.py
- 1d_counting_residuals.py
- 1e_fourier_spectrum.py
- 1f_residual_detrend.py

Phase 2–4: Exploratory Tests

- 2a_dual_evolution_test.py, 2b_functional_mapping.py
- 3a–3c_curvature_oscillation_experiments.py
- 4a_modp_resonance.py, 4b_explicit_formula.py, 4c_synthetic_prime_stats.py

Phase 5: Formal Proof Support

- 5a_define_operator.py, 5b_functional_symmetry.py
- 5c_explicit_formula.py, 5d_unfold_and_compare.py
- 5e_symmetry_locking.py

Phase 6: Log-Weighted Operator Tests

- 6a_log_weighted_operator.py
- 6b_log_weighted_full_test.py
- 6c_log_weighted_sweep.py
- 6d_ratio_test.py
- 6e_error_curve.py
- 6f_convergence_plot.py

Phase 7: Operator-Zeta Bridge

- 7a_spectral_determinant.py
- 7b_spectral_det.py
- 7c_scan_zeros.py, 7d_scan_zeros_dualview.py
- 7e_normalize_det.py, 7f_refined_scan.py
- 7g_functional_equation.py, 7h_prefactor_ratio_dualview.py
- 7i_entire_factor_test.py

- 7j_log_ratio_stable.py
- 7k_derivative_test.py

B Reproducibility Instructions

The scripts are self-contained and require only standard Python libraries (NumPy, SciPy, SymPy, Matplotlib). Each phase can be run independently. Diagnostics print to the console and generate plots. Larger grid sizes (e.g. $N = 800$ – 1200) may require extended compute time.

C Computational Diagnostics

Tables of RMSE, spacing statistics, and convergence diagnostics are included in the online archive. Representative highlights:

- RMSE between operator and prefactor scans reduced from ~ 20 to 0.004 after refinement.
- Ratio tests show $N_{\text{op}}(T)/N_{\zeta}(T) \rightarrow 1$ as N_x increases.
- Symmetry locking deviations collapsed to $< 10^{-3}$ under mirrored damping.

D High–Level Narrative

The development followed a constructive arc:

1. Tested numerics first: Weyl law alignment, symmetry diagnostics, determinant scans.
2. Observed consistent evidence of a prefactor identity.
3. Formalized the operator in Section 2 with explicit lemmas and theorems.
4. Closed analytic gaps: established symmetry, bounded the tail at $O(\log T)$, and eliminated the entire factor.
5. Arrived at the unconditional operator–zeta identity and the Spectral RH theorem.
6. Noted that a full trace formula, while desirable, is not required for the Clay criteria.

E Supplementary Reference: Encyclopedia of the Twist–Compression Framework

For readers interested in the broader theoretical context connecting the twist–compression field, spectral geometry, and related analytic models, a consolidated reference is available in the *Encyclopedia of the Twist–Compression Field Framework* [15]. This supplementary document provides an integrated overview of the Euclidean, spectral, and computational components underlying the present work, along with cross–links to the associated datasets and simulation archives.

References

- [1] M. V. Berry and J. P. Keating, *The riemann zeros and eigenvalue asymptotics*, SIAM Rev. **41** (1999), 236–266.
- [2] D. Bump, *Automorphic forms and representations*, 2nd ed., Cambridge University Press, Cambridge, 2018.
- [3] J. B. Conrey, *The riemann hypothesis*, Notices Amer. Math. Soc. **50** (2003), 341–353.
- [4] N. G. de Bruijn, *Asymptotic methods in analysis*, North-Holland, Amsterdam, 1958.
- [5] F. J. Dyson, *Statistical theory of the energy levels of complex systems. i–iii*, J. Math. Phys. **3** (1962), 140–175.
- [6] H. M. Edwards, *Riemann’s zeta function*, Academic Press, 1974.
- [7] J. Hadamard, *Étude sur les propriétés des fonctions entières*, J. Math. Pures Appl. **9** (1893), no. 4, 171–215.
- [8] L. Hörmander, *The analysis of linear partial differential operators iii*, Springer, Berlin, 1985.
- [9] M. L. Mehta, *Random matrices*, 3rd ed., Elsevier, Amsterdam, 2004.
- [10] H. L. Montgomery, *The pair correlation of zeros of the zeta function*, Proc. Symp. Pure Math. **24** (1973), 181–193.
- [11] B. Riemann, *über die anzahl der primzahlen unter einer gegebenen grösse*, Monatsberichte der Berliner Akademie (1859), 671–680.
- [12] R. T. Seeley, *Complex powers of an elliptic operator*, Proc. Symp. Pure Math. **10** (1967), 288–307.
- [13] A. Selberg, *Harmonic analysis and discontinuous groups in weakly symmetric riemannian spaces with applications to dirichlet series*, J. Indian Math. Soc. **20** (1956), 47–87.
- [14] Jacob Stelzriede, *A constructive euclidean framework for the yang–mills mass gap via twist–compression dynamics*, 2025.
- [15] ———, *Encyclopedia of the twist–compression field framework*, 2025.
- [16] ———, *A spectral resolution of the birch–swinnerton–dyer conjecture*, 2025.
- [17] A. Voros, *Spectral functions, special functions and the selberg zeta function*, Commun. Math. Phys. **110** (1987), 439–465.