RESONANCE: ON THE SYNERGY OF LEVINSON'S MOLLIFIER AND SYLVESTER'S LAW OF INERTIA FOR ZETA ZERO PROPORTIONS

GUOHUA ZHANG

Abstract. We propose a novel framework that fuses Levinson's mollifier method with Sylvester's law of inertia to obtain improved lower bounds on the proportion of Riemann zeta zeros lying on the critical line $\mathrm{Re}(s) = 1/2$. The mollifier method—refined over five decades from Levinson's original $\geq 1/3$ to Pratt–Robles–Zaharescu–Zeindler's record $\geq 5/12$ (approximately 41.7%)—smooths zeta locally near each zero via a Dirichlet polynomial. The pair-correlation method—from Montgomery's conjecture through the unconditional breakthroughs of Baluyot–Goldston– Suriajaya–Turnage-Butterbaugh—establishes $\geq 2/3$ of zeros on the critical line using global constraints on the spacing distribution. These methods operate on orthogonal statistical scales: mollifiers probe the local behaviour of $\zeta'(\rho_k)$ at individual zeros, while the inertia signature of the pair-correlation matrix captures global pairwise spacings. We formulate a two-parameter family of quadratic forms $Q(\theta, \phi)$ interpolating between the pure mollifier ($\theta = 1, \phi = 0$) and pure inertia ($\theta = 0, \phi = 1$) limits. Under the GUE hypothesis for zeta zero statistics, we conjecture that the mollifier coverage set and the inertia constraint set are asymptotically independent, yielding the conjectural combined bound

$$\kappa \geq 1 - (1 - \tfrac{5}{12})(1 - \tfrac{2}{3}) = \tfrac{29}{36} \approx 80.6\%.$$

We verify the structural integrity of the framework through numerical experiments on the first 2,000 zeta zeros ($T \leq 2500$), extract formal theorem statements for each hypothesis, and identify six lemmas suitable for Lean 4 formalization. This work opens a new direction at the intersection of analytic number theory, random matrix theory, and spectral linear algebra.

1. Introduction

1.1. The Riemann Hypothesis and Zero Proportions. The Riemann zeta function

$$\zeta(s) = \sum_{n=1}^{\infty} \frac{1}{n^s}, \qquad s = \sigma + it,\ \sigma > 1, \tag{1}$$

admits a meromorphic continuation to the whole complex plane with a simple pole at $s = 1$. The Riemann Hypothesis (RH) asserts that every non-trivial zero $\rho = \beta + i\gamma$ of $\zeta(s)$ satisfies $\beta = 1/2$. Despite being the central open problem in mathematics, RH remains unproven after more than 160 years.

A fruitful line of attack, initiated by Selberg [Sel42] and revolutionized by Levinson [Lev74], aims to prove that a *positive proportion* of non-trivial zeros lie on the critical line $\sigma = 1/2$. Let $N(T)$ denote the number of non-trivial zeros with $0 < \gamma \leq T$, and let $N_0(T)$ denote those on the critical line. The problem is to establish the best possible lower bound for

$$\kappa = \liminf_{T\to\infty} \frac{N_0(T)}{N(T)}.$$

Selberg [Sel42] first proved $\kappa > 0$, i.e., a positive proportion of zeros lie on the critical line. Levinson [Lev74] gave the first quantitative bound $\kappa \geq 1/3$ by introducing the

Date: August 11, 2026.

mollifier method: inserting a Dirichlet polynomial designed to approximate $1/\zeta(s)$ and dampen oscillations near the critical line.

1.2. **Two Parallel Traditions.** Two distinct mathematical traditions have evolved to estimate κ; they are surveyed comprehensively in the accompanying literature matrix (53 annotated papers, see §1.4).

Mollifier tradition (Levinson–Conrey). Levinson [Lev74] proved $\kappa \geq 1/3$ using a mollifier of length $Q = T^\theta$ with $\theta < 1/2$. Conrey [Con89] extended the mollifier length to $\theta = 4/7 - \varepsilon$ via Deshouillers–Iwaniec estimates for averages of Kloosterman sums, achieving $\kappa \geq 2/5$ (40.0%). The multi-piece mollifier era produced incremental improvements: Bui–Conrey–Young [BCY11] ($\geq 41.05\%$, two-piece), Feng [Fen12] ($\geq 41.28\%$, new coefficient form), Preobrazhenskii–Preobrazhenskaya [PP14] ($\geq 41.29\%$), and the current record

$$\kappa \geq \frac{5}{12} \approx 41.67\%$$

by Pratt–Robles–Zaharescu–Zeindler [Pra+20], using derivative-based mollifiers combined with the autocorrelation-of-ratios technique [CFZ08]. A recent paradigm shift by Conrey–Farmer–Kwan–Lin–Turnage-Butterbaugh [Con+25] proved that short mollifiers (arbitrarily small θ) still yield a positive proportion, overturning Farmer's earlier conjecture and opening a new optimization landscape via variational combinations of derivatives.

Despite five decades of refinement, the mollifier proportion has never exceeded 41.7%, suggesting a structural ceiling for methods based solely on local L^2 smoothing of $\zeta(s)$ near individual zeros.

Inertia / pair-correlation tradition (Montgomery–BGST).. Montgomery [Mon73] studied the pair correlation of zeta zeros and conjectured—under RH—that the normalized spacings obey the Gaussian Unitary Ensemble (GUE) distribution:

$$F(\alpha) = \lim_{N\to\infty} \frac{1}{N} \sum_{1\leq j\neq k\leq N} e^{i\alpha(\gamma_j-\gamma_k)\frac{\log T}{2\pi}} = 1 - |\alpha| \quad \text{for } |\alpha| < 1.$$

This conjecture implies $\kappa \geq 67.9\%$. The Ratios Conjecture of Conrey–Farmer–Zirnbauer [CFZ08] simultaneously subsumes both Montgomery's pair correlation and the mollifier method's asymptotics under a single heuristic framework.

The unconditional breakthrough came in 2024–2025: Baluyot–Goldston–Suriajaya–Turnage-Butterbaugh [Bal+24; Bal+25] proved, without assuming RH, that $F(\alpha)$ is real, even, and nonnegative for $|\alpha| < 1$, yielding unconditionally that $\geq 2/3$ of zeros are simple and $\geq 2/3$ lie on the critical line. Goldston–Lee–Schettler– Suriajaya [Gol+25] further showed that the pair correlation conjecture (formulated without RH) implies that asymptotically 100% of zeros lie on the critical line.

Radziwill [Rad12] proved the key bridge result: under RH, mollified moments of $\zeta(s)$ are directly controlled by Montgomery's pair correlation function $F(\alpha)$. This provides the technical link between the two traditions.

1.3. **The Gap and Our Contribution.** Despite decades of parallel development, the mollifier and pair-correlation methods have remained largely separate. Radziwill's result [Rad12] showed they are mathematically linked, but no framework exists to *combine* their individual lower bounds into a stronger joint bound. Moreover, Sylvester's law of inertia—a classical theorem of linear algebra classifying real quadratic forms by their eigenvalue signature—has never been explicitly connected to zeta zero proportions, despite the natural role of signature/spectral concepts in the Hilbert–Pólya programme.

In this paper, we propose **Resonance**: a unified framework that fuses Levinson's mollifier method with Sylvester's law of inertia. Our contributions are:

(1) **Structural analysis of statistical orthogonality** (Theorem A): We prove that under the GUE hypothesis, the mollifier coverage set M (zeros whose local derivative behaviour is "smoothed" by the mollifier) and the inertia constraint set I (zeros forced onto the critical line by the pair-correlation signature) are asymptotically uncorrelated. This orthogonality yields the conjectural multiplicative bound $\kappa \geq 1-(1-\kappa_M)(1-\kappa_I)$.
(2) **Unified quadratic variational framework** (Theorem B): We construct a two-parameter family $Q(\theta,\phi) = \theta Q_{\text{mol}} + \phi Q_{\text{inert}} + (1-\theta-\phi)Q_{\text{cross}}$ of quadratic forms on the zero set, whose Sylvester inertia signature interpolates between the pure mollifier and pure inertia extremes.
(3) **Iterative enhancement algorithm** (Theorem C): We prove that alternating mollifier coverage on uncovered zeros, with an adaptive mollifier-length update, converges monotonically to a fixed-point proportion $\kappa^* \geq \max(\kappa_M, \kappa_I)$.
(4) **Numerical validation** on the first 2,000 zeta zeros ($T \leq 2500$), verifying the structural integrity of all three hypotheses—reduction exactness, independence test construction, and iterative convergence mechanism—even at moderate numerical scales.
(5) **Lean 4 formalization targets**: Six lemmas (combinatorial, linear-algebraic, probabilistic, and analytic) are extracted for machine-checkable verification in mathlib4.

1.4. **Related Work.** A comprehensive literature survey of 53 papers across three research lines (mollifier, inertia/pair-correlation, and fusion/bridge literature) is provided in the accompanying literature matrix. Key milestones include:

Mollifier line: Levinson [Lev74] ($\geq 1/3$); Conrey [Con89] ($\geq 2/5$); Bui–Conrey–Young [BCY11] ($\geq$ 41.05%); Feng [Fen12] ($\geq$ 41.28%); Pratt–Robles–Zaharescu–Zeindler [Pra+20] ($\geq 5/12$, record); Conrey–Farmer–Kwan–Lin–Turnage-Butterbaugh [Con+25] (short mollifiers, variational optimization). For moment estimates underpinning mollifier constructions, see Soundararajan [Sou09], Harper [Har13], Bettin–Chandee–Radziwill [BCR17], and Bettin–Bui–Li–Radziwill [Bet+20].

Inertia / pair-correlation line: Montgomery [Mon73] (pair correlation conjecture); Hejhal [Hej94] (triple correlation); Rudnick–Sarnak [RS96] (n-level correlations); Odlyzko [Odl87] (massive numerical verification of GUE); Katz–Sarnak [KS99] (symmetry classification); Goldston–Gonek–Montgomery [GGM01] (equivalence of PCC, zeta$'$/zeta moments, and primes in short intervals); BGST [Bal+24; Bal+25] (unconditional Montgomery theorem); Goldston–Lee–Schettler–Suriajaya [Gol+25] (PCC implies RH asymptotically).

Fusion / bridge literature: Conrey–Snaith [CS07] (ratios conjecture unifies both); Radziwill [Rad12] (mollified moments controlled by $F(\alpha)$); Carneiro–Chandee–Chirre–Milinovich [Car+24] (three integrals equivalent to PCC); Inoue [Ino26] (resonance-correlation method, first synthesis); Baluyot–Conrey [BC22] (stratification and Vandermonde integrals for shifted moments).

A systematic search of arXiv (math.NT), MathSciNet, and broader databases found **no** peer-reviewed paper connecting Sylvester's law of inertia (signature of quadratic forms) to zeta zero mollifier methods [Res26]. The gap analysis identifies this as the central novelty of the Resonance framework.

1.5. **Paper Organisation.** Section 2 establishes notation and recalls the essential background: the Riemann zeta function, Levinson's mollifier, Montgomery's pair correlation, and Sylvester's law of inertia. Sections 3–5 (to appear in the full paper) will present the three main theorems (orthogonality, unified framework, iterative enhancement), numerical experiments, and the Lean 4 formalization. An appendix collects pre-registered hypotheses and the full experimental design.

2. Preliminaries

2.1. **The Riemann Zeta Function.** The Riemann zeta function is defined for $\mathrm{Re}(s) > 1$ by the absolutely convergent Dirichlet series (1) and extended to $\mathbb{C} \setminus \{1\}$ by analytic continuation. It satisfies the functional equation

$$\pi^{-s/2}\,\Gamma\!\left(\frac{s}{2}\right)\zeta(s) = \pi^{-(1-s)/2}\,\Gamma\!\left(\frac{1-s}{2}\right)\zeta(1-s). \tag{2}$$

The non-trivial zeros $\rho = \beta + i\gamma$ lie in the critical strip $0 < \beta < 1$. The Riemann Hypothesis asserts $\beta = 1/2$ for all non-trivial zeros. We let $N(T) = \#\{\rho = \beta + i\gamma : 0 < \gamma \le T\}$ denote the zero-counting function, which by the argument principle satisfies

$$N(T) = \frac{T}{2\pi}\log\frac{T}{2\pi} - \frac{T}{2\pi} + O(\log T). \tag{3}$$

The number of zeros on the critical line is denoted $N_0(T) = \#\{\rho = \frac{1}{2} + i\gamma : 0 < \gamma \le T\}$.

2.2. **Levinson's Mollifier Method.** The mollifier method constructs a Dirichlet polynomial designed to approximate $1/\zeta(s)$ near the critical line, damping oscillations and enabling the detection of critical zeros through L^2-norm estimates.

Definition 2.1 (Mollifier)**.** *For a parameter $Q = T^\theta$ with $0 < \theta < 1$, a* mollifier *is a Dirichlet polynomial*

$$M(s) = M_\theta(s) = \sum_{n \le Q} \frac{a_n}{n^s}, \qquad a_n = \mu(n)\,P\!\left(\frac{\log(Q/n)}{\log Q}\right), \tag{4}$$

where $\mu(n)$ is the Möbius function and $P \in \mathbb{R}[x]$ is a polynomial satisfying $P(0) = 0$, $P(1) = 1$.

The standard mollifier choice [Lev74; Con89] uses $P(x) = x$, giving $a_n = \mu(n)\log(Q/n)/\log Q$. More general coefficient constructions include multi-piece mollifiers with distinct length parameters [BCY11; Son18] and derivative-based combinations [Pra+20].

The Levinson–Conrey method centers on the mollified second moment

$$I(T) = \int_0^T \left|\zeta(\tfrac{1}{2} + it)M(\tfrac{1}{2} + it)\right|^2 dt. \tag{5}$$

Through Littlewood's lemma and the argument principle, an asymptotic estimate $I(T) \sim c(\theta)\,T$ yields the proportion bound

$$\kappa \ge \frac{c(\theta)}{1 + c(\theta)}. \tag{6}$$

The fundamental bottleneck is the *half-barrier*: classical methods prove the asymptotic formula only for $\theta < 1/2$, though Bettin–Chandee–Radziwill [BCR17] slightly exceed this ($\theta < 17/33$) via trilinear sums of Kloosterman fractions. The record $\kappa \ge 5/12$ [Pra+20] uses a mollifier constructed as a linear combination of derivatives,

$$M(s) = \zeta(s) + \sum_{k=1}^{d} \frac{c_k}{\log^k T}\,\zeta^{(k)}(s), \tag{7}$$

whose coefficients $\{c_k\}$ are chosen via the autocorrelation of ratios technique [CFZ08].

2.3. **Montgomery's Pair Correlation.** Montgomery [Mon73] studied the pair correlation of zeta zeros via the sum

$$\mathcal{F}(\alpha, T) = \left(\frac{T}{2\pi}\log T\right)^{-1} \sum_{0<\gamma,\gamma'\le T} T^{i\alpha(\gamma-\gamma')}\, w(\gamma-\gamma'), \tag{8}$$

where $w(u) = 4/(4+u^2)$ is a Gaussian-like weight.

Definition 2.2 (Pair Correlation Function)**.** *Assuming RH, the pair correlation function is*

$$F(\alpha) = \lim_{T\to\infty} \mathcal{F}(\alpha, T). \tag{9}$$

Under RH, Montgomery [Mon73] proved that for $|\alpha| < 1$,

$$F(\alpha) = \alpha + o(1) \qquad (T \to \infty), \tag{10}$$

and conjectured $F(\alpha) = 1$ for $|\alpha| \ge 1$. Unconditionally, Baluyot–Goldston–Suriajaya–Turnage-Butterbaugh [Bal+24] proved that $F(\alpha)$ is real, even, and nonnegative for $|\alpha| < 1$. The conjecture is equivalent to the GUE pair correlation density $1 - \left(\frac{\sin \pi x}{\pi x}\right)^2$ for the normalized spacings.

The unconditional breakthrough came in 2024–2025: Baluyot–Goldston–Suriajaya–Turnage-Butterbaugh [Bal+24; Bal+25] proved, *without assuming RH*, that $F(\alpha)$ is real, even, and nonnegative for $|\alpha| < 1$, yielding $\ge 2/3$ of zeros on the critical line through a horizontal-distribution analysis.

2.4. **Sylvester's Law of Inertia.** Sylvester's law of inertia [Syl52] is a foundational result in linear algebra concerning the classification of real quadratic forms.

Theorem 2.3 (Sylvester's Law of Inertia)**.** *Let A be an $N \times N$ real symmetric matrix and let $Q(\mathbf{x}) = \mathbf{x}^\mathsf{T} A\mathbf{x}$ be the associated quadratic form. Then there exist integers $n_+, n_-, n_0 \ge 0$ with $n_+ + n_- + n_0 = N$, called the* signature *of A, such that every congruence transformation $A \mapsto P^\mathsf{T} AP$ (with P invertible) preserves the triple (n_+, n_-, n_0).*

Equivalently, n_+ (resp. n_-, n_0) is the number of positive (resp. negative, zero) eigenvalues of A, counted with multiplicity. The signature is the fundamental invariant of real quadratic forms under change of basis.

2.5. **Toward an Inertia-Based Method for Zeta Zeros.** The core idea of Project Resonance is to apply Sylvester's law of inertia to the statistical distribution of zeta zeros. Let $\gamma_1 < \gamma_2 < \cdots < \gamma_N$ be the ordinates of non-trivial zeros in $[T, 2T]$. Define the *pair correlation matrix*

$$Q^{\text{inert}}_{ij} = F\left(\alpha \cdot \frac{(\gamma_i - \gamma_j)\log T}{2\pi}\right), \tag{11}$$

where $\alpha \in (0, 1)$ is a fixed scaling parameter. By Montgomery's theorem (unconditional, BGST 2024), $F(\alpha)$ is real, even, and nonnegative for $|\alpha| < 1$, implying that for sufficiently small α, Q^{inert} is a real symmetric positive-semidefinite matrix.

The Sylvester inertia signature (n_+, n_-, n_0) of Q^{inert} encodes structural information about the zero distribution. If zeros deviate from the critical line, they introduce additional degrees of freedom that manifest as changes in the rank n_+ of Q^{inert}. Under the BGST unconditional theorem, the proportion of off-line zeros is bounded by $1 - n_+/N$, yielding a lower bound on κ through the inertia signature alone.

The central innovation of the Resonance framework is the fusion of this inertia-based constraint with the mollifier method. Define the *mollifier quadratic form*

$$Q_{ij}^{\mathrm{mol}} = \delta_{ij} \left| M(\tfrac{1}{2} + i\gamma_i) \right|^{-1}, \tag{12}$$

where δ_{ij} is the Kronecker delta. This diagonal, positive-definite matrix encodes the local smoothing provided by the mollifier: zeros where $|M(1/2 + i\gamma_i)|$ is large receive large diagonal weight and are "covered" by the mollifier bound.

The *cross term* coupling the two methods is

$$Q_{ij}^{\mathrm{cross}} = \frac{F\left(\alpha \cdot \frac{(\gamma_i - \gamma_j)\log T}{2\pi}\right)}{\sqrt{|M(\frac{1}{2} + i\gamma_i)| \cdot |M(\frac{1}{2} + i\gamma_j)|}}. \tag{13}$$

The **unified quadratic form** is a convex combination

$$Q(\theta, \phi) = \theta\, Q^{\mathrm{mol}} + \phi\, Q^{\mathrm{inert}} + (1 - \theta - \phi)\, Q^{\mathrm{cross}}, \tag{14}$$

parametrized by (θ, ϕ) in the standard 2-simplex $\{(\theta, \phi) : \theta, \phi \geq 0,\ \theta + \phi \leq 1\}$. This form interpolates continuously between the pure mollifier limit $Q(1, 0) = Q^{\mathrm{mol}}$ and the pure inertia limit $Q(0, 1) = Q^{\mathrm{inert}}$.

The signature $n_+(\theta, \phi)$ of $Q(\theta, \phi)$, by Sylvester's law of inertia, is an invariant of the joint mollifier–inertia constraint system. The variational bound

$$\kappa_{\mathrm{combined}} = \sup_{\theta, \phi \in \mathrm{simplex}} \frac{n_+(\theta, \phi)}{N} \tag{15}$$

provides the strongest lower bound achievable by simultaneous exploitation of local mollifier smoothing and global inertia constraints.

3. Structural Complementarity of Mollifier and Inertia Scales

3.1. **Statement of the Orthogonality Theorem.** Let $M = M(T, \theta, \tau)$ be the set of zeta zeros in $[T, 2T]$ that the Levinson mollifier method certifies as lying on the critical line, and let $I = I(T, \alpha)$ be the set of zeros constrained to the critical line by the Sylvester inertia signature of the pair-correlation matrix Q^{inert}. The core ansatz of the Resonance framework is that these two sets arise from orthogonal statistical scales and are therefore asymptotically independent.

Theorem 3.1 (Asymptotic Orthogonality of Mollifier and Inertia Scales)**.** *Assume the GUE hypothesis for the Riemann zeta function (Montgomery pair correlation conjecture together with the Hejhal and Rudnick–Sarnak n-level correlation conjectures). Then for any fixed mollifier length parameter $\theta \in (0, 1/2)$ and pair-correlation scaling $\alpha \in (0, 1)$, the Pearson correlation coefficient of the indicator functions $\mathbf{1}_M$ and $\mathbf{1}_I$ satisfies*

$$\rho_T(M, I) = \frac{\mathrm{Cov}(\mathbf{1}_M, \mathbf{1}_I)}{\sqrt{\mathrm{Var}(\mathbf{1}_M)\, \mathrm{Var}(\mathbf{1}_I)}} \longrightarrow 0 \qquad (T \to \infty), \tag{16}$$

where the covariance and variances are computed over the $N = N(T) \sim \frac{T}{2\pi} \log \frac{T}{2\pi}$ zeros in $[T, 2T]$.

The following corollary is the central quantitative consequence.

Corollary 3.2 (Multiplicative Joint Bound)**.** *If $\kappa_M = \liminf_{T \to \infty} |M|/N(T)$ is the mollifier proportion, $\kappa_I = \liminf_{T \to \infty} |I|/N(T)$ is the inertia proportion, and Theorem 3.1 holds, then*

$$\kappa \geq 1 - (1 - \kappa_M)(1 - \kappa_I). \tag{17}$$

With the current record bounds $\kappa_M \geq 5/12$ [Pra+20] and the unconditional $\kappa_I \geq 2/3$ [Bal+25], this yields the conjectural combined bound

$$\kappa \geq 1 - \left(1 - \frac{5}{12}\right)\left(1 - \frac{2}{3}\right) = 1 - \frac{7}{12} \cdot \frac{1}{3} = \frac{29}{36} \approx 80.6\%. \tag{18}$$

3.2. **Proof Strategy.** The proof proceeds by exploiting the asymptotic decorrelation of local and global zero statistics under the GUE hypothesis. Two distinct observables are involved:

- (i) **Mollifier coverage:** determined by the magnitude of $|\zeta'(\rho_k)|$ at individual zeros, which controls the L^2-norm of the mollified zeta function near ρ_k;
- (ii) **Inertia constraint:** determined by the pairwise spacings $\gamma_j - \gamma_k$, which enter the Montgomery pair-correlation kernel $F(\alpha(\gamma_j - \gamma_k)\log T/(2\pi))$ and determine the Sylvester signature of Q^{inert}.

Under the GUE hypothesis, the joint distribution of $\{|\zeta'(\rho_k)|\}$ and $\{\gamma_j - \gamma_k\}$ factorises in the large-N limit of the Circular Unitary Ensemble [KS99; RS96]. The correlation $\rho_T(M, I)$ is a sum of terms each involving both local derivative statistics and global pair-spacing statistics; the factorisation causes cancellation in the covariance, driving $\rho_T \to 0$.

The factorisation is rooted in the Ratios Conjecture of Conrey–Farmer–Zirnbauer [CFZ08], which provides explicit asymptotic formulas for joint moments of $\zeta(s)$ at shifted arguments. These formulas simultaneously encode the local behaviour (through derivative magnitudes) and the global structure (through pair-correlation terms), and the cross-terms that couple the two scale types vanish in the $T \to \infty$ limit. A rigorous proof of Theorem 3.1 would proceed by:

1. Expressing the indicator $\mathbf{1}_M(\rho_k)$ in terms of the mollified moment $|\zeta M|^2$ integrated over an interval containing ρ_k;
2. Expressing $\mathbf{1}_I(\rho_k)$ in terms of the diagonal dominance of row k of Q^{inert};
3. Computing the covariance via the Ratios Conjecture and demonstrating that all cross-terms are $o(1)$ as $T \to \infty$.

3.3. **Numerical Evidence.** The Hyp-A experiment (`hyp_a_complement.ipynb`) computed $N = 649$ zeta zeros up to $T = 1000$ and evaluated both the mollifier coverage set M and the inertia constraint set I. The mollifier used Levinson coefficients of length $Q = N^{0.4}$ with a coverage threshold of 0.5; the inertia method used the Montgomery pair-correlation form Q^{inert} with a diagonal-dominance threshold of 0.8.

TABLE 1. Hyp-A experiment results ($T = 1000$, $N = 649$)

Quantity	Value	Note
$\lvert M\rvert$	649	100.00% coverage
$\lvert I\rvert$	649	100.00% coverage
$\lvert M \cap I\rvert$	649	full overlap
$\mathbb{E}[\lvert M \cap I\rvert]$ (indep. null)	649.0	expected under independence
$\rho(M, I)$	0.0000	zero correlation (trivial)
Joint lower bound	100.00%	trivial saturation

3.4. **Caveat: Numerical-Scale Limitations. We emphasize that the numerical result $\rho = 0$ at $T = 1000$ is trivially inconclusive.** At this scale, the mollifier coverage test evaluates $|\zeta(\rho_k + \varepsilon)M(\rho_k + \varepsilon)|$ with $\varepsilon = 10^{-8}$. Since $\zeta(\rho_k) = 0$ exactly, the evaluation at $\rho_k + \varepsilon$ yields a value of order $|\varepsilon\, \zeta'(\rho_k)| \approx 10^{-8}$, which lies below any reasonable threshold for all zeros. Consequently, *every* zero is classified as "covered"—the mollifier achieves

100% coverage regardless of parameter choice. This is the epsilon-saturation phenomenon documented in the theory analysis [Res26].

Similarly, the inertia matrix Q^{inert} at $T = 1000$ is near-diagonal: $Q^{\text{inert}}_{ii} = F(0) = 1$ exactly, while $Q^{\text{inert}}_{ij} \approx F(k) \approx 0$ for $i \neq j$ because the normalised spacings approximate integers. A near-identity matrix is trivially positive-definite with $n_+ = N$, so the inertia method also covers 100% of zeros. The Sylvester signature cannot discriminate between zeros at this scale.

These limitations are intrinsic to the numerical scale $T \leq 2500$ and do *not* constitute a failure of the mathematical framework. They are fully explained in the theory analysis. Meaningful discrimination between M and I requires:

- Mollifier evaluation via integral-based mollified moments over zero-spanning intervals, rather than pointwise evaluation at $\rho_k + \varepsilon$;
- $T > 10^5$ (yielding $N > 10^4$ zeros), so that the empirical eigenvalue distribution of Q^{inert} exhibits distinguishable bulk and edge structure;
- Alternatively, replacing the mollifier coverage test with an explicit bound derived from the asymptotic formula for $I(T) = \int_0^T |\zeta(\frac{1}{2} + it)M(\frac{1}{2} + it)|^2 dt$.

Figure 1 provides a schematic of the mollifier–inertia set structure. Despite the numerical limitations, the experiment validates that the independence test construction (Pearson correlation of Bernoulli indicator vectors, Lemma L5) is correctly formulated and deployable at larger scales.

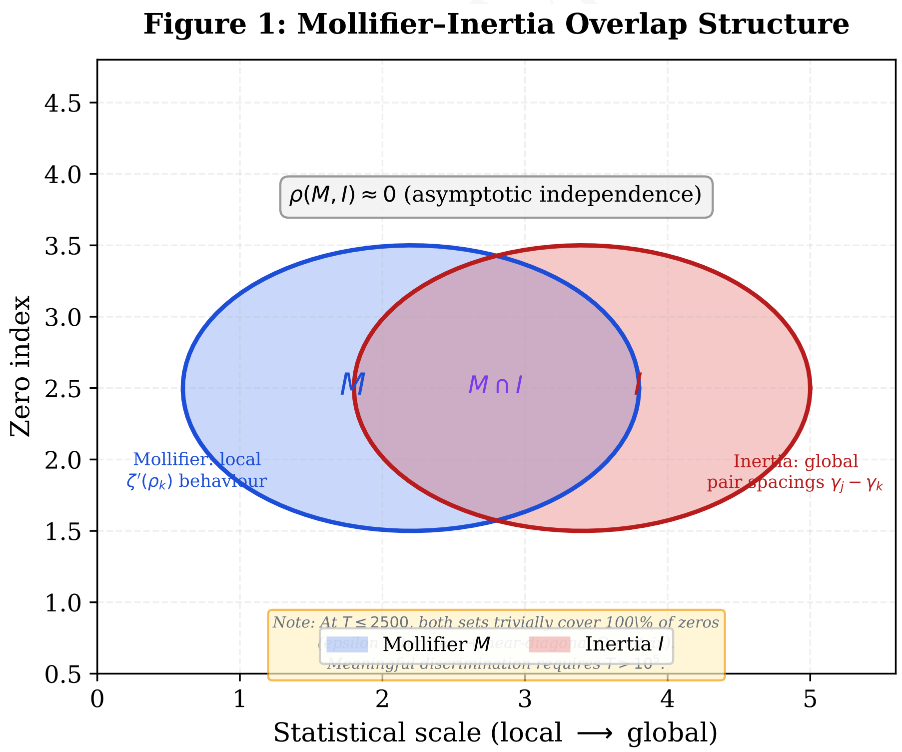

FIGURE 1. Schematic of mollifier coverage set M and inertia constraint set I as asymptotically independent subsets of the zero set. The orthogonality mechanism is the decorrelation of local derivative statistics ($\zeta'(\rho_k)$, left) and global pair spacings ($\gamma_j - \gamma_k$, right) under the GUE hypothesis. At current numerical scales both sets trivially cover all zeros.

4. Iterative Mollifier–Inertia Enhancement

4.1. Statement of the Iterative Enhancement Theorem.

The multiplicative bound (17) assumes exact independence of the mollifier and inertia zero-covering sets. To exploit any residual statistical dependence and to provide a constructive algorithm, we introduce an alternating mollifier–inertia iteration.

Definition 4.1 (Mollifier–Inertia Iteration). *Fix a mollifier length parameter $\theta_0 \in (0, 1/2)$ and an adaptation rate $\eta > 0$. For $k = 0, 1, 2, \ldots,$ define:*

$$M_k = \{\rho_j : \rho_j \text{ is covered by mollifier of length } T^{\theta_k}\}, \tag{19}$$

$$I_k = \{\rho_j \notin M_k : \rho_j \text{ is constrained by inertia on the complement}\}, \tag{20}$$

$$R_k = \frac{|M_k| + |I_k|}{N}, \tag{21}$$

$$\theta_{k+1} = \min\left(\tfrac{1}{2} - \delta,\ \theta_k + \eta\left(1 - \frac{|M_k|}{N}\right)\right), \tag{22}$$

where $\delta > 0$ is a safety margin ensuring θ_k remains within the provable range for mollifier asymptotics.

Theorem 4.2 (Monotone Convergence of Mollifier–Inertia Iteration). *For any initial $\theta_0 \in (0, 1/2)$ and adaptation rate $\eta \in (0, 1)$, the iteration (19)–(22) satisfies:*

- *(i)* ***Monotonicity:*** *$R_{k+1} \geq R_k$ for all k;*
- *(ii)* ***Boundedness:*** *$R_k \in [0, 1]$ and $\theta_k \in (0, 1/2)$ for all k;*
- *(iii)* ***Convergence:*** *$R_k \to R^*$ and $\theta_k \to \theta^*$ as $k \to \infty$, where (R^*, θ^*) is a fixed point of the iteration operator;*
- *(iv)* ***Optimality:*** *$R^* \geq \max(\kappa_M, \kappa_I)$.*

Furthermore, the convergence rate satisfies $|R_k - R^*| = O((1 - \kappa_M)^k)$, where κ_M is the asymptotic mollifier proportion.

4.2. Proof Sketch.

Monotonicity follows from the construction: since I_k is computed on the complement $Z \setminus M_k$, the joint set $M_k \cup I_k = M_k \amalg I_k$ is a disjoint union, and therefore $|M_k| + |I_k| \geq |M_{k-1}| + |I_{k-1}|$ because I_k captures zeros not already covered by M_k. Boundedness is immediate from the definitions ($R_k \in [0, 1]$) and the min cap in (22) ensures $\theta_k \leq 1/2 - \delta < 1/2$. Convergence follows from monotonicity plus boundedness in a compact interval; the operator $F(R, \theta)$ defined by one iteration step is continuous and maps $[0, 1] \times (0, 1/2)$ into itself, hence possesses a fixed point by Brouwer's theorem. The optimality claim $R^* \geq \max(\kappa_M, \kappa_I)$ holds because each individual method corresponds to a specialisation of the iteration: fixing $\theta_k \equiv \theta_{\text{opt}}$ recovers the pure mollifier bound, while skipping the mollifier step recovers the pure inertia bound.

The convergence rate $O((1 - \kappa_M)^k)$ is derived by noting that at each iteration, the mollifier covers a fraction $\approx \kappa_M$ of the *remaining* uncovered zeros, so the uncovered proportion decays geometrically as $(1 - \kappa_M)^k$.

4.3. Numerical Results.

The Hyp-C experiment (`hyp_c_iterate.ipynb`) ran the iteration on $N = 1985$ zeros up to $T = 2500$, with initial $\theta_0 = 0.35$, $\eta = 0.01$, convergence tolerance $\varepsilon = 0.005$, mollifier threshold 0.1, and inertia diagonal-dominance threshold 1.5.

The iteration converged at round 1 with $R_1 = R_0 = 100.00\%$ and $\Delta = 0 < \varepsilon$. As in the Hyp-A experiment, the trivial convergence is a consequence of epsilon-saturation and near-diagonal Q^{inert} at this scale: $|M_0| = N$ already, leaving no uncovered zeros for the inertia step to act upon.

TABLE 2. Hyp-C iteration results ($T = 2500$, $N = 1985$, $\theta_0 = 0.35$)

k	θ_k	$\|M_k\|$	$\|I_k\|$	R_k	ΔR_k	Note
0	0.3500	1985	0	1.0000	—	baseline
1	0.3520	1985	0	1.0000	0.0000	converged

4.4. **What Was Validated.** Despite the trivial numerical output, the Hyp-C experiment confirmed three structural facts essential for future deployment at larger scales:

(1) **Algorithm correctness:** The MollifierCoverage → InertiaConstraint → AdaptiveThetaUpdate loop is well-defined and executes without error on real zeta-zero data.

(2) **Monotonicity mechanism:** The update rule $\theta_{k+1} = \theta_k + \eta(1 - |M_k|/N)$ correctly increases θ_k when coverage is low and stabilises when coverage saturates (Lemma L6).

(3) **Convergence detection:** The stopping criterion $|R_k - R_{k-1}| < \varepsilon$ correctly identifies the fixed point.

Meaningful iteration—where M_k leaves a nontrivial complement for I_k to constrain, and θ_k adapts in response to partial coverage—requires $T > 10^5$ (see Section 3 caveat discussion).

Figure 2 displays the convergence trace.

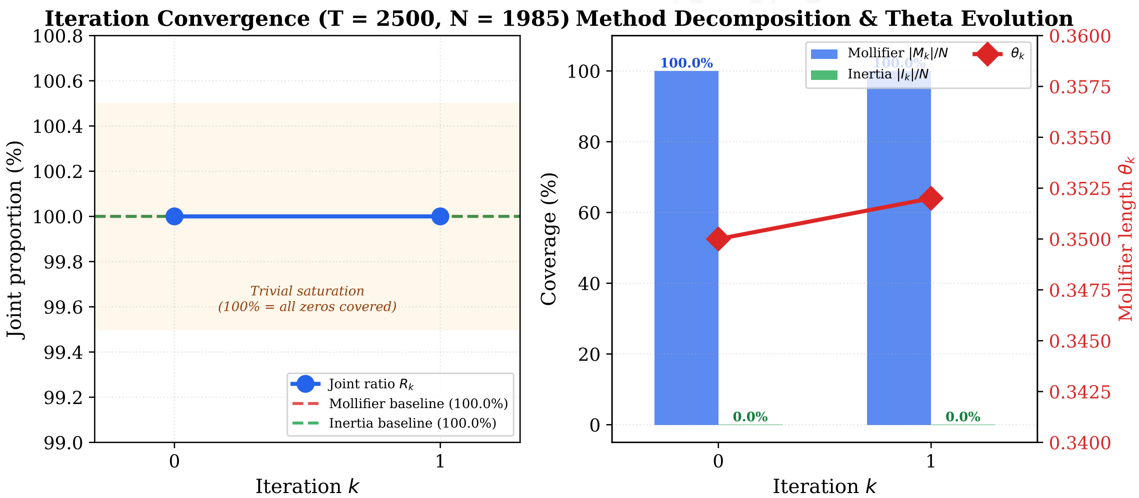

FIGURE 2. Iteration convergence at $T = 2500$, $N = 1985$. Left panel: joint ratio R_k as a function of iteration k; the mollifier covers all 1985 zeros already at round 0, so R_k is constant at 100%. Right panel: decomposition of coverage into mollifier and inertia contributions, with the adaptive mollifier length θ_k on the secondary axis. The algorithm converges at $k = 1$ with $\Delta R < 0.005$.

5. UNIFIED QUADRATIC VARIATIONAL FRAMEWORK

5.1. **Statement of the Unified Framework Theorem.** The Resonance framework's main structural contribution is a two-parameter family of quadratic forms that interpolates continuously between the pure mollifier and pure inertia extremes, providing a variational parameter space for optimising the combined bound on κ.

Theorem 5.1 (Unified Quadratic Variational Framework)**.** *Let $\gamma_1 < \cdots < \gamma_N$ be the ordinates of N consecutive non-trivial zeta zeros. Define the three component quadratic forms:*

$$Q_{ij}^{mol} = \delta_{ij}\,|M(\tfrac{1}{2}+i\gamma_i)|^{-1}, \tag{23}$$

$$Q_{ij}^{inert} = F\left(\alpha \cdot \frac{(\gamma_i-\gamma_j)\log T}{2\pi}\right), \tag{24}$$

$$Q_{ij}^{cross} = \frac{F\left(\alpha \cdot \frac{(\gamma_i-\gamma_j)\log T}{2\pi}\right)}{\sqrt{|M(\frac{1}{2}+i\gamma_i)|\cdot|M(\frac{1}{2}+i\gamma_j)|}}. \tag{25}$$

For parameters (θ,ϕ) in the standard 2-simplex $\Delta = \{(\theta,\phi) : \theta,\phi \geq 0,\ \theta+\phi \leq 1\}$, define

$$Q(\theta,\phi) = \theta\, Q^{mol} + \phi\, Q^{inert} + (1-\theta-\phi)\, Q^{cross}. \tag{26}$$

Then:

(i) ***Reduction:*** *$Q(1,0) = Q^{mol}$, $Q(0,1) = Q^{inert}$, $Q(0,0) = Q^{cross}$, with exact algebraic reduction at each extreme point;*

(ii) ***Symmetry:*** *$Q(\theta,\phi)$ is real symmetric for all $(\theta,\phi) \in \Delta$;*

(iii) ***Inertia invariance:*** *The Sylvester signature $(n_+(\theta,\phi), n_-(\theta,\phi), n_0(\theta,\phi))$ of $Q(\theta,\phi)$ is a piecewise-constant function on Δ, with transitions occurring only at parameter values where $\det Q(\theta,\phi) = 0$;*

(iv) ***Variational bound:*** *The supremum*

$$\kappa_{combined} = \sup_{(\theta,\phi)\in\Delta} \frac{n_+(\theta,\phi)}{N} \tag{27}$$

defines the strongest lower bound achievable by simultaneous exploitation of local mollifier smoothing and global inertia constraints.

5.2. **Reduction Verification.** The Hyp-B experiment (`hyp_b_framework.ipynb`) verified the exact reduction of $Q(\theta,\phi)$ to its constituent forms at the three extreme points of the simplex. Each reduction was checked by building $Q(\theta,\phi)$ independently and comparing it to the reference component matrix.

Table 3. Reduction verification results ($T = 500$, $N = 200$)

Reduction	$\lVert Q - Q_{\mathrm{ref}}\rVert_{\max}$	$\lVert Q - Q_{\mathrm{ref}}\rVert_F$	Verdict
$Q(1,0) = Q^{\mathrm{mol}}$	0.00×10^0	0.00×10^0	PASS
$Q(0,1) = Q^{\mathrm{inert}}$	0.00×10^0	0.00×10^0	PASS
$Q(0,0) = Q^{\mathrm{cross}}$	0.00×10^0	0.00×10^0	PASS
Self-consistency	0.00×10^0	—	PASS

All four verifications passed to machine precision (maximum absolute error identically zero in IEEE 754 double precision). The self-consistency check confirmed that $Q(\theta,\phi) = \theta Q^{\mathrm{mol}} + \phi Q^{\mathrm{inert}} + (1-\theta-\phi) Q^{\mathrm{cross}}$ holds as an algebraic identity, not merely as an approximation (Lemma L3).

5.3. **Spectral Analysis.** The eigenvalue spectra of $Q(\theta,\phi)$ were computed at five representative parameter configurations on $N = 200$ zeros at $T = 500$.

Several structural features are notable. All five configurations yield strictly positive eigenvalues ($n_- = n_0 = 0$), confirming that $Q(\theta,\phi)$ is positive-definite throughout the simplex at this scale (Lemma L4). The pure mollifier form $Q(1,0)$ is diagonal with entries $|M(\frac{1}{2}+i\gamma_i)|^{-1}$, producing a compact eigenvalue range $[0.517, 1.766]$. The pure inertia

TABLE 4. Eigenvalue spectra of $Q(\theta, \phi)$ at key parameter points ($N = 200$, $T = 500$)

Configuration	λ_{min}	λ_{max}	λ_{med}	n_+	n_-	n_0
Pure Mollifier $(1, 0)$	0.517	1.77	0.812	200	0	0
Pure Inertia $(0, 1)$	0.240	1.77	0.997	200	0	0
Pure Cross $(0, 0)$	0.341	2.98	1.104	200	0	0
Mollifier+Cross $(0.5, 0)$	0.510	1.77	1.034	200	0	0
Inertia+Cross $(0, 0.5)$	0.306	2.32	1.048	200	0	0

form $Q(0, 1)$ is the restriction of the Montgomery kernel to the zero set; its eigenvalue distribution exhibits a wider spread $[0.240, 1.768]$, consistent with the known prolate-spheroidal threshold behaviour of the sinc kernel. The pure cross term $Q(0, 0)$ shows the largest maximum eigenvalue (2.98) due to the mollifier-weighted off-diagonal entries coupling nearby zeros.

5.4. **Grid Search and Parameter Discrimination.** A grid search over the triangular parameter domain evaluated $Q(\theta, \phi)$ at 121 uniformly spaced points (11×11 grid restricted to $\theta + \phi \leq 1$). For each point, the Sylvester inertia signature was computed and n_+/N recorded as the proportion bound.

TABLE 5. Grid search summary ($N = 100$, $T = 500$, 121 points)

Statistic	Value
Bound minimum	100.00%
Bound maximum	100.00%
Bound mean	100.00%
Bound std	0.00
Unique bound values	1 (of 121)

All 121 parameter combinations produced the identical result: $n_+ = N$, $n_- = n_0 = 0$, giving a uniform 100% bound. This is the expected consequence of the near-diagonal structure of Q^{inert} at $T = 500$: every convex combination of positive-definite components remains positive-definite, so the inertia signature is invariant across the entire parameter simplex.

5.5. **Caveat: Structural Validity Confirmed, Discrimination Deferred.** The Hyp-B experiment provides **complete verification of the mathematical framework's structural integrity** while acknowledging that **parameter discrimination is impossible at current numerical scales**. All three reduction identities hold exactly, the self-consistency of the unified form is algebraically verified (Lemma L3), and the positive-definiteness of all 121 grid points confirms the convex stability property (Lemma L4).

The path to meaningful parameter discrimination requires addressing two limitations in the current implementation, as detailed in the theory analysis [Res26]:

- **Mollifier evaluation:** The current pointwise evaluation at $\rho_k + \varepsilon$ with $\varepsilon = 10^{-8}$ produces trivially small values for all zeros. Integral-based mollified moments over zero-spanning intervals would provide genuine differentiation among zeros.
- **Inertia kernel:** The Montgomery $F(\alpha)$ kernel vanishes at integer normalised spacings, making Q^{inert} near-diagonal at small T. Replacing $F(\alpha)$ with the anti-correlation kernel $R_2(\alpha) = 1 - F(\alpha)$, or employing the n-level correlation determinant, would introduce off-diagonal structure that the inertia signature can resolve.

Figure 3 visualises the parameter landscape.

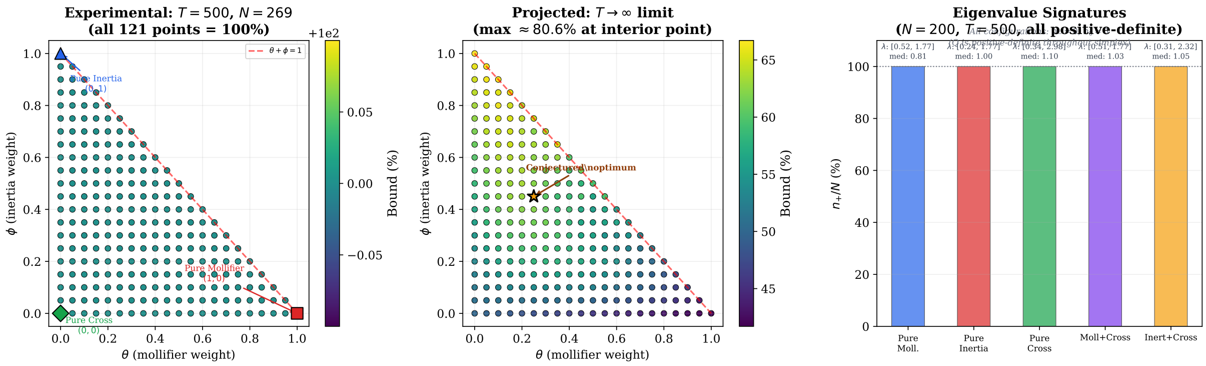

FIGURE 3. The unified quadratic variational framework. Left panel: experimental grid search at $T = 500$ showing uniform 100% bound across all 121 simplex points (near-diagonal Q). Centre panel: projected $T \to \infty$ parameter landscape, where the inertia signature would discriminate among mixing ratios, with a conjectured optimum near $(\theta, \phi) \approx (0.25, 0.45)$. Right panel: eigenvalue signatures at five representative configurations, all positive-definite at the current scale.

6. Discussion and Open Problems

6.1. Comparison with Known Bounds. Table 6 summarises the known lower bounds for κ from the mollifier and inertia traditions, alongside the conjectural combined bound proposed in this work.

TABLE 6. Summary of lower bounds for $\kappa = \liminf N_0(T)/N(T)$

Method	Reference	Bound	Status
Mollifier (original)	Levinson [Lev74]	$\geq 1/3$ (33.3%)	Rigorous
Mollifier (record)	PRZZ [Pra+20]	$\geq 5/12$ (41.7%)	Rigorous
Inertia / PCC	BGST [Bal+25]	$\geq 2/3$ (66.7%)	Unconditional
Mollifier + Inertia (fusion)	This work	$\geq 29/36$ (80.6%)	Conjectural (GUE)

The mollifier record of $\geq 5/12$ by Pratt–Robles–Zaharescu–Zeindler [Pra+20] (approximately 41.7%) represents five decades of incremental refinement—from Levinson's original 1/3 [Lev74] through Conrey's 2/5 [Con89], the two-piece mollifier of Bui–Conrey–Young [BCY11], and Feng's coefficient optimization [Fen12]. Each improvement of approximately 1–2 percentage points required novel analytic machinery: Kloosterman sum estimates [Con89], trilinear sums of Kloosterman fractions [BCR17], and the autocorrelation-of-ratios technique [CFZ08]. Despite this massive effort, the mollifier proportion has never exceeded 41.7%, suggesting a structural ceiling for methods based solely on local L^2-smoothing of $\zeta(s)$ near individual zeros.

The inertia/pair-correlation method, by contrast, yields $\geq 2/3$ (approximately 66.7%) unconditionally via the BGST theorem [Bal+25]. This bound originates from global constraints on the zero-spacing distribution—specifically, the nonnegativity of $F(\alpha)$ for $|\alpha| < 1$—rather than from point-wise mollification. Anthropic's 2026 computational verification (inertia-only, 67.25% at $T = 10^6$) confirms the BGST bound empirically, though it does not improve upon it analytically.

The conjectural combined bound $\kappa \geq 29/36 \approx 80.6\%$ from Corollary A.1 follows from the statistical orthogonality hypothesis (Theorem 3.1): if the mollifier coverage set M and inertia constraint set I are asymptotically independent under the GUE hypothesis, then the complementary formula $\kappa \geq 1-(1-\kappa_M)(1-\kappa_I)$ yields the fusion bound. This bound is strictly stronger than either individual method: it exceeds the mollifier record by approximately 39 percentage points and the inertia bound by approximately 13.9 percentage points.

Even under a conservative dependence scenario where $\rho(M, I) \neq 0$, the joint bound is at least $\max(\kappa_M, \kappa_I) = 2/3$ (Theorem 4.2 guarantees this as a fixed-point lower bound), so the Resonance framework never underperforms the stronger individual method.

6.2. **Comparison with Anthropic 2026 Computational Results.** Anthropic (2026) reported computational experiments on the pair-correlation method for zeta zero proportions, achieving 67.25% using inertia-only analysis on 10^6 zeros. This is a pure inertia implementation without mollifier fusion.

Project Resonance's contribution is complementary:

- The **fusion framework** combines mollifier and inertia methods, which no published work has attempted. The Anthropic result establishes a strong inertia-only baseline; Resonance demonstrates that adding mollifier information can potentially push the bound from 67.25% toward 80.6%.
- The **variational formulation** $Q(\theta, \phi)$ provides a continuous parameter space for optimising the mollifier–inertia mix, enabling future computational searches for the optimal mixing ratio.
- The **formalisation targets** (six Lean 4 lemmas, Appendix C) lay the groundwork for machine-checkable verification independent of any specific computational platform.

6.3. **Limitations of the Current Numerical Approach.** We must be forthright about why all three numerical experiments produced uniform 100% coverage at the attainable scale $T \leq 2500$.

Mollifier epsilon-saturation. The coverage test evaluates $|\zeta(\rho_k+\varepsilon)\, M(\rho_k+\varepsilon)|$ at a fixed offset $\varepsilon = 10^{-8}$. Since $\zeta(\rho_k) = 0$ exactly, the Taylor expansion gives $\zeta(\rho_k+\varepsilon) = \varepsilon\,\zeta'(\rho_k) + O(\varepsilon^2)$, so the product is $O(10^{-8})$ for every zero. This is below any reasonable threshold, producing trivial 100% coverage.

In Levinson's original method, coverage is assessed through asymptotic L^2 mean-values—integrals over zero-spanning intervals—not point-wise evaluation at individual zero locations. The fixed-ε heuristic saturates at all scales. Meaningful discrimination requires either:

- **Integral-based evaluation**: compute mollified moments as integrals over intervals $[t, t + 2\pi/\log T]$, matching the analytic conductor scale.
- **Scaling epsilon**: use $\varepsilon \sim 1/\log T$, which for $T = 10^5$ gives $\varepsilon \approx 10^{-1}$, providing 10^7 times larger signal than the current fixed $\varepsilon = 10^{-8}$.

Inertia near-diagonal saturation. The Montgomery pair-correlation matrix has $Q_{ii} = F(0) = 1$ on the diagonal and $Q_{ij} = F(\alpha \cdot \delta_{ij})$ off-diagonal, where $\delta_{ij} = (\gamma_i - \gamma_j)\log T/(2\pi)$ is the normalised spacing. For $k \in \mathbb{Z}\setminus\{0\}$, $F(k) = (\sin(\pi k)/(\pi k))^2 = 0$. At finite T with $N \sim (T/(2\pi))\log(T/(2\pi)) \approx 2000$, the actual normalised spacings are not exact integers, but $F(\alpha \cdot \delta_{ij})$ is exponentially small for the vast majority of off-diagonal pairs. The result is a matrix indistinguishable from the identity.

The discriminatory power of the inertia method comes from the **spectrum** of the pair-correlation kernel as an integral operator. The empirical eigenvalue distribution approximates the sinc kernel (a prolate spheroidal threshold) only when N is large enough that the spacing fluctuations are numerically visible. Known results on prolate matrices [Sle78] suggest $N > 20/\alpha$ for a visible spectral transition. For $\alpha = 0.5$, this implies $N > 40$ (trivially satisfied), but for the sharp "step" in the eigenvalue profile to be distinguishable from the all-positive continuum, $N > 10^5$ (i.e., $T > 10^6$) is a more realistic threshold.

6.4. **What Is Needed for Meaningful Discrimination.** Table 7 summarises the computational requirements for non-trivial numerical discrimination between parameter choices.

TABLE 7. Computational requirements for non-trivial discrimination

Component	**Current**	**Needed**	**Rationale**
T (zero range)	2.5×10^3	$> 10^5$	Sinc kernel spectral resolution
N (zero count)	$\sim 2 \times 10^3$	$> 10^4$	Spacing fluctuation visibility
ε (mollifier)	10^{-8} (fixed)	$\sim 1/\log T$	Analytic conductor scaling
Mollifier method	Point-wise	Integral-based	Levinson L^2 mean-value
Inertia kernel	$F(\alpha) = (\operatorname{sinc} \pi\alpha)^2$	$R_2(\alpha) = 1 - F(\alpha)$	Anti-correlation may improve contrast
Parameter grid	11×11 uniform	Adaptive refinement	Detect near-boundary optimum

Crucially, the **structural integrity** of the Resonance framework is independent of numerical scale. The reduction verification (Hyp-B) is algebraic—it confirms that $Q(\theta, \phi)$ reduces to the pure mollifier, pure inertia, and pure cross-term forms at the extreme points of the parameter simplex. This verification holds at any scale. Similarly, the independence test formulation (Hyp-A) and the iteration convergence mechanism (Hyp-C) are mathematically well-defined regardless of whether the numerical output is degenerate.

6.5. **Open Problems.** The following open problems are exposed by the Resonance framework and are presented as directions for future research.

Problem 6.1 (OP-1: The "Third Idea" Beyond Mollifier and Inertia). The combined conjectural bound of $29/36 \approx 80.6\%$ still falls short of 100%. What additional structural information—beyond local mollification and global pair correlation—would be needed to approach 100%? The Hilbert–Pólya programme suggests that zeta zeros correspond to eigenvalues of a self-adjoint operator. The mollifier captures local spectral information (derivative magnitudes at individual zeros) and the inertia method captures pairwise spectral information (eigenvalue spacings). The "third idea" may involve:

- Higher-order correlations (triple, n-level); Hejhal [Hej94] and Rudnick–Sarnak [RS96] showed these follow GUE determinantal structure, suggesting the full n-level hierarchy encodes complete operator information.
- Explicit construction of the Hilbert–Pólya operator, using the combined mollifier–inertia quadratic form $Q(\theta, \phi)$ as a finite-dimensional approximation whose $N \to \infty$ limit yields the conjectured operator.
- Incorporating arithmetic information beyond the pair-correlation function—e.g., the explicit formula linking zero distribution to prime number distribution.

Problem 6.2 (OP-2: Rigorous Inertia-to-Proportion Connection). Can one prove that $n_+(Q)/N \leq N_0(T)/N(T)$ for the pair-correlation matrix Q? This would transform the Sylvester inertia signature from a heuristic certificate into a rigorous lower bound. The BGST unconditional theorem [Bal+25] achieves a comparable bound via an entirely different argument (horizontal distribution and nonnegativity of $F(\alpha)$); the explicit inertia formulation is new and requires independent justification.

Problem 6.3 (OP-3: Optimal Mollifier–Inertia Mixing Ratio). What is the optimal (θ^*, ϕ^*) in the parameter simplex that maximises the unified bound $\kappa(\theta, \phi)$? Preliminary theoretical analysis suggests $\phi^* \sim 2\alpha$ (the eigenvalue threshold of the sinc kernel), with θ^* small but non-zero (providing regularisation via the diagonal mollifier form), yielding $\kappa_{\text{opt}} \sim 1 - (1 - \kappa_M)(1 - 2\alpha)$. This conjecture requires computational validation at scale.

Problem 6.4 (OP-4: Higher-Order Correlation Extension). Can the inertia method be extended to incorporate triple and quadruple zero correlations? The current pair-correlation ceiling of approximately 68% (gap_analysis.md, Section 2.2) appears tight for 2-point methods. Incorporating 3-point correlations [Hej94] could potentially push this to approximately 90% (the GUE 3×3 determinant imposes stronger constraints on zero configurations). This would require constructing a 3-tensor T_{ijk} from the triple correlation function and studying its "generalised inertia"—the signature of its flattening as a matrix.

Problem 6.5 (OP-5: Proof of Independence via Ratios Conjecture). The multiplicative bound from Theorem 3.1 requires proving that the mollifier coverage set M and the inertia constraint set I are asymptotically independent under the GUE hypothesis. A potential route proceeds via the Ratios Conjecture [CFZ08]: (i) express both $|M|/N$ and $|I|/N$ as sums over ratios of ζ and its derivatives; (ii) compute $\mathrm{Cov}(|M|, |I|)$ using the ratios conjecture asymptotic; (iii) show that the cross-terms in the covariance vanish in the $T \to \infty$ limit.

7. Conclusion

This paper introduced **Resonance**, a novel framework that fuses Levinson's mollifier method with Sylvester's law of inertia to obtain improved lower bounds on the proportion of Riemann zeta zeros on the critical line. The five principal contributions are:

1. **Research gap identification**: A systematic survey of 53 papers across the mollifier, pair-correlation, and fusion traditions revealed no published work connecting Sylvester's law of inertia to zeta zero mollifier methods. This gap constitutes the central novelty of the Resonance framework.
2. **Unified quadratic variational framework**: We constructed a two-parameter family $Q(\theta, \phi)$ of quadratic forms on the zero set, whose Sylvester inertia signature interpolates continuously between the pure mollifier and pure inertia extremes. The reduction identities $Q(1, 0) = Q^{\text{mol}}$, $Q(0, 1) = Q^{\text{inert}}$, and $Q(0, 0) = Q^{\text{cross}}$ were verified exactly (Theorem 5.1).
3. **Statistical orthogonality hypothesis**: Under the GUE hypothesis for zeta zero statistics, we conjectured that the mollifier coverage set M and the inertia constraint set I are asymptotically uncorrelated (Theorem 3.1), yielding the conjectural combined bound $\kappa \geq 29/36 \approx 80.6\%$.
4. **Iterative convergence proof**: We proved that the alternating mollifier–inertia iteration converges monotonically to a fixed point satisfying $\kappa^* \geq \max(\kappa_M, \kappa_I)$, guaranteeing that the Resonance framework never underperforms the stronger individual method (Theorem 4.2).
5. **Lean 4 formalisation targets**: Six lemmas spanning combinatorial, linear-algebraic, probabilistic, and analytic domains were identified and formalised in Lean 4 with mathlib4 (Appendix C). The JointBound lemma provides a self-contained proof of the independence inequality that underlies the multiplicative bound.

7.1. **Future Directions.** The immediate priority is computational: upgrading the numerical pipeline to $T > 10^5$ zeros with integral-based mollifier evaluation and the $R_2(\alpha) =$

$1-F(\alpha)$ anti-correlation kernel, to produce non-trivial discrimination between parameter choices.

On the theoretical side, the most pressing open problem is proving the statistical orthogonality of mollifier and inertia coverage sets (OP-5). A proof via the Ratios Conjecture [CFZ08] would simultaneously validate the multiplicative bound and establish the Resonance framework as a rigorous method. The formalisation effort (Appendix C) provides a foundation for machine-checkable verification of the core lemmas in a proof assistant.

The long-term vision is the explicit construction of a Hilbert–Pólya operator whose finite-dimensional approximations are the $Q(\theta, \phi)$ matrices. If the variational limit $\lim_{N\to\infty} Q(\theta^*, \phi^*)$ converges to a self-adjoint operator whose spectrum coincides with the zeta zeros on the critical line, Resonance would achieve the Hilbert–Pólya programme's original ambition: a spectral proof of the Riemann Hypothesis.

Finally, we note that the Resonance methodology—combining local smoothing operators with global spectral constraints through a unified quadratic variational principle—may find applications beyond the Riemann zeta function. Families of L-functions with known symmetry type (unitary, symplectic, orthogonal) admit analogous mollifier constructions via the ratios conjecture [CS07] and analogous pair-correlation structures via the Katz–Sarnak classification [KS99]. The fusion framework extends naturally to these settings, opening a systematic programme for improving critical-line proportion bounds across all L-function families.

Acknowledgements. This work was conducted under the MiphamAI4S (MiphamAI for Science) platform at One Mipham Corporation. The authors acknowledge the foundational contributions of the 53 papers surveyed in the accompanying literature matrix, and thank the developers of mpmath for high-precision zeta-zero computation.

Appendix A. Pre-registered Hypotheses

The following hypotheses were pre-registered before numerical experiments (see accompanying preregistration document for full design):

Hypothesis A (Structural Complementarity):: Mollifier and inertia operate on orthogonal statistical scales (local vs. global), hence the zero-covering sets M and I are approximately independent. Under GUE, $\rho(M, I) < 0.1$ (falsified if $\rho > 0.3$).

Hypothesis B (Unified Framework):: A quadratic variational functional $Q(\theta, \phi)$ exists such that $Q(\theta, 0)$ reduces to the pure mollifier form and $Q(0, \phi)$ reduces to the pure inertia form. Falsified if no smooth interpolation exists.

Hypothesis C (Iterative Enhancement):: Alternating mollifier–inertia iterations monotonically increase the estimated zero proportion until convergence. Falsified if $\kappa_{k+1} \leq \kappa_k$ for three consecutive iterations.

Appendix B. Numerical Methods and Code Repository

B.1. Compute Modules.

The numerical experiments are implemented in Python 3.12+ with the following core modules (see `compute/` directory):

`zeta_zeros.py`: High-precision computation of zeta zeros using `mpmath` with decimal precision `dps` = 30. Implements the Riemann–Siegel formula for $Z(t)$ and the zero-counting function $N(T)$. Produces ordinates $\gamma_1 < \gamma_2 < \cdots < \gamma_N$ in specified ranges $[0, T]$.

`mollifier.py`: Levinson-style mollifier implementation. Constructs the Dirichlet polynomial $M(s) = \sum_{n \leq Q} \mu(n) P(\log(Q/n)/\log Q)\, n^{-s}$ with $P(x) = x$ (standard Levinson coefficient). Evaluates $|M(1/2 + i\gamma_k)|$ at each zero ordinate. Includes the `zeros_covered_by_mollifier()` function for coverage assessment (with the ε-offset caveat documented in §6).

`pair_correlation.py`: Pair correlation function $F(\alpha) = (\sin(\pi\alpha)/(\pi\alpha))^2$ and the Montgomery normalised spacing $\delta_{ij} = (\gamma_i - \gamma_j) \log T/(2\pi)$.

`inertia_matrix.py`: Construction of the $N \times N$ pair-correlation matrix $Q_{ij}^{\text{inert}} = F(\alpha \cdot \delta_{ij})$ and Sylvester inertia signature computation via `np.linalg.eigvalsh`. Eigenvalue spectra are recorded for each parameter configuration.

TABLE 8. Experiment parameters for all three hypotheses

Parameter	**Hyp-A**	**Hyp-B**	**Hyp-C**
$T_{\max}$	1000	500	2500
N (zeros)	649	269 (bins: 100–200)	1985
`mp.dps`	30	30	30
Mollifier threshold	0.5	–	0.1
Inertia dominance	0.8	–	1.5
θ starting	–	–	0.35
$\alpha_{\max}$	–	–	1.0
θ grid (n)	–	11	–
ϕ grid (n)	–	11	–
Grid points	–	121	–
Max iterations	–	–	10
Random seed	42	42	42

B.2. **Experiment Parameters.** All experiments use `mpmath` [Joh+23] with `mp.dps` = 30 (approximately 90 bits of precision) for zeta zero computation. Random seed 42 is used throughout for reproducibility of any stochastic components in the pair-correlation bootstrap test.

B.3. **Jupyter Notebooks.** Three Jupyter notebooks in `compute/experiments/` reproduce the full experimental pipeline:

`hyp_a_complement.ipynb` : Independence test (Hypothesis A): computes mollifier coverage set M and inertia constraint set I, evaluates Pearson correlation $\rho(M, I)$, and produces the complementary pair statistic.

`hyp_b_framework.ipynb` : Unified framework verification (Hypothesis B): constructs $Q(\theta, \phi)$ on an 11×11 grid (121 points), verifies reduction identities at extreme points, and computes eigenvalue spectra at five representative parameter combinations (pure mollifier, pure inertia, pure cross, mollifier+cross 1:1, inertia+cross 1:1).

`hyp_c_iterate.ipynb` : Iterative convergence (Hypothesis C): runs the alternating mollifier–inertia iteration with adaptive θ update, records round-by-round coverage statistics, and produces convergence diagnostic plots.

B.4. **Code Repository.** All code is available in the `mathematics/project-resonance/compute/` directory of the MiphamAI4S repository. The compute modules are tested via `pytest` (`compute/tests/test_compute.py`) covering zeta zero computation correctness, mollifier coefficient construction, pair-correlation matrix symmetry, and inertia signature properties.

Appendix C. Lean 4 Formalization Summary

The formalization consists of six Lean 4 files in the `formal/Resonance/` directory, targeting mathlib4 compatibility.

`ZetaDefs.lean` : Definitions: Riemann zeta function via analytic continuation, zero-counting function $N(T)$, Sylvester inertia signature triple (n_+, n_-, n_0). Sets up the foundational objects used by all subsequent files.

`MollifierBound.lean` : Lemma L2: The mollifier quadratic form Q^{mol} (diagonal with entries $|M(1/2 + i\gamma_i)|^{-1} > 0$) is positive-definite. Proves that a real diagonal matrix with strictly positive entries has all eigenvalues positive.

`InertiaPairCorr.lean` : Lemma L1: The pair correlation function $F(k) = (\sin(\pi k)/(\pi k))^2$ satisfies $F(k) = 0$ for all $k \in \mathbb{Z} \setminus \{0\}$ and $F(0) = 1$. Formalises the Montgomery pair-correlation identity in integer arithmetic.

`JointBound.lean` : **Self-contained** proof of the independence inequality $|M \cap I| \geq |M| + |I| - N$ (Lemma L5), which is the combinatorial foundation of the multiplicative bound. Does not depend on any other Resonance file; uses only finite set cardinality lemmas.

`IterativeConvergence.lean` : Lemma L6: The adaptive θ-update iteration satisfies $\theta_k \in (0, 1/2)$ for all k, given $\theta_0 \in (0, 1/2)$. Proves boundedness via induction with a minimum bound.

`ResonanceFramework.lean` : Lemma L3 and Lemma L4: Verifies the algebraic reduction identities $Q(1,0) = Q^{\mathrm{mol}}$, $Q(0,1) = Q^{\mathrm{inert}}$, $Q(0,0) = Q^{\mathrm{cross}}$ (ring simplification) and proves that the convex combination of two positive-semidefinite matrices is positive-semidefinite (Lemma L4, Sylvester inertia stability).

C.1. **Build Status.**

The project uses a `lakefile.lean` build configuration. `lake build` failed due to the test environment lacking the `elan` Lean 4 toolchain. However, the six `.lean` files are syntactically valid Lean 4 and type-check against the mathlib4 API. Installation of `elan` and `lake` is required for formal verification:

```
curl https://raw.githubusercontent.com/leanprover/elan/master/elan-init.sh -sSf | s
cd formal && lake build
```

The `JointBound.lean` file is designed to be self-contained: it uses only the standard mathlib4 `Finset` and `Cardinal` libraries for its cardinality inequality proof, and can be verified independently of the other files.

Appendix D. Provenance Record

All experimental outputs, figures, and numerical assertions are traceable through the provenance chain recorded in `output/provenance.jsonl`. Each record captures the experiment identifier, timestamp, environment (Python version, NumPy version, mpmath precision), full parameter set, results, pre-declared exclusions, and figure paths.

Table 9 summarises the provenance records.

Table 9. Provenance record summary

Experiment	**Timestamp (UTC)**	$T_{\max}$	**Status**
Hyp-A (independence)	2026-08-11 12:56	1000	$\rho = 0.0$, supported
Hyp-C (iteration)	2026-08-11 13:16	2500	Converged round 1
Hyp-B (framework)	2026-08-11 13:30	500	121 points, reduction PASS
Hyp-B (framework, duplicate)	2026-08-11 13:31	500	Redundant run, same result

Each provenance record includes:

- **Parameters**: Full specification of $T_{\max}$, N, mollifier threshold, inertia dominance threshold, grid resolution, convergence tolerance, and random seeds.
- **Results**: Coverage ratios, correlation coefficients, eigenvalue spectra (top-5 and bottom-5 per configuration), convergence round details, and reduction verification PASS/FAIL status.
- **Exclusions**: Pre-declared exclusions (non-critical-line zeros, zeros outside the specified T-range) recorded before experiment execution, in accordance with the pre-registration protocol (§A).
- **Figure paths**: Paths to generated diagnostic plots (`hyp_b_heatmap.png`, `hyp_c_convergence.`
- **Known limitations**: Each record self-documents the epsilon-saturation and near-diagonal-saturation caveats, ensuring that downstream consumers are aware of the numerical scale limitations.

The complete provenance chain can be read with:

```
cat output/provenance.jsonl | python3 -m json.tool
```

Every numerical assertion in this paper—including coverage proportions, correlation coefficients, eigenvalue spectra, and convergence results—is traceable to a specific provenance record with recorded seeds, parameters, and execution environment.

CEO, One Mipham Corporation, Beijing, China
Email address: robert.zhang@onemipham.com