Toward the Riemann Hypothesis: A Superquadratic-Growth Framework for Zeta Moments and its Limits

Tamás Nagy, Ph.D.

tnagyphd@gmail.com

Working Paper

Build-std v2.0+410559b6c1 | 2026-07-10 13:14 | 80e763b6

Executive Summary (Non-Technical)

The **Riemann Hypothesis** — that all non-trivial zeros of the Riemann zeta function lie on a single vertical line in the complex plane — is the most important open problem in pure mathematics. It governs the distribution of prime numbers, and a proof would settle over a thousand theorems that currently assume it.

This paper develops a **framework** — not a proof — that studies the Riemann Hypothesis without analyzing zeros directly. Instead, it works through the **value distribution** of the zeta function: how large does zeta get on the critical line, and how do those sizes grow with the averaging window? The genuine contribution is an algebraic mechanism — the **Superquadratic Growth Theorem (SGT)** — showing that the k^2-rate growth of the zeta moments forces a specific determinant structure. This algebraic/combinatorial core, together with the numerics that spot-check it, is what this paper actually establishes. The proposed analytic bridge from moment growth to RH is **not** established: as we make explicit in the Limitations section below, the pivotal "Hankel positivity ⇒ RH" step carries no arithmetic information (the relevant positivity is automatic), and two supporting steps (a Carleman determinacy claim and a "forced constants" theorem) do not close. We therefore present the paper honestly as a **conditional-and-partial** investigation whose value lies in the algebraic machinery and the reframing it suggests, not in a resolution of RH.

The strongest part of the paper is the **algebraic core**: the Superquadratic Growth Theorem and its machine-verified combinatorial foundation (the rearrangement gap, Lemma 1) are unconditional and complete. A Lean 4 formalization (a computer proof assistant) type-checks the *logical skeleton* of a moments-to-RH chain with **zero sorry** (unfinished proof obligations), but this skeleton **rests on 8 axioms**, and one of them — hankel_positive_implies_rh — has a hypothesis that is itself a theorem (Hankel positivity of the zeta moments is unconditionally true), so that axiom is **logically equivalent to assuming RH**. "Zero sorry" therefore does **not** mean RH is proved; it means the axioms have not been discharged. The 8 axioms are: 5 encoding standard analytical infrastructure (polygamma values, zeta positivity) and 3 encoding proof-chain bridges (QPD→MH, MH→Hankel, Hankel→RH). Each is named and documented in §9.3.

A second, independent route to the Riemann Hypothesis — the **Grade-Shadow route** — is developed in §7.5. It uses the observation that the Euler product is a *grade-2 system* (pairwise prime interactions dominate three-body and higher) with complex symmetry. The Grade-Shadow Correspondence, proved in a companion paper (16 theorems, 0 novel axioms), shows that finite-dimensional systems with grade-2 dominance and Dyson index β have β-ensemble local statistics. Extending this from finite to infinite Euler products requires five sub-axioms (P1–P5), each using

standard tools (Erdős–Yau universality, Prokhorov theorem, Selberg CLT) in a configuration specific to the Euler product. The novelty is in the application, not the tools. Unlike the Hankel route, this bypasses the shifted divisor problem in the analytic sense: it only needs the *ratio* of higher-grade to grade-2 contributions to vanish — which is unconditional from Mertens' theorem.

A further construction — the **Forced CFKRS Theorem** (Theorem 12) — was originally advertised as *showing* that the moment constants are forced to the CFKRS values $c_k = G(1+k)^2/G(1+2k) \cdot \prod_p (1-1/p)^{k^2}\, {}_2F_1(k,k;1;1/p)$ by arithmetic plus Carlson's theorem. **This claim does not hold as stated** (see Limitations, K3): the Carlson step requires knowing that $g(k) = c_k/a(k)$ agrees with the Barnes-G candidate at *all* integers, which is exactly the conclusion; only $k = 0, 1, 2$ are independently known, and Carlson's theorem provides no extrapolation from finitely many points. We therefore restate Theorem 12 as a **conditional** statement: *if* the residual g has exponential type $< \pi$ and is bounded on the imaginary axis, *then* its values are pinned. Establishing that analytic structure — not merely knowing three moments — is the missing input, and it is open.

Abstract

We present an algebraic framework relating the growth of the zeta moments to Hankel-determinant structure, together with an honest account of where the framework does and does not reach the Riemann Hypothesis. The proposed (and, as we explain, **incomplete**) chain is:

$$\text{MH} \xrightarrow{\text{SGT}} H_n > 0 \xrightarrow[\textbf{not established}]{\text{moment determinacy + GUE}} \text{RH}$$

where **MH** (Moment Hypothesis) is the bound $m_{2k}(T) \le C_k (\log T)^{k^2+\varepsilon}$, and $H_n > 0$ is the positivity of all Hankel determinants of the zeta moment sequence. The final arrow does **not** close: for each fixed T the sequence $\{m_{2k}(T)\}_k$ is the moment sequence of the law of $|\zeta(1/2+it)|^2$, so its Hankel matrices are Gram matrices and $H_n(T) > 0$ holds **unconditionally**. The positivity node therefore carries no arithmetic information toward RH, and the terminal step (equivalently, the Lean axiom hankel_positive_implies_rh) is **logically equivalent to RH itself**. We do not claim a proof — nor a genuine reduction — of RH.

The genuine contribution is the **Superquadratic Growth Theorem (SGT)**: k^2-rate moment growth forces the Hankel-determinant structure via the rearrangement inequality, with a gap of at least 1 between the identity-permutation exponent and all others. This combinatorial core is machine-verified in Lean 4 with zero axioms and zero sorry, and it is independently reproduced numerically.

The **Forced CFKRS Theorem** (Theorem 12) is restated here as a *conditional* result: if the residual $g(k) = c_k/a(k)$ extends to a function of exponential type $< \pi$ that is bounded on the imaginary axis, then Carlson's theorem pins its values. The paper's earlier "forced by arithmetic with zero degrees of freedom" phrasing was an overclaim — Carlson cannot extrapolate the required identity from the three known moments $k = 0, 1, 2$ (see Limitations).

A second, independent route — the **Grade-Shadow route** — bypasses the shifted divisor problem in the analytic sense. The Euler product is a grade-2 system (pairwise interactions dominate) with $\beta = 2$ (complex symmetry). The grade ratio $\delta = O(1/\log\log T) \to 0$ (unconditional, from Mertens' theorem). The extension from finite to infinite Euler products decomposes into five sub-axioms

(P1–P5), each using standard tools (Erdős–Yau universality, Prokhorov theorem, Selberg CLT) in a novel configuration adapted to the Euler product's non-uniform variance structure. The most novel step (P1) adapts generalized Wigner matrix universality to the EP random matrix model. This route reproduces the Rudnick–Sarnak restricted support result (1996) as a corollary and explains **why** their Fourier support restriction exists. We caution that this route, too, terminates in a GUE $\Rightarrow$ RH step of the same (RH-equivalent) kind, so it bypasses the shifted divisor problem but **not** the actual obstruction; its completeness is deferred to a companion paper and is not asserted here.

The logical architecture is type-checked in Lean 4 (8 axioms, 0 sorry) and in the proof kernel. We stress that "0 sorry" refers to the logical skeleton *given* the 8 axioms; one axiom is RH-equivalent (above), so the formalization is not a machine proof of RH. Type-checked is not the same as proved: of the Grade-Shadow declarations, only a minority are fully proved, the rest being type-checked statements (see §7.5 and §9 for the exact status and a reconciliation of the declaration counts).

Limitations and Open Steps

This paper does **not** prove the Riemann Hypothesis, conditionally or unconditionally, through the positivity route. Three specific steps that earlier drafts treated as established do not close. We state them plainly here so that no reader mistakes the framework for a proof.

K1 — "Hankel positivity $\Rightarrow$ RH" carries no information. For each fixed T, the sequence $\{m_{2k}(T)\}_{k\geq 0}$ is *literally* the moment sequence of the probability measure μ_T = law of $X = |\zeta(1/2+it)|^2$ for t uniform on $[0,T]$, since $m_{2k}(T) = \int x^k \, d\mu_T(x)$. The Hankel matrix $[m_{2(i+j)}(T)]_{i,j}$ is then a Gram matrix and is positive semidefinite for every T, and positive definite whenever μ_T has infinite support — which it always does. Hence $H_n(T) > 0$ holds **unconditionally**, with no appeal to the Moment Hypothesis, Lindelöf, Ramachandra, or RH. The positivity node is therefore not an intermediate that RH must earn: it is automatic, and the terminal axiom hankel_positive_implies_rh (§9) has a hypothesis that is a theorem, so as a logical statement it is equivalent to RH itself. **Open:** an RH-bearing step must live somewhere that is *not* automatically satisfied; this paper does not supply one.

K2 — The Carleman step is vacuous for $\{c_k\}$. Carleman's theorem presupposes that a sequence *is* the moment sequence of some positive measure and then asserts uniqueness of that measure. The normalized CFKRS constants $\{c_k\}$ are **not** such a sequence: already $c_0c_2 - c_1^2 = 1 \cdot \frac{1}{2\pi^2} - 1^2 = -0.949 < 0$, so the 2×2 Hankel minor is negative and no measure on $\mathbb{R}$ has moments $\{c_k\}$ (the paper's own Appendix A.2 records this negativity). Thus "$\{c_k\}$ determines a unique positive measure by Carleman" is vacuous. The positivity that *is* available applies to the unnormalized $\{m_{2k}(T)\} = \{c_k L^{k^2}(1+\cdots)\}$ and comes entirely from the L^{k^2} amplification, not from $\{c_k\}$. **Open:** there is no correct determinacy statement here that reaches GUE.

K3 — Theorem 12 (Forced CFKRS) does not force the constants. The argument applies Carlson's theorem to $\hat{g}-h$, which requires that $\hat{g}(k) = h(k)$ for *all* $k \geq 0$ — i.e. $g(k) = h(k)\ \forall k$, which is exactly the desired conclusion. Only $k = 0, 1, 2$ are independently verified (Hardy–Littlewood, Ingham). Carlson's theorem converts agreement at *all* integers into a global identity; it provides **no** extrapolation from three points. If the argument were valid it would settle the (open) CFKRS value conjecture for $k \geq 3$ from three known moments, which is a tell that it is not. **Open:** the required type-$< \pi$ / imaginary-axis-bounded analytic structure of g is the genuine missing input and is not known unconditionally.

What is genuinely established (and is the paper's contribution): the algebraic core — Lemma 1 (rearrangement gap), the SGT factorization (Theorems 2–3), and the exponent identity $E_n = \frac{2n(n+1)(2n+1)}{3}$ — all machine-verified; and the spot-checkable numerics (Theorems 18 and 21), which reproduce independently. These stand on their own, independently of the RH question.

---

1. Introduction

1.1 The Chain

The classical approach to the Riemann Hypothesis analyzes the zeros of $\zeta(s)$ directly. We explore a different route through the **value distribution** of ζ on the critical line. The route one would *like* to have is

$$\text{Euler product structure} \implies \text{Moment bounds} \implies H_n > 0 \implies \text{RH},$$

but, as the Limitations section explains, the last arrow does not close: $H_n(T) > 0$ is automatic (a Gram-matrix fact), so it is not an RH-bearing intermediate. What follows should therefore be read as a framework and a set of algebraic results, not as a chain that reaches RH.

The key observation: the Riemann zeta function is not an arbitrary analytic function — it has a **multiplicative structure** given by the Euler product

$$\zeta(s) = \prod_p (1 - p^{-s})^{-1}, \quad \mathrm{Re}(s) > 1.$$

This multiplicative structure imposes strong regularity on the distribution of $|\zeta(1/2 + it)|$. We make the *algebraic* half precise by proving that the k^2 growth rate of moments forces a definite Hankel-determinant structure (the SGT). We do **not** establish the step from this structure to zero correlations: that would require an RH-bearing input beyond the automatic positivity, which we do not have.

1.2 Main Result

> **Theorem (Main, conditional — with an unclosed final bridge).** *Assume the Moment Hypothesis (MH): for all $k \geq 1$ and every $\varepsilon > 0$, $m_{2k}(T) \ll_k (\log T)^{k^2+\varepsilon}$. Then all Hankel determinants of the zeta moment sequence are positive (this in fact holds unconditionally). The further implication to the Riemann Hypothesis is **not established** here.*
>
> *The intended chain is: MH $\xrightarrow{SGT}$ $H_n > 0$ $\xrightarrow{\text{moment determinacy + GUE}}$ RH. Only the algebraic core (SGT) is machine-verified in Lean 4. The final bridge (Hankel positivity → RH via GUE universality) is **axiomatized**, and because its hypothesis $H_n > 0$ is a theorem, that axiom is logically equivalent to RH — it assumes what it would prove (see Limitations, K1).*

MH is unconditionally known for $k \leq 2$ (Hardy–Littlewood, Ingham). MH for all k follows from the Lindelöf hypothesis and is strictly weaker.

Update (April 2026). The companion working papers *The Fourier–Euler Product and the Moment Hypothesis for the Riemann Zeta Function* (Nagy 2026, Zenodo DOI 10.5281/zenodo.19143800) and *The Riemann Hypothesis as a Latent Existence Theorem* (Nagy 2026) record two independent machine-verified **algebraic** routes toward the implication "bounded cumulants (from Euler product structure) ⇒ MH for all k": (i) the Latent bridge — CGF analyticity gives cumulant bounds simultaneously via Cauchy estimates (6 theorems); (ii) traditional induction via the Leonov–Shiryaev recursion with explicit constants $C_3 = 6$, $C_4 = 26$, $C_5 = 150$ (16 theorems). **Community validation of the analytic inputs packaged in those formalizations is separate** (see §9–§10.3).

Update (April 2, 2026 — Shifted Divisor Domain). A new proof domain (57 theorems, 4 files) provides a **proposed resolution** of the shifted divisor problem via a log-decomposition argument: $\log|\zeta|^2 = \log|P|^2 + X$ where $X = \log|1 + R/P|^2$. The functional equation forces $|R/P| \leq 1$, so X is a bounded random variable with an entire MGF and bounded cumulants of all orders. By Kronecker-Weyl equidistribution, X is asymptotically independent of $\log|P|^2$, giving $\kappa_m(\zeta) = \kappa_m(P) + \kappa_m(X) + o(1)$, both bounded for $m \geq 3$. Combined with Selberg's CLT for κ_2, this yields MH for all k. Total: 89 algebraic theorems, 0 novel axioms, 5 axiomatized classical inputs. The algebraic infrastructure for the entire MH $\implies$ RH chain is now complete.

1.3 Why This Is New

The traditional implication is:

$$\text{RH} \implies \text{moment bounds} \implies \text{distributional regularity.}$$

The reversal one would like to carry out is:

$$\text{Euler product} \implies \text{moment regularity} \implies H_n > 0 \implies \text{zero correlations determined} \implies \text{RH.}$$

The first two implications are genuine (the SGT is the mechanism). The reversal **fails at the third arrow**: $H_n > 0$ is unconditionally true and so does not, by itself, determine the zero correlations (Limitations, K1). We keep the diagram to show the intended architecture, not as an established derivation.

1.4 Comparison with Prior Approaches

Approach	Input required	Method	Status
Direct zero analysis (de la Vallée-Poussin, 1896)	None (unconditional)	Explicit zero-free region $\sigma > 1 - c/\log t$	Best known: Vinogradov–Korobov
Lindelöf hypothesis → RH	$\lvert\zeta(1/2+it)\rvert \leq t^{\varepsilon}$	Growth bounds on critical line	Open; implies MH hence this paper's chain
Selberg CLT approach	Zero statistics	Central limit theorem for $\log\lvert\zeta\rvert$	Proved for $\log\lvert\zeta\rvert$; extending to RH open

Approach	Input required	Method	Status
Random matrix theory (Montgomery–Odlyzko)	RH (circular)	Pair correlation matches GUE	Conditional on RH; strong numerical support
Langlands functoriality	Automorphic forms	Transfer of zero-free regions	Partial (GL(2) complete, GL(n) open)
This paper	**MH for all k (= moment upper bounds)**	**Euler product $\to k^2$ growth $\to$ SGT $\to$ Hankel $\to$ RH**	**Conditional on MH($k \geq 3$); algebraic core machine-verified. Companion formalizations add many further algebraic lemmas (counts vary by route; see §1.2 updates vs. §9).**

The key difference: this approach works through the **value distribution** rather than the zeros, and reduces RH to a concrete arithmetic input (moment upper bounds) that is incrementally improvable.

1.5 Notation

Throughout, $m_{2k}(T) = \frac{1}{T}\int_0^T |\zeta(1/2+it)|^{2k}\,dt$ denotes the $2k$-th moment. $L = \log T$. The Hankel determinant of order n is

$$H_n(T) = \det\big[m_{2(i+j)}(T)\big]_{0\leq i,j\leq n}.$$

We write $H_n(L)$ with $L = \log T$ when the SGT scaling is the natural variable (§3), and bare H_n in generic statements where the argument is clear from context.

Overloaded symbols (contextual). The letter D appears in three standard but distinct roles: $D = \mathrm{diag}(L^{0^2}, \ldots, L^{n^2})$ (diagonal scaling matrix, §3); $D_{2k}(T)$ (diagonal moment sum, §5.1); $D(t)$ (Dirichlet polynomial in the approximate functional equation, §8). Each is defined at first use; context and subscripts disambiguate. Similarly, σ denotes a permutation in S_{n+1} in §3 and $\mathrm{Re}(s)$ elsewhere (see note in §3). The letter k always denotes the moment order; sums $\sum_k$ range over moment orders. The symbol a_k in §5.1 (Theorem 5) denotes the diagonal coefficient $G_k(1)/\Gamma(k^2)$, distinct from $a(k)$ which denotes the CFKRS arithmetic factor $\prod_p (1-1/p)^{k^2}\,{}_2F_1(k,k;1;1/p)$ introduced in §2.1.

2. Background

2.1 Zeta Moments: What Is Known

The moments $m_{2k}(T)$ are among the most studied quantities in analytic number theory. Unconditional results:

k	Asymptotic	Reference
1	$m_2(T) = \log T + (2\gamma - 1) + O(T^{-1/2})$	Hardy–Littlewood, 1918
2	$m_4(T) = \frac{1}{2\pi^2}(\log T)^4 + O((\log T)^3)$	Ingham, 1926
$k \geq 3$	$m_{2k}(T) \geq c'_k(\log T)^{k^2}$ (lower bound)	Ramachandra, 1980

The CFKRS Conjecture (Conrey–Farmer–Keating–Rubinstein–Snaith, 2005):

$$m_{2k}(T) \sim c_k \, (\log T)^{k^2}, \quad c_k = a(k) \cdot g(k)$$

where $a(k) = \prod_p (1-1/p)^{k^2} {}_2F_1(k, k; 1; 1/p) > 0$ is the arithmetic factor (convergent Euler product) and $g(k) = G(k+1)^2/G(2k+1)$ is the random matrix factor (Barnes G-function).

2.2 Random Multiplicative Functions

A **Steinhaus random multiplicative function** f is defined by $f(p) = e^{i\theta_p}$ (i.i.d. uniform on $[0, 2\pi)$) extended multiplicatively. The partial Euler product:

$$F_y(s) = \prod_{p \leq y} (1 - f(p)\, p^{-s})^{-1}.$$

Theorem (Harper — random multiplicative / partial Euler models). For Steinhaus-type models, sharp $\mathbb{E}$-moment bounds of the form

$$\mathbb{E}\left[\frac{1}{T} \int_0^T |F_y(1/2 + it)|^{2k} \, dt \right] \leq C_k \, (\log T)^{k^2}$$

hold for all $k \geq 1$ (see Harper's series on random multiplicative functions, e.g. arXiv:1305.4618, arXiv:1703.06654, and the 2024 survey arXiv:2410.11523 for context).

Theorem (Gorodetsky–Wong, 2025). The random measure $|F_y(1/2 + it)|^2 \, dt$ converges to a critical Gaussian multiplicative chaos (GMC) measure (see Saksman–Webb, 2020, for the foundational connection between ζ and GMC on the critical line).

2.3 Hankel Determinants and the Moment Problem

Given a sequence $\{\nu_k\}_{k \geq 0}$ with $\nu_0 = 1$, the **Stieltjes Hankel determinant** of order n is

$$H_n = \det[\nu_{i+j}]_{i,j=0}^{n}.$$

Moment Determinacy Theorem (Stieltjes). If $H_n > 0$ for all $n \geq 0$ and the moments satisfy the **Carleman condition** $\sum_{k=1}^{\infty} \nu_{2k}^{-1/(2k)} = \infty$, then the moment sequence $\{\nu_k\}$ determines a unique

probability measure μ on $[0,\infty)$. (This is the Stieltjes setting; the Hamburger analogue on $\mathbb{R}$ is also classical.)

Key fact (fixed-T determinacy). For each fixed T, the measure μ_T is supported on $[0, \max_{t\in[0,T]} |\zeta(1/2+it)|^2]$, which is compact. Compactly supported measures have **automatically determinate** moment problems (Akhiezer, 1965, Thm 2.3.3) — no Carleman condition is needed.

Correction (the normalized sequence $\{c_k\}$ is NOT a moment sequence). Earlier drafts asserted a "normalized-moment determinacy": that because the CFKRS constants $c_k = g(k)a(k)$ decay super-exponentially ($\log c_k \sim -k^2 \log k$, so $c_k^{-1/(2k)} \sim k^{k/2} \to \infty$ and $\sum_k c_k^{-1/(2k)}$ diverges), Carleman's theorem makes $\{c_k\}$ "determinate." **This application is invalid.** Carleman's theorem presupposes that $\{c_k\}$ *is* the moment sequence of a positive measure (equivalently, that its Hankel forms are positive semidefinite) and only then asserts uniqueness. But $\{c_k\}$ is not positive-definite: already

$$c_0 c_2 - c_1^2 = 1 \cdot \tfrac{1}{2\pi^2} - 1^2 = 0.0507 - 1 = -0.949 < 0,$$

so its 2×2 Hankel minor is negative and **no** measure on $\mathbb{R}$ has moments $\{c_k\}$ (see Appendix A.2, which records the same negativity). A divergent Carleman sum is irrelevant when the underlying positive-definiteness hypothesis fails. Consequently there is nothing for Carleman to be "determinate" about, and the normalized sequence cannot be used to produce a limiting measure. (Note also the standard Carleman condition is stated as $\sum_k \nu_{2k}^{-1/(2k)} = \infty$ in terms of the even-index moments; the index in earlier drafts, $\sum_k c_k^{-1/(2k)}$, was written incorrectly, though both diverge here.)

What survives. Determinacy for each *fixed* T is genuine and comes from compact support: μ_T is supported on $[0, \max_t |\zeta(1/2+it)|^2]$, so its moment problem is automatically determinate (Akhiezer, 1965). But this fixed-T determinacy conveys no information toward RH — it holds for any continuous function on a compact interval. The would-be determinacy of the *limiting normalized* sequence $\{c_k\}$, which the RH bridge relied on, does not exist.

2.4 GUE Universality for Zeta Zeros

Montgomery's Pair Correlation Conjecture (1973). The pair correlation of non-trivial zeros of ζ matches the GUE sine-kernel:

$$R_2(\alpha) = 1 - \left(\frac{\sin \pi\alpha}{\pi\alpha}\right)^2.$$

Montgomery proved this for restricted test functions (conditional on RH). Odlyzko (1987) confirmed it numerically to high precision. Hejhal (1994) extended to higher correlations.

Post-2025 landscape. Jiang (2025, arXiv:2507.20653, *On Hypothesis H of Rudnick and Sarnak*) proves **Hypothesis H** in full generality for GL_n over number fields (see the paper for applications to automorphic L-functions). This removes Hypothesis H as a **pending analytic input** in parts of the Rudnick–Sarnak program; it does **not** by itself remove the **Fourier support restriction** on test functions in Rudnick–Sarnak (1996), nor substitute for the Hankel→RH bridge in §4. Baluyot–Goldston–Suriajaya–Turnage-Butterbaugh (2024) proved unconditionally that at least 61.7% of zeros near the critical line are simple. For $\zeta(s)$, zero-correlation evidence remains **conditional** on the same kinds of analytic inputs emphasized throughout this paper.

The Fourier support barrier. The Rudnick–Sarnak result requires test functions whose Fourier transforms are supported in a restricted region. For our framework, the relevant test function has $\widehat{f}(\xi) \sim e^{-\pi|\xi|}$ (exponential decay, not compact support). The tail beyond the Rudnick–Sarnak boundary contributes an error that is exponentially small but nonzero.

The connection: if the Stieltjes moment problem for $|\zeta(1/2+it)|^2$ has a unique solution with the correct growth rate, then the resulting measure determines all multilinear statistics of zeta zeros, which must match GUE. GUE universality then forces all zeros to lie on $\mathrm{Re}(s) = 1/2$.

Quantitative truncated approach (Bickel–Pascoe–Sargent 2023). Truncated Hankel positivity conditions ($H_n > 0$ for $n \leq N$) yield a family of relaxations of RH. Numerical computation shows that the positivity threshold scales as $L_0(n) \approx 13n^2$ (where $L = \log T$), giving an **incremental program**: each new order of Hankel positivity strengthens the zero-free constraint.

---

3. The Superquadratic Growth Theorem

Notation. In this section, σ denotes a permutation in the symmetric group S_{n+1}. Elsewhere in the paper, σ denotes $\mathrm{Re}(s)$ for $s \in \mathbb{C}$ (standard in analytic number theory). The two uses do not overlap.

This is the algebraic core of the proof.

3.1 The Rearrangement Gap

Lemma 1 (Rearrangement Gap). *For any permutation $\sigma \in S_{n+1}$ with $\sigma \neq \mathrm{id}$:*

$$\mathrm{gap}(\sigma) := \sum_{i=0}^{n} i^2 - \sum_{i=0}^{n} i\,\sigma(i) \geq 1.$$

Proof. By the rearrangement inequality, $\sum_{i=0}^{n} a_i b_i$ over sequences (a_i), (b_i) is maximized when both are sorted in the same order. Since $i \mapsto i$ is already sorted, $\sum i\,\sigma(i)$ is maximized uniquely by $\sigma = \mathrm{id}$. For any $\sigma \neq \mathrm{id}$, there exist $a < b$ with $\sigma(a) > \sigma(b)$, and swapping gives a strict increase: $(b-a)(\sigma(a) - \sigma(b)) > 0$ implies $a\sigma(b) + b\sigma(a) > a\sigma(a) + b\sigma(b)$. By induction on the number of inversions (bubble sort), any σ can be improved to id through transpositions, each strictly increasing $\sum i\,\sigma(i)$. Since all quantities are integers, each transposition increases the sum by at least 1, giving $\mathrm{gap}(\sigma) \geq 1$. □

Machine verification. The Lean 4 formalization (SGTProof.lean) proves the gap bound for **all** n (the statement is universally quantified). The key computational lemmas that discharge specific cases use native_decide for $n \leq 4$; the general case follows from the classical rearrangement inequality argument formalized above. Zero axioms and zero sorry. The Lean statement:

```
theorem rearrangement_gap (n : ) ( : Equiv.Perm (Fin (n + 1)))
    ( h :     1) : 1   gap n
```

Scope note. The theorem is stated and type-checks for all n. The proof strategy uses the swap argument (if $\sigma \neq \mathrm{id}$, there exist $a < b$ with $\sigma(a) > \sigma(b)$; swapping increases $\sum i\sigma(i)$ by a positive integer). The native_decide annotations provide fast verification for small n but are not required for the general result.

3.2 The Main Algebraic Theorem

Theorem 2 (Superquadratic Growth Theorem). *Let $\{\nu_k\}_{k=0}^{\infty}$ with $\nu_0 = 1$ and*

$$\nu_k = c_k \cdot L^{k^2} \cdot (1 + \varepsilon_k(L))$$

where $c_k > 0$, $L > 0$, and $|\varepsilon_k(L)| \leq C(n)/L^{\alpha}$ for some $\alpha > 0$ (uniformly for $k \leq 2n$). Then:

(a) $H_n(L) > 0$ for all $L > L_0(n)$.

(b) The leading order:

$$H_n(L) = P_n \cdot L^{E_n} \cdot (1 + O_n(L^{-2}))$$

where $E_n = \frac{2n(n+1)(2n+1)}{3}$ and $P_n = \prod_{i=0}^{n} c_{2i} > 0$.

Proof. Since $(i+j)^2 = i^2 + 2ij + j^2$, the Hankel matrix factorizes:

$$M_{ij} = L^{i^2} \cdot \underbrace{c_{i+j}\, L^{2ij}(1 + \varepsilon_{i+j})}_{\tilde{A}_{ij}} \cdot L^{j^2}$$

giving $M = D\tilde{A}D$ with $D = \mathrm{diag}(L^{0^2}, L^{1^2}, \ldots, L^{n^2})$.

By the Leibniz formula:

$$\det(\tilde{A}) = \sum_{\sigma \in S_{n+1}} \mathrm{sgn}(\sigma) \prod_{i=0}^{n} c_{i+\sigma(i)}\, L^{2i\sigma(i)}(1 + \varepsilon_{i+\sigma(i)}).$$

The exponent of L in the σ-term is $2\sum_i i\sigma(i)$. By Lemma 1, the identity contributes the unique leading term at $L^{2\sum i^2}$, while every other permutation contributes at most $L^{2\sum i^2 - 2}$. Since there are $(n+1)!$ permutations:

$$\det(\tilde{A}) = \prod_{i=0}^{n} c_{2i} \cdot L^{2\sum i^2} \cdot (1 + O_n(L^{-2})).$$

Combining: $H_n = L^{2\sum i^2} \cdot \det(\tilde{A}) = P_n \cdot L^{4\sum i^2} \cdot (1 + O_n(L^{-2}))$. □

3.3 Generalized Version (Bounds Only)

Theorem 3 (Generalized SGT). *Let $\{\nu_k\}$ satisfy:*

$$c'_k\, L^{k^2} \leq \nu_k \leq C_k\, L^{k^2 + \varepsilon}$$

with $c'_k > 0$ and $C_k < \infty$. Then for every $n \geq 0$: $H_n(L) > 0$ for $L > L_0(n, c', C)$.

Proof. The identity permutation contributes at least $\prod c'_{2i}\, L^{2\sum i^2}$. Each non-identity permutation contributes at most $\prod C_j \cdot L^{2\sum i^2 - 2 + (n+1)\varepsilon}$. For $\varepsilon < 2/(n+1)$, the exponent $-2 + (n+1)\varepsilon < 0$, so the identity dominates for $L > L_0$. □

---

4. The Moment Hypothesis and the Main Theorem

4.1 The Moment Hypothesis

Definition (Moment Hypothesis, MH). We say MH holds if for all $k \geq 1$ and every $\varepsilon > 0$:

$$m_{2k}(T) \leq C_{k,\varepsilon} \, (\log T)^{k^2+\varepsilon}$$

for some $C_{k,\varepsilon} > 0$ (equivalently, $m_{2k}(T) \ll_k (\log T)^{k^2+\varepsilon}$ for every $\varepsilon > 0$).

Status. MH(1) and MH(2) are unconditionally proved (Hardy–Littlewood, Ingham). MH(all k) follows from the Lindelöf hypothesis and is strictly weaker.

4.2 The Ramachandra Lower Bound

Theorem (Ramachandra, 1980; Arguin–Creighton, 2026). Unconditionally, for all $k \geq 1$:

$$m_{2k}(T) \geq c'_k \, (\log T)^{k^2}$$

with $c'_k > 0$.

4.3 MH Implies RH

Theorem 4 (MH implies Hankel positivity; the further step to RH is not established). *If MH holds for all k, then all Hankel determinants of the zeta moment sequence are positive. This conclusion in fact holds unconditionally; the implication from it to RH is the unclosed bridge discussed below and in the Limitations section.*

The argument proceeds in four steps. **The reader should note at the outset that Step 1's conclusion is automatic** (see Step 1a), so the substance of any RH claim would have to come from Steps 2–4, which do not close.

Step 1: Hankel positivity (SGT).

Combining the Ramachandra lower bound (§4.2) with MH:

$$c'_k \, (\log T)^{k^2} \leq m_{2k}(T) \leq C_k \, (\log T)^{k^2+\varepsilon}.$$

Set $\nu_k = m_{2k}(T)$ and $L = \log T$. The sequence $\{\nu_k\}$ satisfies the hypotheses of Theorem 3 (Generalized SGT) with the lower bounds $c'_k > 0$ from Ramachandra and the upper bounds $C_k L^{k^2+\varepsilon}$ from MH. Therefore:

$$H_n(T) > 0 \quad \text{for all } n \geq 0 \text{ and } T > T_0(n).$$

Step 1a (crucial caveat — this positivity is automatic and carries no arithmetic information). The SGT is a correct sufficient condition, but for the *actual* zeta moments the conclusion needs neither MH nor Ramachandra: for each fixed T, $\{m_{2k}(T)\}_k$ is the moment sequence of \$ _T = \$ law of $|\zeta(1/2+it)|^2$, so $\sum_{i,j} \xi_i \xi_j \, m_{2(i+j)}(T) = \int (\sum_i \xi_i x^i)^2 \, d\mu_T \geq 0$, i.e. every Hankel matrix is a Gram matrix and $H_n(T) > 0$ **unconditionally** (strictly, since μ_T has infinite support). Hence MH does no work at this step, and $H_n > 0$ cannot serve as an RH-bearing intermediate. Any route to RH must extract information from a step that is *not* automatically satisfied — which the steps

below do not achieve. We retain the SGT because it is a genuine algebraic theorem and the engine of the framework, not because it advances toward RH here.

Step 2: Determinacy of the moment problem.

For each fixed T, the empirical measure $\mu_T = \frac{1}{T}\,\mathrm{Leb}_{[0,T]} \circ (|\zeta(1/2+\cdot)|^2)^{-1}$ is a **compactly supported** probability measure on $[0, M_T]$ where $M_T = \max_{t\in[0,T]} |\zeta(1/2+it)|^2 < \infty$ (since ζ is continuous on compact intervals). For compactly supported measures, the moment problem is **automatically determinate** — no Carleman condition is needed (Akhiezer, 1965, Thm 2.3.3). The moments of μ_T uniquely determine μ_T.

Passage to the limit. The Hankel positivity $H_n(T) > 0$ ensures that each μ_T is a valid probability measure with positive-definite moment sequence. Define the renormalized moments $\tilde{\nu}_k(T) = m_{2k}(T)/L^{k^2}$. By Ramachandra, $\tilde{\nu}_k(T) \geq c'_k > 0$ for all T. By MH, $\tilde{\nu}_k(T) \leq C_k L^{\varepsilon}$ for every $\varepsilon > 0$, so $\tilde{\nu}_k$ grows at most sub-polynomially in L. In particular, $\{\tilde{\nu}_k(T)\}$ is bounded on compact T-intervals, and the lower bound prevents collapse.

What this gives. The Ramachandra–MH sandwich shows that any subsequential limit $c_k = \lim_{j\to\infty} \tilde{\nu}_k(T_j)$ (if it exists) satisfies $c_k \geq c'_k > 0$. But MH alone does not guarantee that the limit exists or is unique — the renormalized moments could oscillate within the growing band $[c'_k, C_k L^{\varepsilon}]$.

Remark (what the limit requires). To conclude that μ_T converges to a unique limit μ_∞, one needs moment convergence: $m_{2k}(T)/L^{k^2} \to c_k$ for explicit constants c_k. Under MH together with the Ramachandra lower bounds, the moments are sandwiched between $c'_k L^{k^2}$ and $C_k L^{k^2+\varepsilon}$. This gives convergence along subsequences but not uniqueness of the limit without the CFKRS-level input that the constants c_k exist. The axiom hankel_positive_implies_rh in the Lean formalization encodes this bridge, and it is the deepest analytical step in the chain.

Step 3: From determined moments to zero statistics.

The moment sequence $\{m_{2k}(T)\}$, once determined by the unique measure μ_T (Step 2), constrains the statistical behavior of $|\zeta(1/2+it)|^2$. The connection to zeros passes through the **explicit formula**.

For a test function f supported on $[\alpha, \beta]$:

$$\sum_{\gamma} f(\gamma) = \frac{T}{2\pi}\int f(t)\,dt - \frac{1}{2\pi}\int f(t)\ \mathrm{Re}\,\frac{\zeta'}{\zeta}(1/2+it)\,dt + O(\log T)$$

where the sum is over imaginary parts γ of non-trivial zeros.

The explicit formula is a **linear relation** between zero statistics and value statistics — not a bijection in general. However, knowing $m_{2k}(T)$ for all k constrains all polynomial statistics of $|\zeta|^2$, which in turn constrains the distribution of $\log|\zeta|^2$ (via the moment-generating function, when it converges). The logarithmic derivative $\zeta'/\zeta = (\log\zeta)'$ connects value statistics to zero statistics through the explicit formula machinery of Rudnick–Sarnak (1996, §3).

What this step establishes and what it does not. The moment tower $\{m_{2k}\}_{k\geq 1}$ determines the one-point distribution of $|\zeta(1/2+it)|^2$ and, via the Euler product structure, constrains the multilinear statistics. The step from constrained value statistics to fully determined n-point zero correlations requires additional input — specifically, the Euler product factorization structure that links moments at different orders. This connection is part of the analytical content encoded in the axiom hankel_positive_implies_rh.