

A Theta-Geometric Framework for the Hankel Positivity Problem in the Riemann Hypothesis

Spectral Folding, Stieltjes Moments, and Exact Theta Differential Kernels

Agum Abdillah

Independent Researcher
agumabdillah97@gmail.com

Abstract

We develop a theta-geometric framework for the Hankel positivity problem underlying the Riemann Hypothesis. Starting from the even centered Riemann xi-function, we fold the spectral symmetry, construct a meromorphic resolvent whose poles are the squared centered zeros, and pass to reciprocal spectral variables. This produces a canonical moment sequence whose Stieltjes moment property is equivalent to positivity of the two associated Hankel families.

We then investigate how the classical theta representation interacts with this exact spectral criterion. The centered xi-function is represented as a positive cosh transform of an explicit theta density, producing a parameter-dependent tilted measure and an exact expectation formula for the resolvent. We derive the corresponding differential hierarchy, an adjoint operator, and a two-variable kernel for the first resolvent derivative. After rescaling, the theta operator yields an exact inhomogeneous two-variable differential identity, together with boundary terms that cannot be discarded.

The paper therefore isolates the decisive remaining theorem: positivity of the complete Hankel hierarchy, for every matrix size, in a representation compatible with the fixed reciprocal spectral moment sequence. We separate exact identities and equivalences from numerical checks and do not claim a proof of the Riemann Hypothesis. The purpose of the framework is to reduce the final question to a precise, auditable positivity problem rather than to replace that missing step by a finite-order or parameter-dependent argument.

Keywords: Riemann Hypothesis; Riemann xi-function; spectral folding; reciprocal zeros; Stieltjes moments; Hankel matrices; theta kernels; Loewner kernels.

MSC 2020: 11M26; 30D35; 44A60; 15B48.

1 Introduction

The central aim of this paper is to reorganize the spectral information carried by the nontrivial zeros of the Riemann zeta function in a form that makes positivity visible. The reorganization is deliberately elementary at its first stage. We center the completed xi-function at the symmetry point, quotient the even involution, and then pass to reciprocal spectral variables. The resulting resolvent and its moments give a canonical bridge between zero geometry and a Stieltjes moment problem. That bridge is exact. The present paper takes the next step: it asks how the theta representation of the xi-function interacts with this positive-moment geometry.

Let

$$\xi(s) = \frac{1}{2}s(s-1)\pi^{-s/2}\Gamma\left(\frac{s}{2}\right)\zeta(s) \quad (1.1)$$

and define the centered function

$$F(z) = \xi\left(\frac{1}{2} + z\right). \quad (1.2)$$

The functional equation is equivalent to the elementary symmetry

$$F(-z) = F(z). \quad (1.3)$$

The nontrivial zeros in the original variable occur in the familiar symmetry pattern, while the centered variable records the distance from the critical line as its real part. Thus a centered zero is on the critical line precisely when it is purely imaginary.

The first geometric operation is to quotient the involution $z \mapsto -z$. Since F is even, there is an entire function G satisfying $F(z) = G(z^2)$. It is more convenient to use the sign-reversed folded coordinate $w = -z^2$ and to define

$$K(w) = G(-w) = F(i\sqrt{w}). \quad (1.4)$$

The square-root notation is only a mnemonic: evenness makes the resulting function single-valued and entire. If ρ is a centered zero of multiplicity m_ρ , then the corresponding zero of K is at $-\rho^2$ with the same multiplicity. Consequently the critical line, written as $\rho = i\gamma$, is sent to the positive real axis $w = \gamma^2$.

The logarithmic derivative of the folded function is the natural resolvent. We use the convention

$$R(w) = -\frac{K'(w)}{K(w)}, \quad S(x) = R(-x). \quad (1.5)$$

Then

$$S(x) = \sum_{\rho} \frac{m_{\rho}}{x - \rho^2}, \quad (1.6)$$

where one representative is chosen from each pair $\{\rho, -\rho\}$. Passing once more to the reciprocal coordinate

$$q_{\rho} = -\frac{1}{\rho^2} \quad (1.7)$$

turns the critical-line condition into the elementary positivity condition $q_{\rho} > 0$. Expanding $S(-u)$ at the origin therefore produces the reciprocal spectral moments

$$\mu_n = \sum_{\rho} m_{\rho} q_{\rho}^{n+1}. \quad (1.8)$$

The resulting moment problem is not merely analogous to the zero problem. It is exactly equivalent to it in the relevant sense. The two Hankel families

$$H_N = (\mu_{i+j})_{i,j=0}^N, \quad \tilde{H}_N = (\mu_{i+j+1})_{i,j=0}^N \quad (1.9)$$

are positive semidefinite for every N precisely when the reciprocal poles are nonnegative, and comparison with the meromorphic representation of S identifies this condition with the Riemann Hypothesis. This equivalence is the starting point of the present article, not its final result.

The new question is how one can see the same positivity geometry directly from the theta side. The theta representation gives an exact integral formula

$$F(r) = 4 \int_0^{\infty} \Phi(u) \cosh(ru) du, \quad (1.10)$$

where

$$\Phi(u) = \sum_{n \geq 1} \pi n^2 e^{5u/2} (2\pi n^2 e^{2u} - 3) e^{-\pi n^2 e^{2u}}. \quad (1.11)$$

The density Φ is positive on the relevant half-line. The important point, however, is that positivity of Φ is not by itself the desired spectral positivity: the resolvent is a logarithmic derivative of the cosh transform, and differentiation introduces nonlinear interactions between the theta modes.

We therefore develop the paper around that interaction. The first ingredient is a tilted probability measure

$$dP_r(u) = \frac{\Phi(u) \cosh(ru)}{A(r)} du, \quad A(r) = \int_0^\infty \Phi(u) \cosh(ru) du. \quad (1.12)$$

In this measure the resolvent becomes the expectation of the explicit observable $u \tanh(ru)$. Its derivatives can then be written as variance-type expressions. This yields a two-variable kernel whose failure to be pointwise positive is completely explicit.

The theta density itself also has a differential origin. If

$$\Theta(u) = \sum_{n \geq 1} e^{-\pi n^2 e^{2u}}, \quad Y(u) = e^{u/2} \Theta(u), \quad (1.13)$$

then a direct calculation gives

$$\Phi(u) = \frac{1}{2} \left(Y''(u) - \frac{1}{4} Y(u) \right). \quad (1.14)$$

This makes integration by parts natural. The adjoint operator induced by the tilted cosh weight is

$$\mathcal{L}_r f = f'' + 2r \tanh(ru) f' + \left(r^2 - \frac{1}{4} \right) f. \quad (1.15)$$

A two-variable application of this operator to the first-derivative kernel yields an exact inhomogeneous differential identity. This is the main structural observation of the operator analysis.

The operator identity does not finish the positivity problem. Boundary terms remain, and the transformed kernel contains a defect that is not pointwise nonnegative. We do not suppress this fact. Instead we isolate it, express it explicitly, and formulate the all-orders positivity problem in a way that separates exact identities from the missing analytic input. The resulting formulation is intended to be useful even before the final positivity theorem is known.

The organization is as follows. Section 2 fixes the spectral-geometric setup and all conventions. Section 3 derives the folded resolvent and reciprocal variables. Section 4 establishes the Stieltjes and Hankel formulation. Section 5 develops the theta representation and its geometric measure. Section 6 studies the tilted measure and the expectation form of the resolvent. Section 7 derives the differential and kernel identities. Section 8 carries out the operator reduction for the first positivity condition and isolates the remaining boundary problem. Section 9 explains the exact content of the remaining all-orders problem. Section 10 records consequences and a concrete research program. Section 11 concludes. The appendices collect product identities, convergence details, low-order cumulant relations, and algebraic checks used throughout the paper.

2 Spectral-Geometric Setup

We begin by fixing notation carefully, because several later sign choices depend on the same two changes of variables. The completed xi-function will always be denoted by ξ , its centered version by F , the folded entire function by K , the folded logarithmic derivative by R , the reflected

resolvent by S , and the reciprocal generating function by T .

2.1 The centered function

The completion

$$\xi(s) = \frac{1}{2}s(s-1)\pi^{-s/2}\Gamma\left(\frac{s}{2}\right)\zeta(s) \quad (2.1)$$

is entire and satisfies $\xi(s) = \xi(1-s)$. With

$$F(z) = \xi\left(\frac{1}{2} + z\right), \quad (2.2)$$

we obtain

$$F(-z) = F(z). \quad (2.3)$$

Hence F is an entire even function. Its nonzero zeros occur in pairs $\{\rho, -\rho\}$.

The notation ρ in this paper refers to a zero in the centered variable unless the original variable is explicitly displayed. Thus the statement $\rho = i\gamma$ means that the corresponding zero of ζ is $1/2 + i\gamma$. Complex conjugation is not identified with the folding involution; if ρ and $\bar{\rho}$ are distinct representatives, they are retained separately.

Definition 2.1. Let $\mathcal{Z}_c$ be a choice of one representative from each pair $\{\rho, -\rho\}$ of nontrivial centered zeros of F . If m_ρ is the multiplicity of ρ , then m_ρ also denotes the multiplicity attached to the representative in $\mathcal{Z}_c$.

The elementary equivalence

$$\Re \rho = 0 \iff -\rho^2 \in (0, \infty) \iff -\frac{1}{\rho^2} \in (0, \infty) \quad (2.4)$$

will be used repeatedly. It is important that the sign is fixed: the positive coordinate is $-\rho^2$, not ρ^2 .

2.2 Folding the even geometry

Since F is even, there exists an entire function G such that

$$F(z) = G(z^2). \quad (2.5)$$

Define

$$K(w) = G(-w). \quad (2.6)$$

Equivalently, $K(w) = F(i\sqrt{w})$ in the sense explained above. The map $z \mapsto w = -z^2$ identifies z and $-z$ and moves the critical line to the positive real axis.

Proposition 2.2. If $\rho \in \mathcal{Z}_c$ is a zero of F of multiplicity m_ρ , then K has a zero at

$$\lambda_\rho = -\rho^2 \quad (2.7)$$

of the same multiplicity.

Proof. Near $z = \rho$, write $F(z) = (z - \rho)^{m_\rho} h(z)$ with $h(\rho) \neq 0$. Evenness supplies the corresponding factor at $-\rho$. Their product is a nonvanishing even factor times $(z^2 - \rho^2)^{m_\rho}$. Substituting

$z^2 = -w$ gives a nonvanishing factor times $(-w - \rho^2)^{m_\rho}$. Hence the order at $w = -\rho^2$ is exactly m_ρ . $\square$

The proposition contains the geometric core of the construction. A point on the critical line is not merely distinguished in the original plane; after folding it becomes a point on the ordinary positive real axis. This is the reason reciprocal spectral variables are useful below.

2.3 The geometry of the spectral coordinates

The three coordinates

$$z, \quad w = -z^2, \quad q = -z^{-2} \quad (2.8)$$

should be regarded as different descriptions of one geometric operation rather than as unrelated substitutions. The centered coordinate z retains the complex location of a zero. The folded coordinate w removes the involution $z \leftrightarrow -z$. The reciprocal coordinate q reverses the scale so that the zeros nearest the origin in the folded plane become the largest reciprocal spectral points.

This last feature is useful analytically. The high zeros of the xi-function correspond to small values of q , and hence the reciprocal moment series has a natural expansion around the origin. In particular, the moments emphasize the low part of the reciprocal spectrum while still recording the entire zero set. The parameter that appears in the radius of convergence of T is consequently a geometric quantity: it is controlled by the first folded spectral point rather than by the growth at infinity.

There is also a simple sign dictionary. If $z = i\gamma$, then

$$w = \gamma^2 > 0, \quad q = \gamma^{-2} > 0. \quad (2.9)$$

If $z = a + ib$ with $a \neq 0$, then

$$-z^2 = b^2 - a^2 - 2iab, \quad (2.10)$$

which is not positive real unless $a = 0$. Thus positivity of the folded or reciprocal coordinate is not an approximation to the critical-line condition; it is exactly the critical-line condition written in a different coordinate.

2.4 Multiplicity and conjugation

It is worth being explicit about multiplicities because folding can otherwise lead to a factor-of-two error. Suppose ρ is a centered zero of multiplicity m_ρ . Evenness gives a zero of the same multiplicity at $-\rho$. These two factors are identified by the folding map, and the resulting folded zero still has multiplicity m_ρ . Complex conjugation is a separate symmetry. If $\bar{\rho}$ is not equal to either ρ or $-\rho$, then it gives a distinct folded point $-\bar{\rho}^2$, again with multiplicity m_ρ .

Consequently a sum over $\mathcal{Z}_c$ is a sum over the quotient by the involution $\rho \mapsto -\rho$, not a sum over unordered conjugate pairs. This convention is used in all spectral sums below. It is also the reason the measure obtained under RH has one atom for each positive ordinate with its original multiplicity.

2.5 Reality on the real axis

The completed xi-function is real on the real axis. Hence $F(r)$ and $A(r)$ are real for real r . The folded function inherits the corresponding reality property on the real w -axis where it is

evaluated away from its zeros. The reflected resolvent $S(x)$ is therefore real for real x away from poles.

This elementary reality statement is useful later in the Stieltjes inversion discussion. If a positive Stieltjes representation existed, its representing measure could not acquire an arbitrary continuous component from a jump of S across a cut where the meromorphic function has no imaginary boundary value. The only possible positive mass is concentrated at the poles required by the spectral divisor.

2.6 A local expansion at the origin

Because F is even and nonzero at the origin, one may write

$$F(z) = F(0) \left(1 + c_1 z^2 + c_2 z^4 + \cdots \right) \quad (2.11)$$

in a neighborhood of $z = 0$. The folded function therefore has an ordinary Taylor expansion

$$K(w) = F(0) \left(1 - c_1 w + c_2 w^2 - \cdots \right). \quad (2.12)$$

The logarithmic derivative at the origin is finite, and hence $S(0)$ is finite. In terms of the reciprocal moments,

$$S(x) = \sum_{n \geq 0} (-1)^n \mu_n x^n, \quad (2.13)$$

so

$$\mu_n = \frac{(-1)^n}{n!} S^{(n)}(0) \quad (2.14)$$

whenever the derivatives are taken inside the disk of analyticity. This gives a direct local interpretation of the moment sequence without referring to the individual zeros. The exact theta representation in Section 5 also gives $F(0) = 4A(0) > 0$, so the nonvanishing used here is independently visible within the present framework.

2.7 Normalization and branch-free notation

Whenever $x > 0$, we use the positive real square root $\sqrt{x}$. When x is complex, the expression $F(\sqrt{x})$ will only appear inside a branch-independent even combination. In particular,

$$\frac{F'(\sqrt{x})}{2\sqrt{x}F(\sqrt{x})} \quad (2.15)$$

is invariant under the replacement $\sqrt{x} \mapsto -\sqrt{x}$ and therefore defines a meromorphic function of x .

The notation

$$A(r) = \int_0^\infty \Phi(u) \cosh(ru) \, du \quad (2.16)$$

will be reserved for real r in the theta part of the paper. The variable r is not the same as the resolvent variable x ; throughout the theta analysis we use $x = r^2$.

2.8 Canonicity of the folded variable

There is a small but useful uniqueness statement behind the choice $w = -z^2$. Any holomorphic invariant of the involution $z \mapsto -z$ is a holomorphic function of z^2 . Thus the quotient coordinate

is unique up to a holomorphic reparametrization. The sign in $w = -z^2$ is selected so that the critical line is represented by positive real values. This is the only sign choice needed later.

The construction also explains why the square is preferable to a direct projection onto $\Re z$. A projection would lose analyticity and would not preserve the divisor in a meromorphic coordinate. The map $z \mapsto -z^2$ is holomorphic, separates the off-axis geometry through an imaginary part, and retains multiplicities. It is therefore compatible with the entire-function structure of F .

2.9 The resolvent near and away from the spectrum

The meromorphic function S is regular at the origin and has isolated poles at the folded spectral points. Away from those poles, the logarithmic derivative satisfies the usual resolvent identity

$$S'(x) = - \sum_{\rho} \frac{m_{\rho}}{(x - \rho^2)^2}, \quad (2.17)$$

whenever the series is interpreted in a normally convergent region. Higher derivatives satisfy

$$S^{(n)}(x) = (-1)^n n! \sum_{\rho} \frac{m_{\rho}}{(x - \rho^2)^{n+1}}. \quad (2.18)$$

These formulas are formal consequences of the product representation and become genuine identities by local uniform convergence.

Under RH the denominators are $x + \gamma^2$, which gives alternating signs immediately. Away from RH there is no comparable order on the complex poles. This is another way of seeing why complete monotonicity is naturally downstream from the critical-line condition.

2.10 Logical status of the framework

The constructions in Sections 2–4 are exact. They do not rely on an assumed location of the zeros. The implication from the Riemann Hypothesis to positivity of the moment matrices is equally exact. The converse uses the fixed meromorphic function S and the fact that a positive Stieltjes representation has support on the nonnegative real axis. No numerical approximation is needed at this stage.

The theta calculations in Sections 5–8 are also exact. What is not asserted is that the resulting theta identities by themselves already imply positivity of the complete Hankel hierarchy. The distinction is essential. An exact reduction can expose the precise object that must be positive without itself establishing that positivity.

3 Folded Resolvent and Reciprocal Spectral Variables

3.1 The folded logarithmic derivative

The zero counting for the xi-function gives convergence of the reciprocal square sum

$$\sum_{\rho \in \mathcal{Z}_c} \frac{m_{\rho}}{|\rho|^2} < \infty. \quad (3.1)$$

Consequently the folded product can be represented in genus-zero form,

$$K(w) = K(0) \prod_{\rho \in \mathcal{Z}_c} \left(1 + \frac{w}{\rho^2}\right)^{m_\rho}, \quad (3.2)$$

with locally uniform convergence away from the zeros. Logarithmic differentiation gives

$$R(w) = -\frac{K'(w)}{K(w)} = -\sum_{\rho \in \mathcal{Z}_c} \frac{m_\rho}{w + \rho^2}. \quad (3.3)$$

Setting $w = -x$ yields

$$\boxed{S(x) = R(-x) = \sum_{\rho \in \mathcal{Z}_c} \frac{m_\rho}{x - \rho^2}.} \quad (3.4)$$

Direct differentiation of $K(w) = F(i\sqrt{w})$ gives the equivalent identity

$$S(x) = \frac{F'(\sqrt{x})}{2\sqrt{x}F(\sqrt{x})}. \quad (3.5)$$

The apparent singularity at $x = 0$ is removable because F'/F is odd.

Under the Riemann Hypothesis, every centered zero is $\rho = i\gamma$, and therefore

$$S(x) = \sum_{\gamma > 0} \frac{m_\gamma}{x + \gamma^2}, \quad x > 0, \quad (3.6)$$

where each conjugate pair is represented once in $\gamma > 0$ and multiplicities are retained.

Proposition 3.1. *Assuming RH, S is a Stieltjes transform of the positive discrete measure*

$$\nu = \sum_{\gamma > 0} m_\gamma \delta_{\gamma^2}. \quad (3.7)$$

In particular,

$$S(x) = \int_0^\infty \frac{d\nu(t)}{x + t}. \quad (3.8)$$

Proof. Substituting $t = \gamma^2$ in the displayed sum gives the integral representation. Positivity and local finiteness follow from the discreteness and multiplicities of the zeros together with the convergence of the corresponding resolvent series. $\square$

The proposition also gives the complete-monotonicity hierarchy

$$(-1)^n S^{(n)}(x) = n! \sum_{\gamma > 0} \frac{m_\gamma}{(x + \gamma^2)^{n+1}} > 0, \quad n \geq 0. \quad (3.9)$$

This is stronger than a first-derivative test, but it is still only one consequence of the positive Stieltjes representation.

3.2 Reciprocal variables

Define

$$q_\rho = -\frac{1}{\rho^2}. \quad (3.10)$$

Then

$$S(-u) = \sum_{\rho \in \mathcal{Z}_c} \frac{m_\rho q_\rho}{1 - q_\rho u} =: T(u) \quad (3.11)$$

for u in a neighborhood of the origin. Expanding the geometric series gives

$$T(u) = \sum_{n=0}^{\infty} \mu_n u^n, \quad \mu_n = \sum_{\rho \in \mathcal{Z}_c} m_\rho q_\rho^{n+1}. \quad (3.12)$$

The absolute convergence of each moment follows from the convergence of the reciprocal-square sum and the rapid growth of the zero ordinates.

When RH holds, $q_\rho = 1/\gamma^2 > 0$, and therefore

$$\mu_n = \sum_{\gamma > 0} m_\gamma \gamma^{-2n-2} > 0. \quad (3.13)$$

The sequence is then a moment sequence for the positive measure

$$d\sigma(t) = t d\nu_{\text{rec}}(t), \quad \nu_{\text{rec}} = \sum_{\gamma > 0} m_\gamma \delta_{1/\gamma^2}, \quad (3.14)$$

where the shorthand is intended only to emphasize the reciprocal support.

3.3 Pole geometry and the Stieltjes cut

The relation between poles of S and zeros of F is exact:

$$x = \rho^2 \quad (3.15)$$

is a pole location, with the same multiplicity. A Stieltjes function has an analytic continuation to $\mathbb{C} \setminus (-\infty, 0]$ with all singularities lying on the nonpositive real axis. Therefore a positive Stieltjes representation for the fixed meromorphic function S forces

$$\rho^2 \in (-\infty, 0] \quad (3.16)$$

for every pole, hence $\rho \in i\mathbb{R}$.

This observation is useful because it shows that there is no hidden freedom in the measure. If one obtains a positive Stieltjes representation for S , the support is determined by the poles of S and hence by the squared centered zeros. The analytic representation cannot simply choose a different positive measure with unrelated support.

3.4 The resolvent as a divisor-sensitive function

The logarithmic derivative remembers the zero divisor but forgets the nonzero multiplicative normalization of K . This is exactly what is needed here. If K is replaced by cK with $c \neq 0$, then $R = -K'/K$ is unchanged. The spectral information is therefore concentrated in the quotient and not in the normalization of the xi-function.

For later use, note that the poles of S have positive integer residues when counted with the convention adopted here. Indeed,

$$S(x) = \frac{m_\rho}{x - \rho^2} + \text{holomorphic part} \quad (3.17)$$

near a pole associated with ρ . Hence the residue is m_ρ . This positivity of residues is independent of RH. The difficulty is the location of the poles, not the sign of the residues.

3.5 Meromorphic continuation and support rigidity

Suppose that one has a positive measure σ for which

$$T(u) = \int_0^\infty \frac{d\sigma(t)}{1-tu} \quad (3.18)$$

holds near $u = 0$. Then, after substituting $u = -1/x$, one has $T(-1/x) = xS_\sigma(x)$ with $S_\sigma(x) = \int_0^\infty (x+t)^{-1} d\sigma(t)$. Hence the singularities of S can only occur on the negative real x -axis. The important point is that this is a statement about analytic continuation, not merely about the signs of Taylor coefficients.

On the other hand, the spectral expression gives the poles explicitly at $x = \rho^2$. Thus a positive representation cannot coexist with an off-axis pole. The support constraint is therefore rigid enough to read the critical-line condition directly from the complex geometry of the resolvent.

3.6 Residues and the forced atomic measure

The Stieltjes inversion principle also gives a more concrete description. If S is represented as

$$S(x) = \int_0^\infty \frac{d\nu(t)}{x+t}, \quad (3.19)$$

then an atom of mass m at $t = t_0$ contributes $m/(x+t_0)$ and hence produces a simple pole at $x = -t_0$ with residue m . Because the spectral continuation of S is meromorphic and has only isolated poles, any positive representing measure must be supported on the corresponding pole set; in particular, no additional continuous component can produce a branch cut. Its possible atoms therefore have to occur at

$$t_\rho = -\rho^2. \quad (3.20)$$

Positivity of the measure then requires these spectral locations to lie in $[0, \infty)$.

This observation is stronger than a numerical inversion. Numerical inversion can suggest a positive density but cannot by itself rule out small negative spectral components or identify the exact analytic continuation. Here the divisor already fixes the possible atoms, so the only missing ingredient is a proof that all of them lie in the permitted half-line.

3.7 The reciprocal transform and normalization of moments

The generating function T is normalized so that its zeroth coefficient is

$$\mu_0 = \sum_\rho m_\rho q_\rho. \quad (3.21)$$

The corresponding positive measure under RH is therefore not simply $\sum m_\rho \delta_{q_\rho}$ but rather

$$\sigma = \sum_\rho m_\rho q_\rho \delta_{q_\rho}. \quad (3.22)$$

This distinction matters for the shifted Hankel family. The shift inserts one additional factor of q_ρ , which is exactly the factor appearing in the second Stieltjes positivity condition.

3.8 Uniqueness of the reciprocal representation

Suppose two positive measures on a compact interval have the same moment sequence. Their polynomial integrals agree, and polynomial approximation then implies equality on continuous functions. In the present problem the reciprocal spectral support accumulates only at zero and has a finite upper endpoint determined by the lowest folded spectral point. The spectral moment representation is therefore naturally rigid.

The same observation can be phrased through generating functions. If two measures have the same moments and both transforms are analytic in a common neighborhood of the origin, their power series agree there. Analytic continuation then identifies the transforms wherever both continuations exist. In the present setting the meromorphic spectral continuation supplies an especially strong uniqueness mechanism.

3.9 A support test independent of individual residues

The pole-location argument does not require detailed knowledge of the multiplicities beyond positivity of the residues. If a positive Stieltjes measure exists, every pole of S must lie on the negative real axis, because the Stieltjes kernel $1/(x+t)$ has its singularity at $x = -t$. Thus even if the residue data were complicated, the support test alone would force the squared centered zeros to be nonpositive.

This separation is useful conceptually. The zero multiplicities control weights in the moment problem, while the real part of the zero controls the geometric location of the support. The Riemann Hypothesis is therefore read as a support statement, not as a weight statement.

3.10 Divided differences

Two kernels package the moment hierarchy conveniently. Define

$$\mathcal{H}_0(u, v) = \frac{uT(u) - vT(v)}{u - v}, \quad (3.23)$$

and

$$\mathcal{H}_1(u, v) = \frac{T(u) - T(v)}{u - v}. \quad (3.24)$$

Expanding at the origin gives

$$\mathcal{H}_0(u, v) = \sum_{i,j \geq 0} \mu_{i+j} u^i v^j, \quad (3.25)$$

and

$$\mathcal{H}_1(u, v) = \sum_{i,j \geq 0} \mu_{i+j+1} u^i v^j. \quad (3.26)$$

Thus the coefficients of the two divided differences are exactly the entries of the two Hankel families. Positivity of those coefficient matrices is therefore equivalent to the corresponding Hankel positivity conditions. These divided differences are the moment-theoretic shadow of the folded logarithmic derivative.

4 Stieltjes Moments and Hankel Positivity

4.1 Reality of the reciprocal moments

Complex conjugation is a second symmetry of the centered zero set. If ρ is a centered zero of multiplicity m_ρ , then $\bar{\rho}$ is a zero of the same multiplicity, and

$$q_{\bar{\rho}} = \overline{q_\rho}. \quad (4.1)$$

Because complex conjugation either preserves a chosen representative or sends it to the negative of another chosen representative, the quotient set $\mathcal{Z}_c$ is paired by conjugation up to the involution $\rho \mapsto -\rho$. Consequently the terms in the reciprocal sums occur in conjugate pairs. Hence

$$\mu_n = \sum_{\rho \in \mathcal{Z}_c} m_\rho q_\rho^{n+1} \in \mathbb{R}. \quad (4.2)$$

This establishes the realness needed for the ordinary real moment problem before any assumption of RH.

4.2 Moment matrices

For a real sequence $(\mu_n)_{n \geq 0}$ define

$$H_N = (\mu_{i+j})_{i,j=0}^N, \quad \tilde{H}_N = (\mu_{i+j+1})_{i,j=0}^N. \quad (4.3)$$

For a polynomial $p(t) = \sum_{j=0}^N a_j t^j$, with coefficient vector $\mathbf{a} = (a_0, \dots, a_N)^T$,

$$\mathbf{a}^* H_N \mathbf{a} = \sum_{i,j=0}^N \mu_{i+j} \bar{a}_i a_j, \quad (4.4)$$

while

$$\mathbf{a}^* \tilde{H}_N \mathbf{a} = \sum_{i,j=0}^N \mu_{i+j+1} \bar{a}_i a_j. \quad (4.5)$$

If $\mu_n = \int_0^\infty t^n d\sigma(t)$ for a positive measure σ , then these quadratic forms are respectively

$$\int_0^\infty |p(t)|^2 d\sigma(t) \quad (4.6)$$

and

$$\int_0^\infty t |p(t)|^2 d\sigma(t), \quad (4.7)$$

so both matrices are positive semidefinite.

The converse is the classical one-dimensional Stieltjes moment criterion. We record it here in the form needed for the present spectral sequence.

Theorem 4.1. *For a real sequence (μ_n) , the following are equivalent:*

- (i) $H_N \succeq 0$ and $\tilde{H}_N \succeq 0$ for every N ;
- (ii) *there exists a positive Borel measure σ on $[0, \infty)$ such that $\mu_n = \int_0^\infty t^n d\sigma(t)$ for every n .*

For the spectral sequence, the reciprocal quantities satisfy $|q_\rho| \rightarrow 0$ and are uniformly bounded. This gives an exponential moment bound for (μ_n) of the form $|\mu_n| \leq CQ^n$ for some finite Q . If a positive representing measure exists, the same bound forces its support to be contained in

$[0, Q]$, since otherwise its moments would grow faster than Q^n . Thus the representing measure is compactly supported and its generating function is analytic in a neighborhood of the origin. Once a positive measure is obtained, the meromorphic continuation of T determines its atoms through the poles.

4.3 Exact equivalence with RH

The preceding moment criterion becomes spectral when applied to the sequence generated by T . The implication $\text{RH} \Rightarrow \text{Hankel positivity}$ is immediate from $q_\rho > 0$. The converse is just as rigid.

Theorem 4.2 (Spectral Stieltjes equivalence). *For the reciprocal spectral moments defined above,*

$$\text{RH} \iff H_N \succeq 0 \text{ and } \tilde{H}_N \succeq 0 \text{ for every } N. \quad (4.8)$$

Equivalently,

$$\text{RH} \iff T(u) = \int_0^\infty \frac{d\sigma(t)}{1-tu} \quad (4.9)$$

for some positive measure σ in a neighborhood of $u = 0$.

Proof. Assume RH. Then $q_\rho > 0$ for every representative and

$$\mu_n = \sum_\rho m_\rho q_\rho^{n+1} = \int_0^\infty t^n d\sigma(t), \quad (4.10)$$

with

$$\sigma = \sum_\rho m_\rho q_\rho \delta_{q_\rho}. \quad (4.11)$$

Hence both Hankel families are positive semidefinite.

Conversely, suppose both Hankel families are positive semidefinite for all N . By the Stieltjes moment criterion there is a positive measure σ on $[0, \infty)$ with moments (μ_n) . Its generating function is a Stieltjes transform in the variable u and its reciprocal form gives a function $S(x)$ analytic away from the negative real axis in the appropriate x -plane. But the independently defined meromorphic continuation of S has poles at $x = \rho^2$. Therefore every pole must lie on $(-\infty, 0]$, so $\rho^2 \leq 0$. Since the centered zeros are nonzero, ρ is purely imaginary. This is RH. $\square$

4.4 Cauchy-Binet and finite minors

For a finite positive discrete measure

$$\sigma_M = \sum_{k=1}^M w_k \delta_{t_k}, \quad w_k > 0, \quad (4.12)$$

the principal Hankel determinant has the Cauchy-Binet expansion

$$\det(\mu_{i+j})_{i,j=0}^N = \sum_{1 \leq k_0 < \dots < k_N \leq M} \left(\prod_{j=0}^N w_{k_j} \right) \prod_{0 \leq a < b \leq N} (t_{k_b} - t_{k_a})^2. \quad (4.13)$$

Every term is nonnegative. The corresponding identity for the shifted family acquires an additional factor $\prod_j t_{k_j}$. This makes the geometric meaning of the two Hankel conditions transparent: they are squared-volume expressions for the evaluation vectors of monomials on the positive spectral support.

For the infinite spectral measure the same identity is interpreted by monotone finite truncation whenever the relevant moments are finite. The determinant positivity is therefore naturally compatible with the discrete zero geometry.

4.5 Low-order conditions

The first two Hankel conditions are

$$\mu_0 \geq 0, \quad \mu_0\mu_2 - \mu_1^2 \geq 0, \quad (4.14)$$

and

$$\mu_1 \geq 0, \quad \mu_1\mu_3 - \mu_2^2 \geq 0. \quad (4.15)$$

Under RH they are ordinary variance inequalities for the positive spectral measure and its first tilt. The difficulty is not the meaning of any fixed determinant; it is the demand that every determinant in both families be nonnegative simultaneously.

4.6 Quadratic forms and polynomial geometry

The Hankel matrices are best viewed through polynomials rather than entries. Let

$$p(t) = a_0 + a_1t + \cdots + a_Nt^N. \quad (4.16)$$

Then

$$\mathbf{a}^* H_N \mathbf{a} = \sum_{i,j=0}^N \mu_{i+j} \overline{a_i} a_j. \quad (4.17)$$

If a representing measure exists, this becomes

$$\int_0^\infty |p(t)|^2 d\sigma(t). \quad (4.18)$$

Likewise, the shifted form is

$$\mathbf{a}^* \tilde{H}_N \mathbf{a} = \int_0^\infty t |p(t)|^2 d\sigma(t). \quad (4.19)$$

Thus the two families test positivity on the cone of squares and on the cone obtained after multiplication by t . The half-line support is encoded in the simultaneous positivity of these two cones.

4.7 A determinantal interpretation

The principal minors of H_N form the Gram determinants of the monomial vectors

$$(1, t, t^2, \dots, t^N) \quad (4.20)$$

under the inner product induced by σ . A vanishing determinant means that the corresponding monomials become linearly dependent in $L^2(\sigma)$. Thus strict positivity of the minors reflects genuine spectral diversity. Under a simple discrete spectrum one expects strict positivity of every principal minor of every finite order.

The Cauchy-Binet formula makes this statement explicit for finite support. If

$$\sigma = \sum_{k=1}^M w_k \delta_{t_k}, \quad t_1 < \cdots < t_M, \quad (4.21)$$

then every nonzero contribution to the $(N+1) \times (N+1)$ determinant is a squared Vandermonde factor. The shifted determinant inserts $\prod t_k$. No sign ambiguity remains once the support is positive.

4.8 The two Hankel families as a single positivity cone

The pair $(H_N, \tilde{H}_N)$ can be combined into a positivity statement for polynomials of the form

$$p(t)^2 + tq(t)^2. \quad (4.22)$$

The corresponding functional is nonnegative when the support lies in $[0, \infty)$. Conversely, positivity on this quadratic module is the natural one-dimensional certificate for a half-line moment problem. This viewpoint is compatible with the geometric folding: the positive half-line is the image of the critical line under $\rho \mapsto -\rho^2$.

4.9 Why the moment formulation is the correct interface

The theta side naturally produces integrals against positive functions of a continuous variable u , while the spectral side naturally produces discrete sums over zeros. The moment functional is the interface between these two descriptions. It is insensitive to the choice of coordinates used to represent the same function, but it is sensitive to positivity of the spectral support. This is precisely the feature needed for a theta-geometric attack.

4.10 Shifted positivity and the half-line condition

The shifted Hankel matrix is not an optional refinement. A sequence can be a Hamburger moment sequence on the whole real line while failing to be a Stieltjes moment sequence on $[0, \infty)$. The extra factor of t in

$$\int_0^\infty t |p(t)|^2 d\sigma(t) \quad (4.23)$$

encodes the orientation of the support. In the folded spectral geometry this orientation is exactly the choice that sends the critical line to positive $-\rho^2$.

4.11 Finite minors as geometric volumes

For a positive discrete measure the Cauchy-Binet formula shows that the determinant of H_N is a sum of squared Vandermonde volumes. The corresponding shifted determinant is the same sum with an additional positive weight. This gives a geometric interpretation of all finite positivity tests: every test measures a squared volume in the positive spectral coordinate.

This picture suggests why a theta proof should seek a Gram representation. If the theta side can be converted into an inner product of functions indexed by monomials in the reciprocal coordinate, the entire Hankel hierarchy becomes positive at once. The difficulty is to construct those functions without already inserting the sign of the spectral support by hand.

4.12 Why first-order information is insufficient

Complete monotonicity of S on $(0, \infty)$ is weaker than the Stieltjes property. Inverse Laplace transforms can remain positive without automatically being completely monotone in the auxiliary variable needed for a Stieltjes representation. Likewise, positivity of the first derivative

$$-S'(x) \geq 0 \quad (4.24)$$

does not by itself imply the Hankel hierarchy. The theta analysis below is therefore organized not merely around one derivative, but around the algebraic structures that may propagate to all orders.

5 Theta Representation and the Geometric Measure

5.1 The theta density

Introduce

$$\Theta(u) = \sum_{n=1}^{\infty} e^{-\pi n^2 e^{2u}}, \quad Y(u) = e^{u/2} \Theta(u). \quad (5.1)$$

The differentiated theta kernel appearing in the xi-integral is

$$\Phi(u) = \frac{1}{2} \left(Y''(u) - \frac{1}{4} Y(u) \right). \quad (5.2)$$

Writing the derivatives term by term gives the explicit series

$$\Phi(u) = \sum_{n=1}^{\infty} \pi n^2 e^{5u/2} \left(2\pi n^2 e^{2u} - 3 \right) e^{-\pi n^2 e^{2u}}. \quad (5.3)$$

For $u \geq 0$ the individual terms are nonnegative because $2\pi n^2 e^{2u} \geq 2\pi > 3$. Thus $\Phi(u) > 0$ on $[0, \infty)$.

The xi-function has the exact cosine representation

$$\Xi(t) = 4 \int_0^{\infty} \Phi(u) \cos(tu) du, \quad (5.4)$$

where $\Xi(t) = \xi(1/2 + it)$, and hence the centered real-axis representation

$$\boxed{F(r) = 4 \int_0^{\infty} \Phi(u) \cosh(ru) du.} \quad (5.5)$$

The integral converges for every real r because the Gaussian-type theta decay dominates the exponential factor from $\cosh(ru)$.

5.2 Mellin normalization

The same kernel has a Mellin description. For $\Re s > 1/2$,

$$\int_{-\infty}^{\infty} Y(u) e^{su} du = \frac{1}{2} \pi^{-s/2-1/4} \Gamma\left(\frac{s}{2} + \frac{1}{4}\right) \zeta\left(s + \frac{1}{2}\right). \quad (5.6)$$

Using the theta symmetry

$$Y(u) - Y(-u) = -\sinh(u/2) \quad (5.7)$$

and splitting the integral at 0, one finds

$$J(s) := \int_0^\infty Y(u) \cosh(su) \, du = \frac{1}{4} \pi^{-s/2-1/4} \Gamma\left(\frac{s}{2} + \frac{1}{4}\right) \zeta\left(s + \frac{1}{2}\right) - \frac{1}{4(s^2 - 1/4)}. \quad (5.8)$$

Since $\Phi = (Y'' - Y/4)/2$, integration by parts gives

$$F(r) = 2 \left(r^2 - \frac{1}{4}\right) J(r) + \frac{1}{2}. \quad (5.9)$$

Substituting the expression for J yields the exact completion formula, with the value at $r = 1/2$ understood by analytic continuation through the removable singularity,

$$\boxed{F(r) = \frac{1}{2} \left(r^2 - \frac{1}{4}\right) \pi^{-r/2-1/4} \Gamma\left(\frac{r}{2} + \frac{1}{4}\right) \zeta\left(r + \frac{1}{2}\right)}. \quad (5.10)$$

The normalization in this formula is fixed by the original definition of ξ and will be used only as a consistency check for the theta representation.

5.3 The geometric measure

Define

$$A(r) = \int_0^\infty \Phi(u) \cosh(ru) \, du, \quad F(r) = 4A(r). \quad (5.11)$$

For real r , $A(r) > 0$. It is therefore natural to regard

$$dP_r(u) = \frac{\Phi(u) \cosh(ru)}{A(r)} \, du \quad (5.12)$$

as a probability measure on $[0, \infty)$. The family $(P_r)_{r \geq 0}$ is the geometric measure underlying the theta side of the resolvent.

Differentiation gives

$$A'(r) = \int_0^\infty u \Phi(u) \sinh(ru) \, du. \quad (5.13)$$

The ratio A'/A is thus an expectation:

$$\frac{A'(r)}{A(r)} = \mathbb{E}_r[U \tanh(rU)]. \quad (5.14)$$

Since

$$S(r^2) = \frac{F'(r)}{2rF(r)} = \frac{A'(r)}{2rA(r)}, \quad (5.15)$$

we obtain

$$\boxed{S(r^2) = \frac{1}{2r} \mathbb{E}_r[U \tanh(rU)]}. \quad (5.16)$$

This is the basic geometric expression used below.

5.4 Theta symmetry and boundary data

The theta transformation gives

$$Y(u) - Y(-u) = -\sinh(u/2). \quad (5.17)$$

Evaluating at $u = 0$ gives no information about $Y(0)$, but differentiating gives

$$Y'(0) = -\frac{1}{4}. \quad (5.18)$$

The value $Y(0)$ remains the positive theta sum

$$Y(0) = \sum_{n=1}^{\infty} e^{-\pi n^2}. \quad (5.19)$$

These two boundary quantities are precisely the ones that reappear when the differential identity for Φ is integrated by parts.

5.5 Decay and differentiation under the integral sign

The theta density has the large- u form

$$\Phi_n(u) = O\left(e^{5u/2} n^4 e^{2u} e^{-\pi n^2 e^{2u}}\right), \quad (5.20)$$

so every polynomial factor in u and every fixed exponential factor e^{au} is dominated as $u \rightarrow \infty$. This justifies differentiation with respect to r in the cosh representation and the repeated integrations by parts used later. Near $u = 0$, the theta series converges normally together with all derivatives.

5.6 The positive density is not the spectral measure

It is tempting to identify $\Phi(u) du$ directly with the desired Stieltjes measure. The variables do not permit this. The spectral measure lives in $t = q_\rho$, whereas u is a continuous theta parameter. The nonlinear map from Φ to S is a logarithmic derivative, and logarithmic differentiation is precisely where the spectral correlations enter.

The tilt P_r makes this distinction explicit. It is a probability measure indexed by the external parameter r , not a fixed measure whose moments are the reciprocal spectral moments. Any proof that identifies the two without accounting for the r -dependence would be proving a different statement.

5.7 Symmetry at the lower endpoint

The half-line integral is not arbitrary. The lower endpoint $u = 0$ is the fixed point of the theta involution. The identity

$$Y'(0) = -\frac{1}{4} \quad (5.21)$$

therefore carries information from the global theta symmetry into a single boundary scalar. This is one reason the operator reduction is accompanied by a boundary defect: quotienting the theta symmetry onto $[0, \infty)$ leaves a trace at its fixed point.

5.8 A branch-free geometric identity

The representation of S through A can also be written in logarithmic form. Set

$$L(r) = \log A(r). \quad (5.22)$$

Then

$$S(r^2) = \frac{L'(r)}{2r}. \quad (5.23)$$

This relation is exact and remains useful when derivatives in $x = r^2$ are converted to the Euler operator $D = r \frac{d}{dr}$.

6 Tilted Theta Geometry

6.1 Termwise theta structure

The explicit mode

$$\Phi_n(u) = \pi n^2 e^{5u/2} (2\pi n^2 e^{2u} - 3) e^{-\pi n^2 e^{2u}} \quad (6.1)$$

contains two competing factors: a positive algebraic prefactor and a rapidly decreasing exponential. On $u \geq 0$, the algebraic factor is positive because $2\pi n^2 e^{2u} > 3$. The exponential term supplies more than enough decay to justify repeated differentiation and integration against polynomially growing test functions.

This decomposition is more than a technical convenience. It shows that the positivity of the theta density is modewise on the half-line, whereas the positivity sought for the spectral moments is global after taking a logarithmic derivative. The gap between these two types of positivity is the central analytic issue of the paper.

6.2 The cosh transform and strict positivity

For real r , every factor in

$$\Phi(u) \cosh(ru) \quad (6.2)$$

is nonnegative on $u \geq 0$, and Φ is not identically zero. Hence

$$A(r) > 0. \quad (6.3)$$

The probability density P_r is therefore well defined for every real r . This strict positivity is one reason the theta representation is well suited to expectation methods.

It is useful to distinguish this statement from positivity of derivatives of A . The derivative A' contains $\sinh(ru)$ and has the same sign as r . Higher derivatives alternate between the two hyperbolic functions and do not produce an immediate total-positivity statement. The nonlinear ratio A'/A is the genuinely interesting object.

6.3 The measure as an exponential tilt

For $r_2 > r_1 \geq 0$, the Radon-Nikodym ratio of the tilted measures is

$$\frac{dP_{r_2}}{dP_{r_1}}(u) = \frac{A(r_1) \cosh(r_2 u)}{A(r_2) \cosh(r_1 u)}. \quad (6.4)$$

The ratio is increasing in u . Thus the family is ordered by an explicit monotone likelihood ratio. This observation explains why ordinary covariance inequalities are available, but it does not by itself yield the higher Hankel positivity.

6.4 A useful monotonicity identity

If f is independent of r , then

$$\frac{d}{dr} \mathbb{E}_r[f(U)] = \text{Cov}_r(f(U), U \tanh(rU)). \quad (6.5)$$

If f is increasing on $[0, \infty)$, then the covariance is nonnegative because $u \mapsto u \tanh(ru)$ is increasing for $r \geq 0$. This gives a simple monotonicity principle for the tilted family.

The resolvent observable is itself r -dependent, so the preceding principle cannot be applied directly to $S(r^2)$. Nevertheless, it explains the appearance of covariance terms and suggests that order properties of the tilt may be useful in controlling them.

6.5 The nonlinear observable

Let

$$b(y) = y \tanh y. \quad (6.6)$$

For $x = r^2$ define

$$G_r(u) = \frac{u \tanh(ru)}{2r}. \quad (6.7)$$

Then

$$S(r^2) = \mathbb{E}_r[G_r(U)]. \quad (6.8)$$

The dependence of both G_r and P_r on r is important. One cannot differentiate only the observable while freezing the measure.

A direct differentiation gives

$$\frac{d}{dr} \log P_r(u) = u \tanh(ru) - \frac{A'(r)}{A(r)}. \quad (6.9)$$

Therefore the derivative of an expectation is

$$\frac{d}{dr} \mathbb{E}_r[f_r(U)] = \mathbb{E}_r[\partial_r f_r(U)] + \text{Cov}_r(f_r(U), U \tanh(rU)). \quad (6.10)$$

This identity is exact and makes the covariance structure unavoidable.

6.6 First derivative of the resolvent

A more symmetric formula is obtained by differentiating $S(r^2) = A'(r)/(2rA(r))$ directly. Define

$$W_r(u) = \frac{u}{4r^3} \left[\tanh(ru) - ru \operatorname{sech}^2(ru) \right]. \quad (6.11)$$

Then a calculation gives

$$-S'(r^2) = \mathbb{E}_r[W_r(U)] - \text{Var}_r(G_r(U)). \quad (6.12)$$

Equivalently,

$$-S'(x) = \mathbb{E}_{\rho_r}[-\partial_x G_x(U^2)] - \text{Var}_{\rho_r}(G_x(U^2)), \quad (6.13)$$

where

$$\rho_x(u) = \frac{\Phi(u) \cosh(u\sqrt{x})}{A(\sqrt{x})} \quad (6.14)$$

and

$$G_x(z) = \frac{\sqrt{z} \tanh(\sqrt{xz})}{2\sqrt{x}}. \quad (6.15)$$

The two forms are the same identity in different variables.

The sign problem is already visible here. The first term is positive, while the variance is subtractive. A proof of $S' \leq 0$ therefore requires a structural inequality rather than pointwise positivity of the observable.

6.7 Fixed-coordinate Stieltjes expansion

For fixed u and $x > 0$, set

$$f_x(u) = \frac{u \tanh(u\sqrt{x})}{2\sqrt{x}}. \quad (6.16)$$

The partial-fraction expansion of the hyperbolic tangent gives

$$\frac{\tanh z}{z} = 8 \sum_{k=0}^{\infty} \frac{1}{\pi^2(2k+1)^2 + 4z^2}. \quad (6.17)$$

Hence

$$f_x(u) = \sum_{k=0}^{\infty} \frac{1}{x + a_k/u^2}, \quad a_k = \frac{\pi^2(2k+1)^2}{4}. \quad (6.18)$$

For each fixed u , this is a genuine Stieltjes function of x . The difficulty is that the theta probability measure itself depends on x . Thus the fixed-coordinate decomposition cannot be integrated termwise into a fixed positive Stieltjes measure without controlling the parameter dependence of the tilt.

6.8 A useful rescaling

Set

$$G_x(z) = f_x(\sqrt{z}) = \sum_{k=0}^{\infty} \frac{z}{xz + a_k}. \quad (6.19)$$

Then

$$G_{x,z}(z) > 0, \quad G_{x,zz}(z) < 0, \quad (6.20)$$

and

$$-G_{x,x}(z) = \sum_{k=0}^{\infty} \frac{z^2}{(xz + a_k)^2}. \quad (6.21)$$

An exact algebraic relation follows:

$$-G_{x,x}(z) = \frac{G_x(z) - zG_{x,z}(z)}{x}. \quad (6.22)$$

This identity is occasionally useful for converting x -derivatives into z -derivatives, but it does not by itself remove the variance term created by the theta tilt.

6.9 Covariance as the source of nonlocality

For the tilted family, differentiation produces

$$\frac{d}{dr} \mathbb{E}_r[f_r] = \mathbb{E}_r[f'_r] + \text{Cov}_r(f_r, h_r), \quad h_r(u) = u \tanh(ru). \quad (6.23)$$

The second term is intrinsically two-point. No pointwise estimate on Φ can remove it. The first derivative of S is the first place where the logarithmic nature of the resolvent becomes visible in the theta variables.

This observation also explains why the first derivative is the natural place to test a proposed positivity mechanism. If a method cannot handle the covariance term correctly at first order, there is little reason to expect it to control the higher Hankel hierarchy. Conversely, a clean treatment of the covariance may reveal the structure that needs to be iterated.

6.10 An integral identity for the variance

For any square-integrable function g under P_r ,

$$\text{Var}_r(g(U)) = \frac{1}{2} \iint (g(u) - g(v))^2 dP_r(u) dP_r(v). \quad (6.24)$$

Applying this with the appropriate rescaled g gives the two-variable kernel representation of $-S'$. The square difference is therefore not an ad hoc algebraic device; it is exactly the standard two-point form of a variance.

The remaining one-point term in the kernel comes from differentiating the observable itself. This separation is useful because it identifies which part is automatically positive and which part must be compared against it.

6.11 The logarithmic geometry

The logarithmic derivative has a particularly compact form. Since

$$S(r^2) = \frac{L'(r)}{2r}, \quad (6.25)$$

we obtain

$$-4r^4 S'(r^2) = r(L'(r) - rL''(r)). \quad (6.26)$$

Using the tilted measure and $b(y) = y \tanh y$, the same quantity can be written as

$$\boxed{-4r^4 S'(r^2) = \mathbb{E}_r[b(Ur)]^2 + \mathbb{E}_r[b(Ur)] - \mathbb{E}_r[(Ur)^2]}. \quad (6.27)$$

The right-hand side is an exact expression for the first derivative sign problem. It isolates a squared mean, a positive mean, and a quadratic moment. No pointwise argument is hidden inside the formula.

6.12 A failed local pointwise strategy

One might try to solve a Stein equation for the tilted measure and dominate the variance by a pointwise derivative term. Near $u = 0$ the relevant expansion is

$$G_x(u^2) = \frac{u^2}{2} - \frac{xu^4}{6} + O(u^6), \quad (6.28)$$

and hence

$$-\partial_x G_x(u^2) = \frac{u^4}{6} + O(u^6). \quad (6.29)$$

A Stein solution with the natural regularity at the origin has $h'(u) = O(u)$, so $h'(u)^2 = O(u^2)$. Thus a pointwise bound of the form $h'(u)^2 \leq -\partial_x G_x(u^2)$ cannot hold near $u = 0$. This rules out a tempting sufficient condition and explains why the true mechanism, if it exists, must be nonlocal.

7 Differential Identities and Kernel Structures

7.1 A direct covariance derivation

Starting from

$$S(r^2) = \frac{1}{2r} \mathbb{E}_r[U \tanh(rU)], \quad (7.1)$$

write $h_r(u) = u \tanh(ru)$. Then

$$\frac{d}{dr} S(r^2) = -\frac{1}{2r^2} \mathbb{E}_r[h_r(U)] + \frac{1}{2r} \mathbb{E}_r[\partial_r h_r(U)] + \frac{1}{2r} \text{Cov}_r(h_r(U), h_r(U)). \quad (7.2)$$

Since the covariance of a variable with itself is its variance,

$$\frac{d}{dr} S(r^2) = -\frac{1}{2r^2} \mathbb{E}_r[h_r] + \frac{1}{2r} \mathbb{E}_r[U^2 \text{sech}^2(rU)] + \frac{1}{2r} \text{Var}_r(h_r). \quad (7.3)$$

Converting from d/dr to d/dx with $x = r^2$ gives an equivalent first-derivative formula. The variance form obtained earlier is preferable because the cancellation between the two explicit terms is already encoded in W_r .

7.2 Dimensionless variables

The substitution $y = ru$ removes the mixed scales from the hyperbolic functions. One has

$$G_r(u) = \frac{y \tanh y}{2r^2}, \quad (7.4)$$

and

$$W_r(u) = \frac{y}{4r^4} (\tanh y - y \text{sech}^2 y). \quad (7.5)$$

This explains the natural factor r^{-4} in the two-variable kernel. The dimensionless kernel $K(y, z)$ is independent of r ; all r -dependence has been moved into the tilted measure and the scalar parameter

$$c = 1 - \frac{1}{4r^2}. \quad (7.6)$$

This separation is useful when studying either the small- r or large- r limits.

7.3 Elementary bounds for the hyperbolic factors

For $y \geq 0$,

$$0 \leq \tanh y < 1, \quad 0 < \operatorname{sech}^2 y \leq 1, \quad (7.7)$$

and

$$0 \leq b(y) = y \tanh y \leq y. \quad (7.8)$$

Also,

$$0 \leq a(y) = y \tanh y - y^2 \operatorname{sech}^2 y. \quad (7.9)$$

The last inequality follows from

$$\tanh y \geq y \operatorname{sech}^2 y, \quad (7.10)$$

which is equivalent to $\sinh y \cosh y \geq y$. This elementary hyperbolic inequality follows by differentiation because the left-hand side has derivative $\cosh(2y) \geq 1$ and vanishes at 0.

Thus the diagonal terms of the original kernel are nonnegative:

$$K(y, y) = 2a(y) \geq 0. \quad (7.11)$$

The difficulty occurs off the diagonal, where the square difference subtracts a nontrivial quantity.

7.4 The theta adjoint operator

Start from

$$\Phi = \frac{1}{2} \left(Y'' - \frac{1}{4} Y \right) \quad (7.12)$$

and consider a smooth test function f with sufficient decay. Two integrations by parts give

$$\int_0^\infty f(u) \Phi(u) \cosh(ru) \, du = \frac{1}{2} \int_0^\infty Y(u) \mathcal{L}_r f(u) \cosh(ru) \, du \quad (7.13)$$

$$+ \frac{1}{2} [f'(0)Y(0) - f(0)Y'(0)], \quad (7.14)$$

where

$$\boxed{\mathcal{L}_r f = f'' + 2r \tanh(ru) f' + \left(r^2 - \frac{1}{4} \right) f.} \quad (7.15)$$

When $f(0) = f'(0) = 0$, the boundary term vanishes and the shorter adjoint identity used in intermediate calculations is recovered. In the present paper the full boundary formula is retained whenever a kernel does not vanish at the origin.

7.5 The first-derivative kernel

Introduce dimensionless variables

$$y = ru, \quad z = rv, \quad (7.16)$$

and set

$$a(y) = y \tanh y - y^2 \operatorname{sech}^2 y, \quad b(y) = y \tanh y. \quad (7.17)$$

The exact two-variable kernel is

$$\boxed{K(y, z) = a(y) + a(z) - (b(y) - b(z))^2.} \quad (7.18)$$

Therefore

$$-S'(x) = \frac{1}{2A(r)^2} \iint \mathcal{K}_r(u, v) \Phi(u) \Phi(v) \cosh(ru) \cosh(rv) \, du \, dv, \quad (7.19)$$

where

$$\mathcal{K}_r(u, v) = \frac{1}{4r^4} K(ru, rv). \quad (7.20)$$

An equivalent algebraic form is

$$K(y, z) = b(y) + b(z) + 2b(y)b(z) - y^2 - z^2. \quad (7.21)$$

This decomposition makes the rank-one positive component $2b(y)b(z)$ explicit. The residual term is a signed defect rather than a pointwise-positive kernel.

7.6 Boundary values

At the origin,

$$b(0) = 0, \quad a(0) = 0, \quad (7.22)$$

and therefore

$$K(0, z) = b(z) - z^2. \quad (7.23)$$

Moreover,

$$\partial_y K(0, z) = 0, \quad \partial_y^2 K(0, z) = 4b(z). \quad (7.24)$$

These formulas are the exact boundary data required in the double integration-by-parts calculation.

7.7 Operator identities

Let

$$\mathcal{L} = \partial_y^2 + 2 \tanh y \, \partial_y + c, \quad c = 1 - \frac{1}{4r^2}. \quad (7.25)$$

Then

$$\mathcal{L}b = 2 + cb. \quad (7.26)$$

A direct calculation gives the exact identity

$$\boxed{\mathcal{L}a - \mathcal{L}(b^2) = -4b + c(b - y^2)}. \quad (7.27)$$

Since

$$K(y, z) = a(y) + a(z) - (b(y) - b(z))^2, \quad (7.28)$$

substitution of the preceding identities yields

$$\boxed{\mathcal{L}_y \mathcal{L}_z K(y, z) = 8 + c^2 K(y, z)}. \quad (7.29)$$

Equivalently,

$$(\mathcal{L}_y \mathcal{L}_z - c^2)K(y, z) = 8. \quad (7.30)$$

This is the exact operator relation used below. In particular, the two-variable operator does not reduce K to a finite-rank kernel in general.

7.8 Interpretation of the corrected operator identity

The identity

$$(\mathcal{L}_y \mathcal{L}_z - c^2)K = 8 \quad (7.31)$$

is an exact fourth-order constraint on the kernel. It does not imply pointwise positivity, because K changes sign off the diagonal. Its useful role is instead to couple the interior kernel to the explicit boundary traces. Any positivity argument must retain this coupling rather than replace it by a finite-dimensional Gram approximation.

7.9 Boundary traces as finite-dimensional data

For the kernel K , the trace at the origin is

$$K(0, z) = b(z) - z^2, \quad (7.32)$$

while the first normal derivative vanishes. Hence the first nontrivial trace data occur at order zero and order two. The theta adjoint maps these traces into combinations of $Y(0)$, $Y'(0)$, and integrals of $Yb \cosh(ru)$. This explains why the final obstruction is expected to be finite dimensional.

7.10 A warning about pointwise positivity

The kernel K is positive on the diagonal but not on the whole quadrant. In fact, with z large and y near zero,

$$K(y, z) \sim -z^2 + z \quad (7.33)$$

up to exponentially small corrections. Therefore any argument that replaces the quadratic form by pointwise kernel positivity is incompatible with the exact asymptotics of the kernel.

7.11 What the corrected operator identity does and does not say

The corrected identity is a strong structural constraint on the theta kernel: after the natural second-order operator is applied in both variables, the result is determined by the original kernel and a constant forcing term. Even for $r \geq 1/2$, the term $c^2 K$ in the corrected identity inherits the sign changes of K , so the operator identity alone gives no positivity conclusion.

However, the original quadratic form also contains boundary terms because $K(0, z)$ does not vanish. It is therefore not valid to conclude pointwise or global positivity merely from the positivity of the transformed bulk. The exact boundary contribution must be retained. This is one of the places where a careless integration-by-parts argument would create a false proof.

7.12 The defect decomposition

The kernel can be decomposed as

$$K(y, z) = K_0(y) + K_0(z) + 2b(y)b(z), \quad (7.34)$$

where

$$K_0(y) = b(y) - y^2. \quad (7.35)$$

The rank-one term is positive as a Gram kernel. The defect satisfies

$$K_0(0) = K'_0(0) = 0, \quad (7.36)$$

but

$$K_0(y) < 0 \quad (7.37)$$

for sufficiently large y because $b(y) \sim y$ while y^2 dominates. Thus any positivity statement for the complete quadratic form must come from a global interaction between the rank-one term, the boundary structure, and the theta weight.

7.13 The boundary term for a general test function

The full boundary formula is especially important when $f(0) \neq 0$. If f has sufficient decay at infinity, then

$$\int_0^\infty f \Phi \cosh(ru) \, du = \frac{1}{2} \int_0^\infty Y \mathcal{L}_r f \cosh(ru) \, du + \frac{1}{2} f'(0) Y(0) - \frac{1}{2} f(0) Y'(0). \quad (7.38)$$

Because $Y'(0) = -1/4$, the final term is actually

$$+ \frac{1}{8} f(0). \quad (7.39)$$

This is easy to miss if one imports the homogeneous integration-by-parts formula from a class of test functions that vanish at the origin.

7.14 The kernel on the diagonal

The exact kernel satisfies

$$K(y, y) = 2a(y). \quad (7.40)$$

Since $a(y) \geq 0$, the diagonal is positive. In fact, near the origin,

$$a(y) = \frac{2}{3} y^4 - \frac{8}{15} y^6 + O(y^8), \quad (7.41)$$

while for large y ,

$$a(y) = y + O(y^2 e^{-2y}). \quad (7.42)$$

Thus the negative part of the off-diagonal kernel is not caused by negative diagonal values. It comes entirely from the interaction term $(b(y) - b(z))^2$.

7.15 A symmetric quadratic-form representation

Let $Y = rU$ under the push-forward of P_r to $[0, \infty)$, and denote this push-forward probability measure by ω_r . Thus

$$d\omega_r(y) = \frac{1}{rA(r)} \Phi\left(\frac{y}{r}\right) \cosh(y) \, dy, \quad r > 0. \quad (7.43)$$

Then the kernel identity takes the form

$$\iint K(y, z) \, d\omega_r(y) \, d\omega_r(z) = 2\mathbb{E}_r[a(Y)] - 2\operatorname{Var}_r(b(Y)). \quad (7.44)$$

The first derivative inequality is therefore equivalent to a comparison between the mean of a and the variance of b . This is the dimensionless version of the earlier resolvent identity.

The point of writing it this way is that it separates a one-point quantity from a two-point fluctuation. A proof may therefore come from a Poincare inequality, a convexity estimate, or a direct finite-dimensional representation. The earlier numerical tests show that a naive pointwise Stein bound is not the correct mechanism, but the exact variance formulation remains a useful starting point.

7.16 Operator action on the boundary defect

Let

$$K_0(y) = b(y) - y^2. \quad (7.45)$$

Then

$$\mathcal{L}K_0 = y[(c - 4)\tanh y - cy]. \quad (7.46)$$

For $r \geq 1/2$, one has $c \geq 0$. The expression is then nonpositive for $y > 0$ because $\tanh y < y$ and the coefficient of the linear term is already negative. Thus the operator itself does not turn the defect into a positive kernel. The positivity must come from its interaction with the rank-one component and the boundary contribution.

8 Corrected Operator Constraint for the First Positivity Condition

8.1 Mode decomposition

Write

$$\Phi(u) = \sum_{n=1}^{\infty} \Phi_n(u), \quad (8.1)$$

with

$$\Phi_n(u) = \pi n^2 e^{5u/2} (2\pi n^2 e^{2u} - 3) e^{-\pi n^2 e^{2u}}. \quad (8.2)$$

Substitution into the double-kernel representation gives

$$-S'(x) = \frac{1}{2A(r)^2} \sum_{m,n \geq 1} Q_{mn}(r), \quad (8.3)$$

where

$$Q_{mn}(r) = \iint \Phi_m(u) \Phi_n(v) \cosh(ru) \cosh(rv) \mathcal{K}_r(u, v) \, du \, dv. \quad (8.4)$$

The matrix $Q(r) = (Q_{mn}(r))_{m,n \geq 1}$ is symmetric. The individual entries are not required to be nonnegative; the relevant object is the complete quadratic form generated by the full theta sum.

8.2 The corrected interior identity

The natural dimensionless operator is

$$\mathcal{L} = \partial_y^2 + 2 \tanh y \partial_y + c, \quad c = 1 - \frac{1}{4r^2}. \quad (8.5)$$

For

$$b(y) = y \tanh y, \quad a(y) = y \tanh y - y^2 \operatorname{sech}^2 y, \quad (8.6)$$

we have

$$\mathcal{L}b = 2 + cb \quad (8.7)$$

and

$$\mathcal{L}a - \mathcal{L}(b^2) = -4b + c(b - y^2). \quad (8.8)$$

Consequently the exact kernel

$$K(y, z) = a(y) + a(z) - (b(y) - b(z))^2 \quad (8.9)$$

satisfies

$$\boxed{\mathcal{L}_y \mathcal{L}_z K(y, z) = 8 + c^2 K(y, z).} \quad (8.10)$$

Equivalently,

$$\boxed{(\mathcal{L}_y \mathcal{L}_z - c^2)K(y, z) = 8.} \quad (8.11)$$

This is an exact differential constraint. It is not a finite-rank representation of K , nor does it imply a positive-semidefinite matrix factorization.

8.3 What the operator identity actually gives

Applying the corrected operator relation in a double integration-by-parts calculation leaves the kernel K itself together with explicit boundary terms. Thus the interior and boundary contributions cannot be separated into a proved positive bulk matrix plus a residual finite-dimensional defect on the basis of the differential identity alone.

The exact boundary data are

$$K(0, z) = b(z) - z^2, \quad \partial_y K(0, z) = 0, \quad \partial_y^2 K(0, z) = 4b(z). \quad (8.12)$$

These traces must be retained in any subsequent positivity argument. In particular, since $K(0, z) < 0$ for sufficiently large z , no proof can proceed by asserting pointwise nonnegativity of the kernel or by discarding its boundary contribution.

8.4 The signed defect and the positive rank-one component

The algebraic decomposition

$$K(y, z) = K_0(y) + K_0(z) + 2b(y)b(z), \quad K_0(y) = b(y) - y^2, \quad (8.13)$$

is exact. Because $b(y) = y \tanh y \geq 0$ on $[0, \infty)$, the term

$$2b(y)b(z) \quad (8.14)$$

is a positive rank-one Gram kernel. This fact is useful, but it does not imply positivity of the full kernel: the one-variable defect K_0 is negative for sufficiently large arguments.

Thus the first-order problem is a global quadratic-form inequality in which the positive rank-one component, the signed defect, the theta weight, and the boundary traces interact. No positive finite-dimensional completion has been established here.

8.5 Finite mode truncations and passage to the limit

For

$$\Phi^{(M)}(u) = \sum_{n=1}^M \Phi_n(u), \quad A_M(r) = \int_0^\infty \Phi^{(M)}(u) \cosh(ru) du, \quad (8.15)$$

the same kernel identity holds for the truncated quadratic form. Local uniform convergence of the theta series and its differentiated forms, together with the superexponential decay in u , justifies passage to the infinite sum in the integral identities used here. What passes to the limit is the exact operator constraint, not a finite-rank compression.

8.6 Corrected Schur-complement perspective

A Schur-complement organization remains possible as a bookkeeping device, but it must begin from the full boundary-coupled quadratic form. The present identity does not supply a positive-semidefinite bulk block whose Schur complement could then be tested. Any such block decomposition would require an additional exact derivation and a separate proof of its sign.

Accordingly, the precise first-order question is

$$\boxed{-S'(x) \geq 0 \quad (x > 0)}, \quad (8.16)$$

with the left-hand side represented by the exact theta quadratic form above. The corrected operator identity constrains that form but does not settle its sign.

8.7 Status of the first-order route

The corrected calculation removes the earlier finite-rank claim completely. What remains is a mathematically well-defined boundary-coupled positivity problem. Solving that first-order inequality would be a genuine new result, but it would still fall short of the all-orders Hankel positivity required for the Riemann Hypothesis.

9 Complete Hankel Positivity: The Decisive Remaining Problem

9.1 Euler-operator hierarchy

Let

$$R(r) = r \frac{A'(r)}{A(r)}. \quad (9.1)$$

Then

$$S(r^2) = \frac{R(r)}{2r^2}. \quad (9.2)$$

With the Euler operator

$$D = r \frac{d}{dr}, \quad (9.3)$$

one obtains, for $n \geq 0$,

$$\boxed{x^{n+1} S^{(n)}(x) = \frac{1}{2^{n+1}} \prod_{k=1}^n (D - 2k) R(r), \quad x = r^2.} \quad (9.4)$$

Thus complete monotonicity of S is equivalent to the hierarchy

$$(-1)^n \prod_{k=1}^n (D - 2k) R(r) \geq 0. \quad (9.5)$$

The formula is exact and gives a differential-operator version of the moment hierarchy.

9.2 The all-orders Hankel problem

The Stieltjes condition requires more than complete monotonicity. One needs a positive measure representation with the correct spectral support. Equivalently one needs

$$H_N \succeq 0, \quad \tilde{H}_N \succeq 0 \quad (9.6)$$

for every N . The first-order operator reduction does not automatically iterate to all orders because higher derivatives generate increasingly complicated covariance tensors. A successful theory must find a mechanism that propagates positivity through these higher interactions.

One possible formulation is to seek a direct positive kernel representation for the two divided differences

$$\mathcal{H}_0(u, v), \quad \mathcal{H}_1(u, v), \quad (9.7)$$

induced by the theta geometry. Another is to identify a positive representation for the reciprocal spectral measure itself. These are equivalent targets at different levels of resolution.

9.3 The operator hierarchy behind higher derivatives

The first derivative is governed by one application of the logarithmic differentiation and one covariance operation. At higher order, repeated differentiation generates iterated covariance tensors. The Euler-operator formula provides a way to organize these terms without expanding all of them individually.

Let

$$R(r) = r \frac{A'(r)}{A(r)}. \quad (9.8)$$

Since $x = r^2$, the differential operator $x \frac{d}{dx}$ becomes $\frac{1}{2}D$. This is why the factors $(D - 2k)$ appear naturally in the hierarchy. The operator formula is exact and independent of any probabilistic interpretation.

One may therefore view the all-orders problem as a positivity problem for the sequence of operators

$$\mathcal{D}_n = \prod_{k=1}^n (D - 2k). \quad (9.9)$$

The desired signs are

$$(-1)^n \mathcal{D}_n R(r) \geq 0. \quad (9.10)$$

For a Stieltjes function these inequalities are automatic, but in the theta representation they become nontrivial differential inequalities for the logarithm of a positive cosh transform.

9.4 A possible bridge to complete Bernstein structure

The identity

$$xS(x) = \frac{1}{2}R(r) \quad (9.11)$$

suggests considering the product $xS(x)$ rather than S alone. In the classical Stieltjes/Bernstein correspondence, a Stieltjes function becomes especially rigid after multiplication by x . In the present framework this quantity is directly the logarithmic derivative of the theta integral with respect to the Euler scale.

No complete Bernstein theorem is asserted here. The point is to identify a possible intermediate target: if one could establish that $xS(x)$ has the appropriate Bernstein structure directly from the theta measure, then the Stieltjes property would follow from the standard algebraic correspondence. The required proof would still have to control the parameter dependence of the tilt.

9.5 The divided-difference route

The moment characterization also suggests bypassing derivatives altogether. The two kernels

$$\mathcal{H}_0(u, v) = \frac{uT(u) - vT(v)}{u - v}, \quad \mathcal{H}_1(u, v) = \frac{T(u) - T(v)}{u - v} \quad (9.12)$$

are positive semidefinite precisely when the corresponding Hankel families are positive. A theta representation of these divided differences would therefore be enough. Such a representation would have to preserve the two-variable positivity after integrating out the theta variable.

This route has an advantage over the Euler hierarchy: the half-line moment problem is built into the kernel from the start. Its disadvantage is that the corrected two-variable differential identity does not make positivity immediate. A further structural representation is still required.

9.6 The fixed positive measure target

Suppose one could construct, from the theta representation alone, a positive measure σ on $[0, \infty)$ such that

$$\mu_n = \int_0^\infty t^n d\sigma(t) \quad (9.13)$$

for every n . Then the Stieltjes theorem would immediately give both Hankel hierarchies, and the exact spectral pole comparison would imply RH. Thus the problem can be stated in one sentence:

Construct a theta-induced positive measure realizing all reciprocal moments.

(9.14)

The theta probability measures P_r are not yet that fixed measure; they depend on r . Understanding how to remove this parameter dependence is the central remaining issue.

9.7 Relation to the reciprocal moment coefficients

The Taylor coefficients of S at 0 and the cumulants of F carry the same information. In particular,

$$S(x) = \mu_0 - \mu_1 x + \mu_2 x^2 - \cdots \quad (9.15)$$

inside the disk of convergence. Thus any all-orders inequality on the derivatives of S can be translated into an inequality among cumulants, and conversely. The Hankel formulation has the advantage that it packages these coefficient inequalities into a single quadratic-form statement at each degree.

9.8 Why the Euler hierarchy is not a substitute for Hankel positivity

The inequalities

$$(-1)^n S^{(n)}(x) \geq 0 \quad (9.16)$$

for all n express complete monotonicity. They do not, by themselves, identify the support of a representing measure. The Hankel conditions retain additional information because they compare different moments simultaneously. This is why the higher-order section is formulated around the Stieltjes hierarchy rather than only the derivative hierarchy.

9.9 A hierarchy of possible positive kernels

For each N , positivity of H_N is equivalent to positivity of the corresponding Hankel quadratic form; no corrected kernel representation is asserted here.

$$(u, v) \mapsto \sum_{i,j=0}^N \mu_{i+j} u^i v^j \quad (9.17)$$

on polynomial test functions. Likewise the shifted family corresponds to multiplication by the spectral coordinate. A theta proof that constructs these kernels directly would establish all finite Hankel conditions simultaneously. The first-order kernel studied earlier should then be interpreted as the first member of a larger family rather than as an isolated object.

9.10 Why the inverse-Laplace test is not enough

A numerical inverse-Laplace calculation can test positivity of a candidate density, and numerical experiments can test the signs of several derivatives of S . Such evidence is useful for locating possible obstructions, but it does not establish the Stieltjes property. The correct analytic requirement is positivity of a measure together with support on the nonnegative axis, or an equivalent complete positivity statement for the Hankel forms.

9.11 Numerical diagnostics and their exact role

Finite computations were used during the development of the framework to test signs and detect algebraic mistakes. For the full theta representation, direct numerical differentiation gives positive values for the alternating derivatives of S over a wide range of positive arguments. Similar calculations of inverse-Laplace candidates give positive numerical values at sampled points.

These computations are deliberately not promoted to theorems. Their role is diagnostic: they test whether an alleged sign identity is compatible with the exact formula and whether a proposed obstruction is likely to be genuine. In particular, the mode decomposition shows that individual cross-mode terms need not be positive even when the total quadratic form is positive. This is an important warning against termwise positivity arguments.

9.12 A precise standard for completion

A mathematically complete proof of RH through the present framework would require one additional theorem, stated entirely on the theta side, which establishes positivity of the fixed reciprocal moment functional. Once that theorem exists, the spectral conclusion is immediate from the exact equivalence in Section 4.

This separation is useful for future work. It prevents the proof from being blurred across several representations and makes every new lemma testable at the level where it is introduced. In particular, a successful higher-order argument must preserve all boundary contributions and must not silently replace the parameter-dependent family P_r by a parameter-independent measure.

9.13 No hidden conclusion

The paper therefore stops at the exact reduction. It does not assert that the theta kernel already proves RH. The point is not to weaken the goal but to isolate it cleanly: once the all-orders positive representation is supplied, the remainder of the argument is formal and exact. Conversely, any proposed proof must explain the boundary defect and the parameter dependence of the tilted measure; neither can be discarded without changing the problem.

10 Consequences and Research Program

10.1 A hierarchy of equivalent targets

The framework gives several mathematically equivalent targets. The most spectral is positivity of the reciprocal poles. The most algebraic is positivity of the two Hankel families at every finite order. The most measure-theoretic is existence of a positive Stieltjes measure. The most analytic is a positive Stieltjes representation for S . The most geometric is a theta-induced positive structure that survives the logarithmic derivative.

The chain

$$\rho = i\gamma \implies q_\rho > 0 \implies (\mu_n)_{n \geq 0} \text{ is a Stieltjes moment sequence} \implies H_N, \tilde{H}_N \succeq 0 \quad (10.1)$$

is elementary. The reverse implication is equally exact once a positive measure is established. What is difficult is deriving the positivity from the theta geometry without importing the desired zero location.

10.2 A concrete sequence of next steps

The first priority is to complete the boundary-coupled quadratic-form analysis for the first derivative. The corrected operator identity constrains the interior term, but it does not produce a finite-rank bulk block. Proving the resulting sign inequality for all $r > 0$ would provide a nontrivial unconditional first-order theorem, even though it would not by itself imply RH.

The second priority is to search for an operator identity that lifts the first-order differential relation to arbitrary divided differences. The ideal outcome would be a family of positive semidefinite kernels $\mathcal{K}_N$ whose moments are exactly the Hankel forms. Such a construction would bypass the growth in complexity of direct higher derivatives.

The third priority is to understand the tilted family (P_r) as a path in a positive cone. If

the logarithmic derivative of its moment generating function has a total-positivity or variation-diminishing property, that property could potentially propagate to the reciprocal moment hierarchy. The exact Euler-operator identities above suggest that an operator-theoretic formulation is more natural than direct coefficient estimates.

A fourth possibility is to characterize the desired fixed measure through the Stieltjes inversion formula. For a candidate Stieltjes function, the absolutely continuous density is recovered from the boundary values on the negative axis, while atoms arise from poles. In the present problem the continuation of S is meromorphic and real on the negative axis away from poles. The meromorphic continuation leaves no room for an additional continuous component across an open interval of the cut. The residues at the spectral poles are already the positive multiplicities m_ρ ; the unresolved issue is instead the location of those poles. The support condition requires $t_\rho = -\rho^2 \geq 0$, which is precisely the critical-line condition.

10.3 What a completed proof would look like

A successful completion would contain a theorem of the following form.

There exists a theta-geometric positive functional $\mathcal{P}$ on polynomials such that $\mathcal{P}(t^n) = \mu_n$ for all $n \geq 0$.

The Stieltjes moment theorem would then imply positivity of both Hankel families. The exact spectral equivalence would force all centered zeros onto the imaginary axis. The new content of the proof would therefore reside entirely in the construction of $\mathcal{P}$ and its positivity.

10.4 Scope and limitations

The present manuscript is designed as a structural paper. Its purpose is to put the analytic, spectral, and moment sides of the problem into one coherent framework and to derive the corrected theta differential-kernel identity exactly. It deliberately does not turn numerical evidence into a theorem, and it does not use a stronger statement at one stage than has been proved at another. This discipline is particularly important for a problem in which a seemingly small sign error changes the logical status of the conclusion.

11 Conclusion

We have reorganized the reciprocal spectral problem of the Riemann xi-function around a theta-geometric measure. The folding $w = -z^2$ turns the critical line into the positive real axis, the logarithmic derivative produces a canonical resolvent, and reciprocal poles produce an exact moment sequence. The corresponding Stieltjes moment criterion is equivalent to the Riemann Hypothesis through the two Hankel families.

On the theta side, the centered xi-function is an exact cosh transform of the positive density Φ . The induced tilted measure turns the resolvent into an expectation of the explicit function $u \tanh(ru)$. Differentiation yields a signed variance structure, and the natural theta adjoint converts the associated two-variable kernel into the exact two-variable differential identity above. The remaining terms are boundary contributions and a defect that cannot be discarded.

The main conclusion is therefore structural rather than final: the spectral positivity problem has been reduced to a precise theta-geometric positivity problem. The all-orders Hankel hierarchy remains the decisive step. Any future proof must provide that positivity, either directly at the

level of the moment functional or through an equivalent positive kernel construction. Once that step is established, the exact spectral reduction already developed here supplies the implication to the Riemann Hypothesis.

A Product and convergence details

The folded product used in Section 3 is justified by the reciprocal-square summability of the centered zeros. The relevant estimate is

$$\sum_{\rho \in \mathcal{Z}_c} \frac{m_\rho}{|\rho|^2} < \infty. \quad (\text{A.1})$$

For compact subsets of the w -plane away from the zeros, the product factors satisfy

$$\left| \frac{w}{\rho^2} \right| \leq C_K |\rho|^{-2} \quad (\text{A.2})$$

for ρ sufficiently large and a compact-dependent constant C_K . Uniform convergence follows from the Weierstrass product criterion. Logarithmic differentiation is then legitimate locally uniformly away from the zero set.

The series for S satisfies the same convergence statement. On any compact subset of the x -plane away from the poles, the summands admit the bound

$$\left| \frac{m_\rho}{x - \rho^2} \right| \leq C_K \frac{m_\rho}{|\rho|^2} \quad (\text{A.3})$$

for sufficiently large $|\rho|$. This supplies normal convergence and hence meromorphicity of the resolvent.

For the reciprocal generating series, one has

$$|q_\rho| = |\rho|^{-2}, \quad (\text{A.4})$$

so $q_\rho \rightarrow 0$. The power series defining T therefore has positive radius of convergence, with the exact radius determined by the largest reciprocal spectral modulus.

B Theta differential identities

We record the termwise calculation behind

$$\Phi = \frac{1}{2}(Y'' - Y/4). \quad (\text{B.1})$$

For one theta mode,

$$Y_n(u) = e^{u/2} e^{-\pi n^2 e^{2u}}. \quad (\text{B.2})$$

Then

$$Y'_n(u) = Y_n(u) \left(\frac{1}{2} - 2\pi n^2 e^{2u} \right), \quad (\text{B.3})$$

$$Y''_n(u) = Y_n(u) \left(\frac{1}{4} - 10\pi n^2 e^{2u} + 4\pi^2 n^4 e^{4u} \right). \quad (\text{B.4})$$

Therefore

$$\frac{1}{2} \left(Y_n'' - \frac{1}{4} Y_n \right) = \pi n^2 e^{5u/2} \left(2\pi n^2 e^{2u} - 3 \right) e^{-\pi n^2 e^{2u}}, \quad (\text{B.5})$$

which is the stated Φ_n .

The theta symmetry used in the boundary calculation is inherited from the standard theta transformation after the substitution $e^{2u} \mapsto e^{-2u}$. Differentiation at the origin gives

$$Y'(0) = -\frac{1}{4}. \quad (\text{B.6})$$

No numerical approximation is used in this identity.

C Mellin calculation

Let

$$M(s) = \int_{-\infty}^{\infty} Y(u) e^{su} du. \quad (\text{C.1})$$

For $\Re s > 1/2$, substituting $t = \pi n^2 e^{2u}$ mode by mode gives the gamma factor and the zeta sum. The change of variables is

$$e^u = \left(\frac{t}{\pi n^2} \right)^{1/2}, \quad du = \frac{dt}{2t}, \quad (\text{C.2})$$

and therefore

$$\int_{-\infty}^{\infty} Y_n(u) e^{su} du = \frac{1}{2} \pi^{-s/2-1/4} n^{-s-1/2} \Gamma\left(\frac{s}{2} + \frac{1}{4}\right). \quad (\text{C.3})$$

Summing over n gives

$$M(s) = \frac{1}{2} \pi^{-s/2-1/4} \Gamma\left(\frac{s}{2} + \frac{1}{4}\right) \zeta\left(s + \frac{1}{2}\right). \quad (\text{C.4})$$

The half-line formula is obtained by writing

$$M(s) = \int_0^{\infty} Y(u) e^{su} du + \int_0^{\infty} Y(-u) e^{-su} du \quad (\text{C.5})$$

and using

$$Y(-u) = Y(u) + \sinh(u/2). \quad (\text{C.6})$$

Taking even parts produces the stated expression for $J(s)$.

D Cumulant identities

Because F is even, its logarithm has an even expansion near the origin,

$$\log \frac{F(r)}{F(0)} = \sum_{n \geq 1} \frac{\kappa_{2n}}{(2n)!} r^{2n}. \quad (\text{D.1})$$

On the other hand,

$$\frac{F'(r)}{F(r)} = 2rS(r^2). \quad (\text{D.2})$$

Differentiating the logarithmic series gives

$$\frac{F'(r)}{F(r)} = \sum_{n \geq 1} \frac{\kappa_{2n}}{(2n-1)!} r^{2n-1}. \quad (\text{D.3})$$

Comparing with

$$2rS(r^2) = 2 \sum_{m \geq 0} (-1)^m \mu_m r^{2m+1} \quad (\text{D.4})$$

in the reciprocal expansion yields

$$\mu_{n-1} = \frac{(-1)^{n-1} \kappa_{2n}}{2(2n-1)!}. \quad (\text{D.5})$$

The first two examples are

$$\mu_0 = \frac{\kappa_2}{2}, \quad \mu_1 = -\frac{\kappa_4}{12}, \quad (\text{D.6})$$

and

$$\mu_2 = \frac{\kappa_6}{240}. \quad (\text{D.7})$$

Substitution into $\mu_0 \mu_2 - \mu_1^2$ gives

$$\det \begin{pmatrix} \mu_0 & \mu_1 \\ \mu_1 & \mu_2 \end{pmatrix} = \frac{3\kappa_2 \kappa_6 - 10\kappa_4^2}{1440}. \quad (\text{D.8})$$

E Boundary integration formula

For later reference we record the full identity in a form that makes all signs visible. Let f be smooth and rapidly decaying. Starting from

$$\int_0^\infty f \Phi \cosh(ru) \, du = \frac{1}{2} \int_0^\infty f(Y'' - Y/4) \cosh(ru) \, du, \quad (\text{E.1})$$

integrate the Y'' term once:

$$\int_0^\infty f Y'' \cosh(ru) \, du = [f Y' \cosh(ru)]_0^\infty - \int_0^\infty Y' (f' \cosh(ru) + r f \sinh(ru)) \, du. \quad (\text{E.2})$$

A second integration by parts produces

$$\int_0^\infty f Y'' \cosh(ru) \, du = f'(0)Y(0) - f(0)Y'(0) \quad (\text{E.3})$$

$$+ \int_0^\infty Y(f'' + 2r \tanh(ru)f' + r^2 f) \cosh(ru) \, du. \quad (\text{E.4})$$

Subtracting the $Y/4$ term yields exactly

$$\int_0^\infty f \Phi \cosh(ru) \, du = \frac{1}{2} \int_0^\infty Y \mathcal{L}_r f \cosh(ru) \, du + \frac{1}{2} [f'(0)Y(0) - f(0)Y'(0)]. \quad (\text{E.5})$$

No boundary term is omitted.

F Algebraic checks for the corrected kernel identity

For

$$b(y) = y \tanh y, \quad (\text{F.1})$$

one has

$$b'(y) = \tanh y + y \operatorname{sech}^2 y, \quad (\text{F.2})$$

and

$$b''(y) = 2 \operatorname{sech}^2 y - 2y \operatorname{sech}^2 y \tanh y. \quad (\text{F.3})$$

A direct substitution into

$$\mathcal{L}b = b'' + 2 \tanh y b' + cb \quad (\text{F.4})$$

gives

$$\mathcal{L}b = 2 + cb. \quad (\text{F.5})$$

Next,

$$a(y) = b(y) - y^2 \operatorname{sech}^2 y. \quad (\text{F.6})$$

Direct differentiation gives

$$\mathcal{L}a - \mathcal{L}(b^2) = -4b + c(b - y^2). \quad (\text{F.7})$$

Using

$$K(y, z) = a(y) + a(z) - b(y)^2 + 2b(y)b(z) - b(z)^2 \quad (\text{F.8})$$

and the two identities above gives

$$\mathcal{L}_y \mathcal{L}_z K = c \left[-4b(y) + c(b(y) - y^2) \right] + c \left[-4b(z) + c(b(z) - z^2) \right] \quad (\text{F.9})$$

$$+ 2(2 + cb(y))(2 + cb(z)). \quad (\text{F.10})$$

The terms linear in c cancel, leaving

$$8 + c^2(b(y) + b(z) - y^2 - z^2 + 2b(y)b(z)). \quad (\text{F.11})$$

Finally,

$$K(y, z) = b(y) + b(z) - y^2 - z^2 + 2b(y)b(z), \quad (\text{F.12})$$

by $\tanh^2 y + \operatorname{sech}^2 y = 1$, and hence

$$\boxed{\mathcal{L}_y \mathcal{L}_z K(y, z) = 8 + c^2 K(y, z).} \quad (\text{F.13})$$

Thus the corrected identity is not a finite-rank expansion. The appendix is included to make the replacement calculation independently auditable.

G A compact statement of the remaining obstruction

The exact first-order analysis can be summarized as follows. Let $\mathcal{Q}_r$ denote the quadratic functional corresponding to the kernel K . Then

$$\mathcal{Q}_r = \mathcal{Q}_r^{\text{int}} + \mathcal{Q}_r^{\text{bdry}}, \quad (\text{G.1})$$

where the interior term is constrained by

$$(\mathcal{L}_y \mathcal{L}_z - c^2)K(y, z) = 8, \quad (\text{G.2})$$

and the boundary term depends on the explicit traces of K and on $Y(0), Y'(0)$. The corrected relation does not supply a positive matrix factorization by itself.

This formulation also explains why attempts to prove positivity by declaring the original two-variable kernel nonnegative cannot succeed: $K(0, z) = b(z) - z^2$ is negative for large z . The relevant positivity, if true, must be a global quadratic-form statement after integration against the theta weight.

G.1 The role of the present paper

The contribution of the framework is therefore to make the analytic interface precise. The spectral construction supplies an exact target; the theta representation supplies a positive continuous geometry; the tilt converts the logarithmic derivative into a probabilistic observable; and the differential operator exposes an exact interior relation coupled to boundary data. The remaining problem is concrete enough to support further work without reopening the already established spectral reduction.

G.2 Reproducibility of the algebra

All nontrivial transformations used in the manuscript reduce to the displayed differential, product, and integral identities. In particular, the corrected kernel calculation in Section 8 can be checked by direct differentiation of $b(y) = y \tanh y$ and substitution into the operator $\mathcal{L}$. The appendix gives the intermediate identities so that the sign of every contribution can be audited independently of the surrounding discussion.

G.3 A final logical summary

There are three layers in the present construction. The first is spectral and exact: folding, reciprocal variables, resolvent poles, and the Stieltjes-Hankel equivalence. The second is geometric and exact: the theta density, its tilted measures, the adjoint operator, and the corrected two-variable differential identity. The third is the still-open positivity step: an all-orders positive representation of the reciprocal moments. Keeping these layers separate is not a matter of presentation alone; it is what prevents an exact reduction from being mistaken for a completed proof.