Riemann

Heat growth, prime coherence, and the remaining cancellation estimate

Komorebi P.

READER EDITION 1.1 / 8 OCTOBER 2026

A self-contained account of the packet approach to the Riemann hypothesis. Heat growth exposes the zero geometry. Signed arithmetic cancellation remains the decisive estimate.

EXACT IDENTITIES / FINITE EVIDENCE / OPEN CONJECTURE

Contents

1 The conjecture and the research boundary **2**

2 The zeta function and its completed symmetry **2**
2.1 What completion changes 3

3 A Gaussian packet which preserves the zero obstruction **3**
3.1 A usable convergence estimate 5
3.2 Why a nonzero residue cannot be hidden by cancellation 6

4 The exact prime formula and removal of explicit channels **7**
4.1 Following the Gaussian constants 8
4.2 Why the arithmetic channel contains prime powers 9
4.3 From unit bands to a holomorphic transform 11

5 Stable resolution in the physical Gaussian metric **12**
5.1 The metric and a concrete conditioning example 12
5.2 Deriving the coherent coefficients 13

6 The classical Möbius packet and exact prime cubes **15**
6.1 A finite formula and the active scale 15
6.2 Two complete small cubes 17

7 A barrier for discarding the remaining signs **18**
7.1 Which quantity is being bounded 18
7.2 A worked support count 19
7.3 Tracking the proof's scale losses 21

8 Other mechanisms which preserve the obstruction **22**

9 The exact open conjecture in arithmetic coordinates **22**
9.1 What a closing estimate would have to say 23

10 Reading map and notation **24**

11 Evidence, certification, and continuation **25**

1 The conjecture and the research boundary

The Riemann hypothesis asks whether every nontrivial zero of the Riemann zeta function has real part one half. Its statement is short, but the arithmetic condition it imposes is global. A finite list of zeros on the critical line does not settle it. Nor does a convenient identity between primes and zeros. An identity transfers the problem. A proof must control the transferred object for every scale.

This monograph records the Komorebi approach through that distinction. The main objects are a Gaussian packet over the zeros, its exact expression in terms of prime powers, a stable Hermite resolution of the prime packet, and an equivalent classical sequence built from the Möbius function. The analytic changes of representation can be made explicit. They preserve the obstruction carried by a hypothetical zero away from the critical line. The unresolved step is a signed arithmetic estimate.

The most concrete contribution presented here is a barrier theorem for one proposed route. An alternating cube of prime divisions reorganizes a Möbius sum exactly. If the remaining signs are then discarded, its full absolute majorant is too large to meet the Riemann hypothesis criterion. We give the proof with the squarefree support retained, so that the obstruction concerns the actual triangle inequality for the arithmetic sum.

The barrier does not imply that the signed sum is large. It says that a sufficiently small signed sum needs a cancellation theorem which the cube identity itself does not provide. This is a useful boundary. It prevents improvements in bookkeeping, resolution, or transport from being mistaken for a solution of the arithmetic problem.

The saved Komorebi work also reports focused Lean verification of several algebraic obstruction results and metric bookkeeping statements. Those historical certificates are described at the end. The analytic proofs and the revised barrier proof in this monograph are manuscript proofs. No new Lean verification is claimed.

2 The zeta function and its completed symmetry

For a complex number s with real part greater than one, the zeta function is

$$\zeta(s) = \sum_{n=1}^{\infty} n^{-s} = \prod_{p} (1 - p^{-s})^{-1}, \tag{2.1}$$

where n ranges over positive integers and p ranges over primes. Analytic continuation extends it to a meromorphic function on the complex plane with a simple pole at $s = 1$. The completed function

$$\xi(s) = \frac{1}{2} s(s-1) \pi^{-s/2} \Gamma(s/2) \zeta(s) \tag{2.2}$$

is entire and satisfies $\xi(s) = \xi(1-s)$. Here Γ is the gamma function. These are classical facts, with the normalization in (2.2) fixed throughout this article [1].

Let $\mathcal{Z}$ be the multiset of nontrivial zeros, with every zero repeated according to its multiplicity. They lie in the strip $0 < \Re s < 1$. Conjugation and the functional equation imply that a zero ρ is accompanied by $\overline{\rho}$, $1 - \rho$, and $1 - \overline{\rho}$, allowing coincidences on the critical line [2].

Conjecture 2.1 (Riemann hypothesis)**.** *Every $\rho \in \mathcal{Z}$ satisfies $\Re \rho = 1/2$.*

Write $\rho = \beta + i\gamma$, with real numbers β and γ, and define its centered coordinate

$$z_\rho = \rho - \frac{1}{2}. \tag{2.3}$$

The symmetries become $z \mapsto -z$ and $z \mapsto \bar{z}$. Under the conjecture, every centered coordinate is purely imaginary. Away from the line, the real part of z becomes an exponential growth parameter in the packet below.

2.1 What completion changes

The Dirichlet series in (2.1) begins in a half plane where its terms can be summed absolutely. The Euler product there records unique factorization. Expanding each factor as a geometric series chooses a power of every prime, and each positive integer appears once. This gives a direct arithmetic meaning to the function before continuation is needed.

Continuation and completion do different jobs. Continuation lets the same analytic function be studied outside its first convergence region. Completion packages the factors needed for reflection about the critical line. The factor $s(s-1)$ removes the apparent exceptional behavior at zero and one, while the gamma factor balances the trivial zeros at the negative even integers. The resulting entire function has the nontrivial zeros as its zero set. This is why a plot or packet built from ξ can focus on the conjecture without also carrying the pole of ζ.

The reflection is $s \mapsto 1 - s$, not ordinary complex conjugation. These two operations commute, and together they generate the symmetric zero pattern. For example, a hypothetical zero

$$\rho = \frac{1}{2} + \delta + i\gamma, \qquad \delta > 0,$$

would bring centered coordinates

$$\delta + i\gamma, \quad \delta - i\gamma, \quad -\delta + i\gamma, \quad -\delta - i\gamma.$$

The quartet may collapse when $\delta = 0$ or $\gamma = 0$. We sum the actual multiset, so coincidences are handled by multiplicity rather than by pretending that four distinct zeros are always present.

On the critical line, conjugation and reflection identify the same point: $\overline{1/2 + i\gamma} = 1 - (1/2 + i\gamma)$. Therefore $\xi(1/2 + i\gamma)$ is real. Real values on that line explain why one can track sign changes in a one dimensional plot. They do not explain why every zero must lie there. A real analytic function restricted to one line can have additional zeros elsewhere in the complex plane.

The centered coordinate separates this geometry from the notation of the completed function. Its imaginary part is an oscillation frequency. Its real part is a growth or decay rate. In the zero packet, increasing the translation u tests precisely that rate. The heat parameter suppresses large frequencies, which makes the sum convergent, while preserving every mode with a nonzero coefficient.

3 A Gaussian packet which preserves the zero obstruction

Definition 3.1 (Zero heat packet)**.** For a real heat parameter $t > 0$ and a complex translation u, define

$$\Theta_t(u) = \sum_{\rho \in \mathcal{Z}} \exp\left(t z_\rho^2 + u z_\rho\right). \tag{3.1}$$

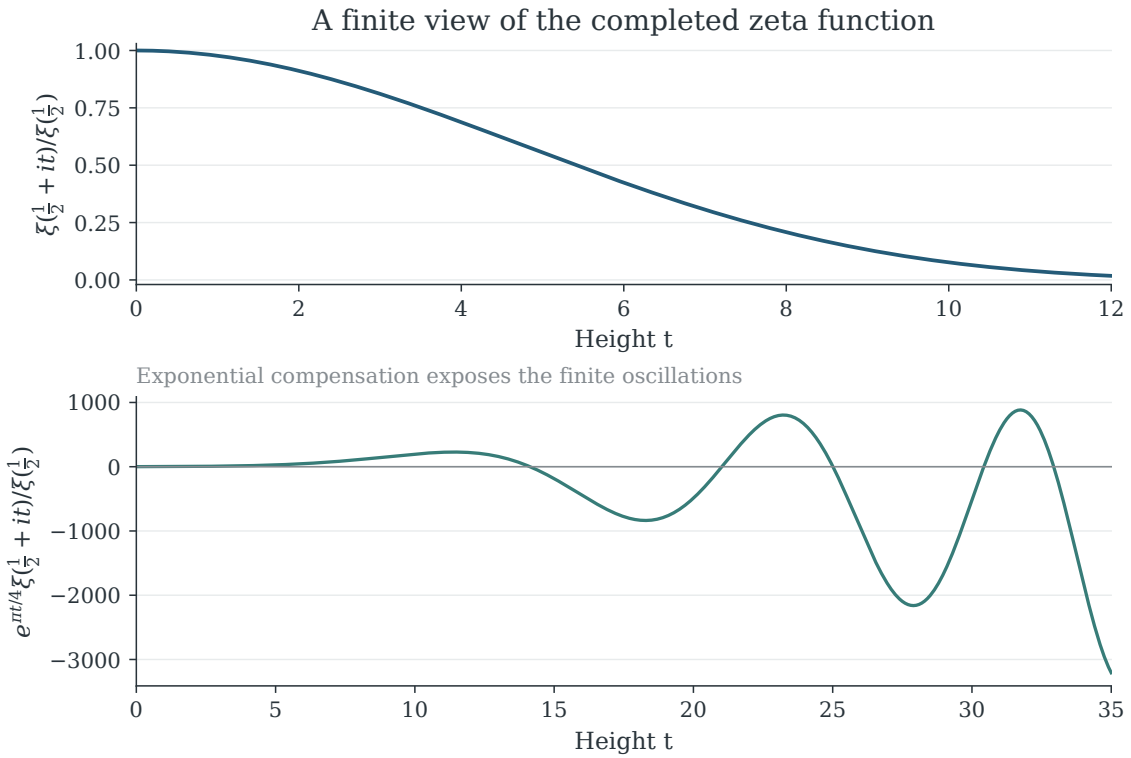


Figure 1: A finite evaluation of the completed function on the critical line. The lower panel multiplies the normalized value by $e^{\pi t/4}$, where t is the plotted height. The symmetry and normalization are exact classical facts. The plotted interval is a finite numerical illustration and supplies no conclusion about all zeros.

This normalization agrees with the recovered Komorebi heat kernel manuscript [6]. It is an exact sum over the actual zeros. A hypothetical shifted mode used in a figure is a separate diagnostic model.

The series converges normally on compact sets of translations and on compact intervals of positive heat times. Indeed, for real u,

$$\left|e^{tz_\rho^2+uz_\rho}\right| \le e^{t/4+|u|/2}e^{-t\gamma^2}. \tag{3.2}$$

For complex u, there is an additional exponential factor linear in $|\gamma|$ on each compact set. Gaussian decay still dominates. The classical zero counting estimate $N(Y) = O(Y \log Y)$, where $N(Y)$ counts zeros with $0 < \Im\rho \le Y$, gives summability [5, Section 2.4, equation (2.11)]. The same domination permits differentiation. Consequently

$$\partial_t\Theta_t = \partial_u^2\Theta_t. \tag{3.3}$$

The zero symmetries make $\Theta_t(u)$ real and even when u is real.

3.1 A usable convergence estimate

It helps to make the normal convergence argument quantitative enough to see which operations it permits. Write $z = x + i\gamma$, where $|x| < 1/2$, and $u = a + ib$. Then

$$\Re(tz^2 + uz) = t(x^2 - \gamma^2) + ax - b\gamma.$$

If $t \in [t_0, T]$ and $|a|, |b| \le B$, this is at most

$$T/4 + B/2 - t_0\gamma^2 + B|\gamma| \le T/4 + B/2 + \frac{B^2}{2t_0} - \frac{t_0}{2}\gamma^2.$$

The last step is the elementary inequality $B|\gamma| \le (t_0/2)\gamma^2 + B^2/(2t_0)$. Thus a compact set of heat times and translations has one Gaussian majorant. A finite number of derivatives only inserts a polynomial in z. Such a polynomial is still summable against this Gaussian.

For completeness, divide the positive ordinates into intervals $[j, j+1)$. The number in an interval is no greater than $N(j+1)$, which is $O((j+1)\log(j+2))$. Even this deliberately loose bound gives

$$\sum_{\rho\in\mathcal{Z}} (1+|\gamma|)^A e^{-c\gamma^2} < \infty \qquad (c > 0,\ A \ge 0).$$

The finitely many small ordinates cause no difficulty, and the negative ones follow by conjugation. This establishes uniform convergence on the compact sets just described. It permits differentiation in u, differentiation in a positive heat interval, and multiplication by the polynomial zero weight introduced later. It does not permit taking $t = 0$, where the Gaussian suppression disappears.

For a single centered mode, $\partial_t e^{tz^2+uz} = z^2 e^{tz^2+uz}$ and $\partial_u^2 e^{tz^2+uz} = z^2 e^{tz^2+uz}$. Normal convergence transfers this equality to the entire packet. The heat equation here is an exact identity for the chosen series, rather than an empirical fit to a finite list of zeros.

Under the Riemann hypothesis the packet is

$$\Theta_t(u) = \sum_{\rho\in\mathcal{Z}} e^{-t\gamma^2}e^{iu\gamma}. \tag{3.4}$$

It is then the Fourier transform of a finite positive discrete measure. In particular, its magnitude is bounded by its value at the origin. The converse can be proved without estimating cancellation among different zeros.

Theorem 3.2 (One heat time criterion)**.** *For any fixed $t > 0$, the Riemann hypothesis is equivalent to boundedness of Θ_t on $[0, \infty)$. It is also equivalent to*

$$\Theta_t(0)^2 - \Theta_t(u)^2 \geq 0 \qquad \text{for every real } u. \tag{3.5}$$

Proof. Under the hypothesis, (3.4) gives the asserted bound and determinant inequality. Conversely, suppose the packet is bounded on the positive half line. Its Laplace transform

$$F_t(w) = \int_0^\infty \Theta_t(u) e^{-wu}\, du$$

is holomorphic for $\Re w > 0$. For $\Re w > 1/2$, absolute convergence gives

$$F_t(w) = \sum_{\rho \in \mathcal{Z}} \frac{e^{tz_\rho^2}}{w - z_\rho}. \tag{3.6}$$

The series on the right is meromorphic, with normal convergence away from its poles. At each distinct centered zero its residue is its positive integer multiplicity times $e^{tz_\rho^2}$, which is nonzero. Analytic continuation therefore excludes a centered zero with positive real part, since it would give a pole in the domain where the left side is holomorphic. Symmetry excludes negative real parts as well. Finally, (3.5) itself bounds the packet by $|\Theta_t(0)|$, so it also gives the converse. □

This criterion replaces an arbitrary collection of test functions by one Gaussian translation family. It still quantifies over every translation. Checking finitely many determinants tests a finite instance of the criterion.

3.2 Why a nonzero residue cannot be hidden by cancellation

An exponential sum may exhibit strong cancellation at many translations. A proof based on one largest term would therefore be unsafe when infinitely many zeros are present. The Laplace transform avoids that issue by separating the modes through their pole locations.

If Θ_t is bounded, then on a compact subset of $\Re w > 0$ the integrand of its Laplace transform is bounded by a constant times e^{-au}, for some $a > 0$. Differentiation with respect to w inserts powers of u, all of which remain integrable. This proves holomorphy on the claimed half plane.

In the smaller initial domain $\Re w > 1/2$, absolute summation justifies integration of each zero term:

$$\int_0^\infty e^{tz^2 + uz - wu}\, du = \frac{e^{tz^2}}{w - z}.$$

For a compact set of w avoiding all centered zeros, large ordinates make $|w - z|$ large. The Gaussian coefficient then gives a summable bound for the meromorphic series. Near one distinct centered zero z_0, isolate all copies of that zero. If its multiplicity is m_0, their sum is

$$\frac{m_0 e^{tz_0^2}}{w - z_0}.$$

Every other term is holomorphic in a sufficiently small disk. The residue cannot vanish because $m_0 \geq 1$ and an exponential is never zero. The reflected zero lies at another pole location unless it coincides with z_0, in which case multiplicity adds rather than cancels.

The two representations agree in their initial half plane. Uniqueness of analytic continuation makes them agree wherever both can be continued. A positive real part of z_0 would place its

nonremovable pole inside the domain of the bounded packet's holomorphic transform. This is the contradiction. The argument is qualitative: even a coefficient too small to be visible in a practical plot still has a nonzero residue.

The determinant formulation has the same content. Under the conjecture, the packet is real and $|\Theta_t(u)| \le \Theta_t(0)$, with $\Theta_t(0) > 0$. Conversely, the displayed determinant inequality forces $\Theta_t(u)^2 \le \Theta_t(0)^2$. No new positivity theorem for an arbitrary zero packet is being assumed.

The packet also makes the obstruction quantitative. Define

$$\delta_\zeta = \sup_{\rho \in \mathcal{Z}} (\Re\rho - 1/2), \qquad M_t(U) = \max\left(1, \sup_{0 \le u \le U} |\Theta_t(u)|\right). \tag{3.7}$$

Then $0 \le \delta_\zeta \le 1/2$. The same Laplace argument proves

$$\limsup_{U \to \infty} \frac{\log M_t(U)}{U} = \delta_\zeta. \tag{3.8}$$

The upper bound follows from summing absolute values. If a smaller exponential rate held, the Laplace transform would be holomorphic beyond that rate, contradicting a pole at a zero further to the right. An off line zero may have a very large ordinate and a very small Gaussian coefficient. Its coefficient remains nonzero, so it eventually determines the possible exponential growth.

To spell out the last implication, suppose the limit superior in (3.8) were below some centered zero's real part. Choose a number a strictly between them. Beyond a finite initial interval the packet would then satisfy $|\Theta_t(u)| \le e^{au}$, after slightly enlarging a if necessary. Its Laplace transform would be holomorphic for $\Re w > a$, but the chosen zero would still place a pole there. This proves the lower bound without requiring that the supremum δ_ζ be attained. When $\delta_\zeta = 0$, the definition $M_t(U) \ge 1$ supplies the matching nonnegative lower bound.

4 The exact prime formula and removal of explicit channels

Let $\Lambda(n)$ denote the von Mangoldt function: it is $\log p$ when $n = p^k$ for a prime p and an integer $k \ge 1$, and it is zero otherwise. Let $\psi = \Gamma'/\Gamma$ be the digamma function, and set

$$\Psi_\infty(v) = -\log\pi + \Re\psi\left(\frac{1}{4} + \frac{iv}{2}\right) \qquad (v \in \mathbb{R}). \tag{4.1}$$

Define the archimedean term and prime packet by

$$\mathfrak{A}_t(u) = \frac{1}{2\pi}\int_{\mathbb{R}} e^{-tv^2}\cos(uv)\Psi_\infty(v)\,dv, \tag{4.2}$$

$$\mathfrak{P}_t(u) = \frac{1}{2\sqrt{\pi t}}\sum_{n \ge 2}\frac{\Lambda(n)}{\sqrt{n}}\left(e^{-(u-\log n)^2/(4t)} + e^{-(u+\log n)^2/(4t)}\right). \tag{4.3}$$

The Gaussian makes the prime series absolutely convergent. The completed explicit formula gives

$$\Theta_t(u) = 2e^{t/4}\cosh(u/2) + \mathfrak{A}_t(u) - \mathfrak{P}_t(u). \tag{4.4}$$

We use the formula in the recovered normalization [6, 5]. With Fourier convention $g(x) = (2\pi)^{-1}\int h(v)e^{-ivx}\,dv$, the even test $h(v) = e^{-tv^2}\cos(uv)$ has transform

$$g(x) = \frac{1}{4\sqrt{\pi t}}\left(e^{-(x-u)^2/(4t)} + e^{-(x+u)^2/(4t)}\right).$$

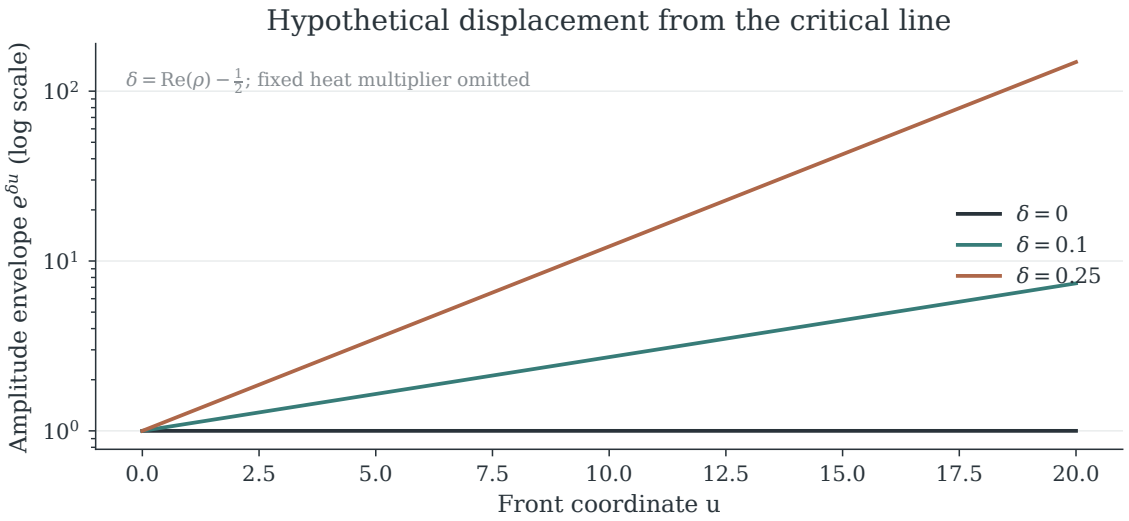


Figure 2: Amplitude envelopes of hypothetical centered modes. A displacement δ gives amplitude growth $e^{\delta u}$, with a fixed nonzero heat multiplier omitted in the illustration. These model curves illustrate the theorem's exponent and are not sampled off line zeta zeros or a complete symmetric zero packet.

The zero symmetry identifies its spectral sum with $\Theta_t(u)$. Its two pole evaluations give $2e^{t/4}\cosh(u/2)$. The prime contribution uses $g(\log n) + g(-\log n)$, which fixes the constants in (4.3). The even archimedean integral reduces to its real digamma part, giving (4.4) with the stated signs.

The prime packet and pole term can each be large. Their exact difference is the zero packet. Bounding the prime packet alone by a positive estimate does not bound that difference at the required scale.

4.1 Following the Gaussian constants

The normalization can be checked without any arithmetic approximation. For real x, the elementary Gaussian Fourier integral is

$$\int_{\mathbb{R}} e^{-tv^2} e^{-ivx}\, dv = \sqrt{\frac{\pi}{t}}\, e^{-x^2/(4t)}.$$

Write $\cos(uv) = (e^{iuv} + e^{-iuv})/2$. The factor one half from this cosine and the factor $(2\pi)^{-1}$ in the Fourier convention give the coefficient $1/(4\sqrt{\pi t})$ in g. The function g is even, so the sum $g(\log n) + g(-\log n)$ doubles it. This produces the coefficient $1/(2\sqrt{\pi t})$ in the prime packet, not twice that coefficient.

The two pole evaluations can be checked just as directly. In the explicit formula the corresponding spectral arguments are $i/2$ and $-i/2$. Both Gaussian values contribute $e^{t/4}$, and the cosine contributes $\cosh(u/2)$. Their sum is the explicit term in (4.4). For a zero, the spectral argument is $(\rho - 1/2)/i$. The Gaussian value becomes $e^{tz_\rho^2}$, while the cosine becomes $\cosh(uz_\rho)$. Summing over the symmetric multiset identifies this sum with $\sum e^{tz_\rho^2 + uz_\rho}$.

At $u = 0$, for example, the prime packet is

$$\mathfrak{P}_t(0) = \frac{1}{\sqrt{\pi t}} \sum_{n \geq 2} \frac{\Lambda(n)}{\sqrt{n}} e^{-(\log n)^2/(4t)}.$$

Every summand is nonnegative. The complete formula is nonetheless a difference involving this positive quantity. At large positive u, the Gaussian centered at $\log n$ selects prime powers with logarithms near u, so the active arithmetic scale is around $n = e^u$. A fixed heat time gives a fixed width in logarithmic coordinates, which corresponds to a multiplicative window in the integers.

Absolute convergence of the prime series requires no prime number theorem. Use the elementary bound $\Lambda(n) \le \log n$ and partition the integers by $j \le \log n < j+1$. There are at most e^{j+1} terms in such a group. The Gaussian in j dominates the resulting exponential factor. The same reasoning covers differentiated packets because their derivatives insert only polynomials in $u - \log n$. This supplies a direct analytic justification for the operations applied to the prime side.

The integral over Ψ_∞, the explicit pole term, and the prime powers have distinct origins. Keeping all three visible prevents a mistaken inference from one individually large or small channel. The exact formula is useful because it states which difference must be controlled.

Definition 4.1 (Pole annihilated packet)**.** Let $L = 1/4 - \partial_u^2$, and define

$$\widetilde{\Theta}_t = L\Theta_t = \sum_{\rho \in \mathcal{Z}} \rho(1-\rho) e^{t z_\rho^2 + u z_\rho}. \tag{4.5}$$

The identity $1/4 - z_\rho^2 = \rho(1-\rho)$ explains the weight. It is nonzero for every nontrivial zero. Thus the operator kills the explicit pole contribution in (4.4) while retaining every zero mode. Differentiated normal convergence gives

$$\widetilde{\Theta}_t = \widetilde{\mathfrak{A}}_t - \widetilde{\mathfrak{P}}_t, \tag{4.6}$$

where tildes on the right mean application of L.

For a translated Gaussian $G_{t,a}(u) = e^{-(u-a)^2/(4t)}$, direct differentiation gives

$$LG_{t,a}(u) = \left(\frac{1}{4} + \frac{1}{2t} - \frac{(u-a)^2}{4t^2} \right) G_{t,a}(u).$$

The annihilated prime packet therefore carries a polynomial factor with both signs. Positivity of the original Gaussian sum is no longer positivity of every transformed term. On the zero side the same operator is simpler because each exponential is an eigenfunction of ∂_u^2. On the explicit pole channel, $\partial_u^2 \cosh(u/2) = \frac{1}{4}\cosh(u/2)$, so cancellation is exact.

The choice of operator is tied to a declared removal target. Its two scalar roots are the centered coordinates $\pm 1/2$ of the exceptional points $s = 0, 1$. Actual nontrivial zeros are strictly inside the critical strip, so none is removed. A different operator which vanished at an unknown nontrivial zero would no longer preserve the obstruction required by the criterion.

4.2 Why the arithmetic channel contains prime powers

The appearance of Λ can be traced directly to the Euler product in its first convergence half plane. There,

$$\log \zeta(s) = \sum_p \sum_{k \ge 1} \frac{p^{-ks}}{k}.$$

Differentiating gives

$$-\frac{\zeta'(s)}{\zeta(s)} = \sum_p \sum_{k \ge 1} (\log p) p^{-ks} = \sum_{n \ge 2} \frac{\Lambda(n)}{n^s}.$$

The factor k from differentiation cancels the denominator k, which is why the coefficient of p^k is $\log p$, rather than $\log(p^k)$. Absolute convergence of these derivative series holds for $\Re s > 1$. For instance, $\Lambda(n) \le \log n$, and $\sum (\log n) n^{-\sigma}$ converges for $\sigma > 1$.

Thus $\Lambda(8) = \log 2$, $\Lambda(9) = \log 3$, and $\Lambda(12) = 0$. The exact prime packet includes these prime powers. Omitting them would change the explicit formula, even if primes supply its largest familiar contribution. The completed formula then carries this arithmetic logarithmic derivative into a Gaussian test against zeros, with the gamma and pole terms accounting for completion.

This derivation explains the arithmetic support. It does not prove the explicit formula by itself, since that formula also involves analytic continuation and a contour argument. Those ingredients are supplied by the cited primary result. Separating the elementary support calculation from that external theorem keeps the source dependence precise.

The archimedean term has a uniform bound on every positive heat strip. A direct proof uses the classical digamma series [3]:

$$\psi(a+iy) - \psi(a) = \sum_{k \ge 0} \frac{iy}{(k+a)(k+a+iy)}, \qquad a > 0.$$

Its absolute value is at most $|y| \sum_{k \ge 0} (k+a)^{-2}$. Therefore $|\Psi_\infty(v)| \le C(1+|v|)$ for a fixed constant C. Since

$$\widetilde{\mathfrak{A}}_t(u) = \frac{1}{2\pi} \int_{\mathbb{R}} (1/4 + v^2) e^{-tv^2} \cos(uv) \Psi_\infty(v)\, dv,$$

its magnitude is bounded independently of u and $t \in [t_0, T]$, whenever $0 < t_0 < T < \infty$. Its energy on a unit translation band is consequently $O(1)$. Removing this explicit channel is an analytic estimate, rather than the missing arithmetic cancellation.

For a nonnegative quantity E_J, call its growth subexponential when

$$\lim_{J \to \infty} \frac{\log(1 + E_J)}{J} = 0. \tag{4.7}$$

Equivalently, for every $\varepsilon > 0$, there is a constant C_ε such that $E_J \le C_\varepsilon e^{\varepsilon J}$ for all positive integers J. This is an upper growth condition. It does not assert a positive lower bound.

Theorem 4.2 (Local energy criterion)**.** *Fix $0 < t_0 < T < \infty$. For positive integers J, set*

$$\widetilde{E}_J = \int_{t_0}^{T} \int_{J}^{J+1} |\widetilde{\Theta}_t(u)|^2 \, du\, dt. \tag{4.8}$$

The Riemann hypothesis is equivalent to subexponential growth of $\widetilde{E}_J$.

Proof. Under the hypothesis, $\rho(1-\rho) = 1/4 + \gamma^2$. Absolute summation in (4.5) gives a uniform bound on the strip, and hence $\widetilde{E}_J = O(1)$.

Conversely, regard $u \mapsto \widetilde{\Theta}_\cdot(u)$ as a function taking values in $H = L^2([t_0, T]; \mathbb{C})$. Cauchy Schwarz on each unit interval and the assumed energy bound imply

$$\int_0^\infty \|\widetilde{\Theta}_\cdot(u)\|_H e^{-au}\, du < \infty \qquad (a > 0).$$

Its H-valued Laplace transform is holomorphic for $\Re w > 0$. Initially for $\Re w > 1/2$, it also equals the normally convergent meromorphic series with terms

$$\frac{\rho(1-\rho) e^{t z_\rho^2}}{w - z_\rho}.$$

The residue at a distinct centered zero is a nonzero function of t in H. Indeed, the multiplicity and $\rho(1-\rho)$ are nonzero, and the exponential never vanishes. A zero with positive real part would therefore give a nonremovable pole in the holomorphic domain. Symmetry now gives the hypothesis. □

The pole annihilator changes the weight but leaves the obstruction intact. Together with the bounded archimedean channel, this criterion locates the required estimate in the centered prime packet.

4.3 From unit bands to a holomorphic transform

The energy statement is deliberately local in u. It does not assume a global square integrable packet, which would already rule out many possible growth behaviors. The Hilbert space H collects only the fixed positive heat interval. Tonelli's theorem identifies

$$\widetilde{E}_J = \int_J^{J+1} \|\widetilde{\Theta}_\cdot(u)\|_H^2 \, du.$$

Let $a > 0$. On a unit band, Cauchy Schwarz gives

$$\int_J^{J+1} \|\widetilde{\Theta}_\cdot(u)\|_H e^{-au} \, du \le e^{-aJ} \widetilde{E}_J^{1/2}.$$

Choose the energy exponent to be $\varepsilon = a$. Subexponential growth then bounds the right side by a constant times $e^{-aJ/2}$. The geometric sum converges. The compact initial band $[0,1]$ is harmless by normal convergence. This proves the weighted integrability used in the theorem for every $a > 0$.

On compact subsets of $\Re w > 0$, the same estimate absorbs any fixed power of u. Thus the Bochner integral defining the H-valued Laplace transform can be differentiated with respect to w, and the transform is holomorphic there. In its initial absolute convergence half plane, one can instead integrate the zero expansion. The compact positive heat interval provides one Gaussian majorant for all the vector coefficients.

Near a distinct centered zero z_0 of multiplicity m_0, the residue vector is

$$r_{z_0}(t) = m_0 \rho_0 (1-\rho_0) e^{t z_0^2}.$$

It is continuous and never zero on $[t_0, T]$, so its H norm is positive. To reduce the contradiction to ordinary complex analysis, choose the continuous linear functional given by pairing with this residue vector. The scalar transform then has nonzero residue $\|r_{z_0}\|_H^2$, whereas the Laplace integral remains holomorphic. This makes explicit why integrating over heat times cannot average the residue away.

The criterion survives the archimedean removal in the form needed for a future estimate. Since that channel has bounded unit band energy, the inequalities $\|f-g\|^2 \le 2\|f\|^2 + 2\|g\|^2$ and $\|f\|^2 \le 2\|f-g\|^2 + 2\|g\|^2$ show that subtracting it preserves the property of subexponential energy. This uses bounded energy, not pointwise arithmetic cancellation. It leaves the full transformed prime expression as the place where an estimate must still be proved.

5 Stable resolution in the physical Gaussian metric

The prime packet is naturally expressed as a sum of translated Gaussians. Treating every integer translate as an independent coordinate produces an unstable frame at large integers. The centers $\log n$ and $\log(n+1)$ differ by about $1/n$, much less than the fixed heat width. Their near coincidence is a coordinate problem which has an exact repair.

For $t > 0$ and a real center a, define the normalized Gaussian

$$g_{t,a}(u) = (2\pi t)^{-1/4} e^{-(u-a)^2/(4t)}. \tag{5.1}$$

It has norm one in $L^2(\mathbb{R}, du)$. For a chosen center c, let $\psi_{r,t,c}$, $r = 0, 1, 2, \ldots$, be the orthonormal scaled Hermite functions

$$\psi_{r,t,c}(u) = \frac{(2\pi t)^{-1/4}}{\sqrt{2^r r!}} H_r\left(\frac{u-c}{\sqrt{2t}}\right) e^{-(u-c)^2/(4t)}. \tag{5.2}$$

Here $H_r(x) = (-1)^r e^{x^2}(d^r/dx^r)e^{-x^2}$ is the physicists' Hermite polynomial. If P_R is orthogonal projection onto indices $0 \le r \le R$, the coherent expansion gives

$$g_{t,a} = e^{-\lambda/2} \sum_{r \ge 0} \frac{\sigma^r}{\sqrt{r!}} \psi_{r,t,c}, \tag{5.3}$$

$$\sigma = \frac{a-c}{2\sqrt{t}}, \qquad \lambda = \sigma^2,$$

$$\|(I - P_R) g_{t,a}\|_2^2 = e^{-\lambda} \sum_{r > R} \frac{\lambda^r}{r!}. \tag{5.4}$$

The last expression is a Poisson upper tail. These formulas and the compression estimates are recovered manuscript results [7].

5.1 The metric and a concrete conditioning example

The word physical refers to the measure du in which the packet is actually evaluated. It does not mean that a separately chosen coefficient norm can be substituted. Completing the square gives the exact overlap

$$\langle g_{t,a}, g_{t,b} \rangle = e^{-(a-b)^2/(8t)}. \tag{5.5}$$

For two translates, the Gram matrix has diagonal entries one and off diagonal entry $\kappa = e^{-(a-b)^2/(8t)}$. Its eigenvalues are $1 + \kappa$ and $1 - \kappa$. When the centers are close, their difference is almost invisible in the physical norm even if the two raw coefficients are individually large.

For example, take $t = 1$, $a = \log 1000$, and $b = \log 1001$. The separation is approximately 0.00099950, the overlap is approximately 0.999999875, and the Gram condition number $(1+\kappa)/(1-\kappa)$ is about 1.60×10^7. These values diagnose a coordinate system, not the size of a zeta coefficient. The asymptotic explanation is exact:

$$1 - e^{-(a-b)^2/(8t)} \sim \frac{(a-b)^2}{8t}, \qquad \log(n+1) - \log n \sim \frac{1}{n}.$$

Consecutive integer centers therefore produce a Gram condition number of order n^2 at fixed heat width.

An orthonormal Hermite basis replaces this almost singular coordinate system. If a finite packet is $f = \sum_r a_r \psi_{r,t,c}$, Parseval's identity gives $\|f\|_2^2 = \sum_r |a_r|^2$ with no unknown Gram loss. This exact equality is the repair. It remains possible for the arithmetic coefficients a_r to be large, or to contain a conjecture level obstruction in one of the first modes.

5.2 Deriving the coherent coefficients

Put $x = (u - c)/\sqrt{2t}$. The ratio of a translated Gaussian to the centered one is

$$\frac{g_{t,a}(u)}{g_{t,c}(u)} = e^{\sqrt{2}\sigma x - \sigma^2}.$$

The Hermite generating identity $e^{2xv-v^2} = \sum_{r\geq 0} H_r(x)v^r/r!$, with $v = \sigma/\sqrt{2}$, turns this ratio into

$$e^{-\sigma^2/2} \sum_{r\geq 0} \frac{H_r(x)\sigma^r}{2^{r/2}r!}.$$

Substitution of the normalized Hermite functions gives (5.3). Squaring its orthogonal coefficients gives $e^{-\lambda}\lambda^r/r!$. Their total is one because the exponential series sums to e^{λ}. The discarded squared norm is therefore exactly a Poisson probability, rather than a bound inferred from a plot.

As finite examples, the following entries are rounded values of that exact tail. The cutoff includes all indices from zero through R.

λ	$R = 4$	$R = 8$	$R = 12$
0	0	0	0
1	3.66×10^{-3}	1.13×10^{-6}	6.36×10^{-11}
4	3.71×10^{-1}	2.14×10^{-2}	2.74×10^{-4}

The table measures squared error. The norm error is its square root. It shows why a cutoff which works for a nearby center need not work for a more distant one. The shell estimate controls the largest permitted displacement before choosing the cutoff.

Suppose a fixed positive width constant $C > 0$ is chosen and the active centers in a shell satisfy $|a - c_J| \leq C\sqrt{J} + O(1)$ and $t \geq t_0 > 0$. Put $\alpha = C^2/(4t_0)$. Then $\lambda \leq \alpha J + O(\sqrt{J})$. For a fixed constant $q > \alpha$, projection through $R_J = \lceil qJ \rceil$ has norm error bounded by

$$\exp\left(-\frac{1}{2}I(q \mid \alpha)J + o(J)\right), \qquad I(q \mid \alpha) = q\log(q/\alpha) - q + \alpha. \tag{5.6}$$

The coefficient mass of the positive prime packet through $X_J = \exp(J + O(\sqrt{J}))$ is at most

$$\sum_{n\leq X_J} \frac{\Lambda(n)}{\sqrt{n}} \leq 2\sqrt{X_J}\log X_J = \exp(J/2 + o(J)).$$

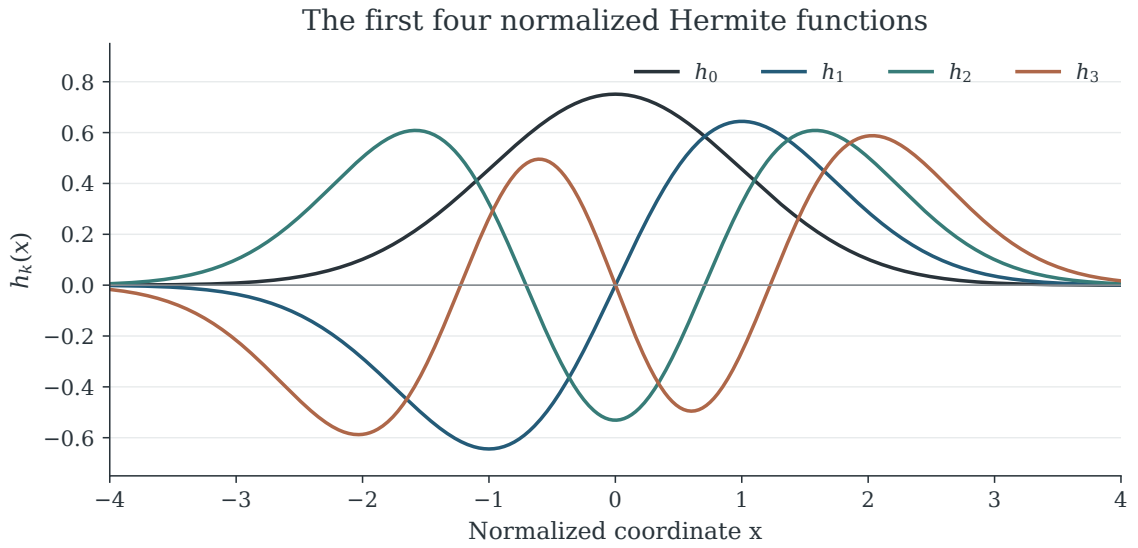
Thus $I(q \mid \alpha) > 1$ makes the aggregate compression error exponentially small. The pole annihilator adds only polynomial factors to the Hermite ladder and leaves this exponential conclusion unchanged.

There are only $O(J)$ retained coordinates, and they are orthogonal in the physical metric. This closes a conditioning problem. It does not make the arithmetic coefficients small.

The exponent in (5.6) comes from a standard elementary exponential bound for a Poisson variable X of mean λ. For $h > 0$,

$$\Pr(X \geq R) \leq e^{-hR}\mathbb{E}(e^{hX}) = \exp\big(-hR + \lambda(e^h - 1)\big).$$

When $R > \lambda > 0$, choose $e^h = R/\lambda$. For $\lambda = 0$, every positive tail is zero. The positive mean case gives the rate $R\log(R/\lambda) - R + \lambda$. Passing to the squared norm tail and then taking a



An example orthonormal basis. No packet approximation is asserted.

Figure 3: The first Hermite modes illustrate an orthogonal Gaussian resolution. The compression theorem concerns the scaled basis, a declared shell of centers, and its Poisson tail. The example basis plot does not measure the unresolved arithmetic coefficients.

square root supplies the factor one half in (5.6). Changing R by one has no effect on its leading shell exponent.

For a concrete choice, let $C = 2\sqrt{t_0}$, so $\alpha = 1$, and take $q = 4$. Then

$$I(4 \mid 1) = 4\log 4 - 3 \approx 2.54518 > 1.$$

The individual norm error decays with exponent approximately $-1.27259J$. Multiplication by a coefficient mass with exponent at most $J/2 + o(J)$ still leaves decay with exponent approximately $-0.77259J$. This example verifies the scale balance in the compression argument. It does not give a bound on the retained coordinates. For different shell widths or heat minima, one must recompute α and choose q accordingly.

The simple coefficient mass bound also has an elementary proof. Since $\Lambda(n) \le \log X$ for $n \le X$, and $\sum_{n \le X} n^{-1/2} \le 2\sqrt{X}$, their product gives the stated estimate. Compression may therefore be arranged without first proving any delicate prime cancellation theorem. The remaining arithmetic estimate lies in the compressed packet, not in the omitted Gaussian tail.

The ground channel shows why the distinction matters. Direct integration gives

$$\int_{\mathbb{R}} g_{s,c}(u)\widetilde{\Theta}_t(u)\,du = K_s\widetilde{\Theta}_{t+s}(c), \qquad K_s = 2\sqrt{\pi s}(2\pi s)^{-1/4}. \tag{5.7}$$

The Gaussian test makes this integral absolutely convergent even without the conjecture. It is a Gaussian pairing, with no assertion that the unweighted zero packet belongs to $L^2(\mathbb{R})$. Completing the square proves the identity term by term. Normal convergence and the Gaussian majorant justify summation [8].

The coefficient of the lowest Gaussian mode is therefore another copy of the same zero field at a later heat time. A stable basis preserves the signal which must be controlled. Removing that coefficient without an exact recovery statement would remove part of the problem.

To verify (5.7), the scalar Gaussian integral is

$$\int_{\mathbb{R}} (2\pi s)^{-1/4} e^{-(u-c)^2/(4s)} e^{uz}\, du = 2\sqrt{\pi s}(2\pi s)^{-1/4} e^{cz+sz^2}.$$

It holds for every complex z, either by completing the square first for real z and continuing analytically, or by the entire Gaussian transform. Multiplying by the mode coefficient at heat time t replaces t by $t+s$. The annihilator weight is unchanged. Absolute values before integration are bounded using $|\Re z| < 1/2$ and Gaussian decay in the ordinate, so the sum and integral can be interchanged without the Riemann hypothesis.

The parameter s in this display is a positive test width, not the complex variable of ζ. At center $c = 0$, the pairing reads the packet at its origin after additional smoothing. At large c, it reads precisely the translations where an off line zero could create growth. Smoothing delays and rescales that mode but cannot make its coefficient vanish. The ground channel is therefore a particularly direct test of whether a proposed coordinate argument has preserved the original question.

6 The classical Möbius packet and exact prime cubes

The Möbius function $\mu(n)$ is zero when a square of a prime divides n. Otherwise it equals $(-1)^k$, where k is the number of distinct prime divisors of n, and $\mu(1) = 1$. For an integer $K \geq 1$, put

$$w_K(n) = \frac{1}{n^2}\left(1 - \frac{1}{n^2}\right)^K, \qquad c_K = \sum_{n\geq 1} \mu(n) w_K(n). \tag{6.1}$$

The series is absolutely convergent. Its natural active scale is n of order $M = \sqrt{K}$. The coefficient c_K is classical. Báez Duarte proved the following exact criterion [4].

Theorem 6.1 (Báez Duarte criterion)**.** *The Riemann hypothesis holds if and only if, for every $\varepsilon > 0$,*

$$|c_K| = O_\varepsilon(K^{-3/4+\varepsilon}). \tag{6.2}$$

Equivalently, for every $\varepsilon > 0$, $|c_{M^2}| = O_\varepsilon(M^{-3/2+\varepsilon})$, with M restricted to square roots of positive integers.

The change in ε between these equivalent formulations is harmless because each is quantified over every positive ε. We do not replace this with the stronger estimate having $\varepsilon = 0$.

6.1 A finite formula and the active scale

Although c_K is written as an infinite arithmetic series, each fixed K has an exact finite expression in even zeta values. Expand its power by the binomial theorem:

$$c_K = \sum_{j=0}^{K} (-1)^j \binom{K}{j} \frac{1}{\zeta(2j+2)}. \tag{6.3}$$

To justify this identity, note first that the expansion has only $K+1$ terms. Each inner series $\sum \mu(n) n^{-2j-2}$ is absolutely convergent. Its Euler product is $\prod_p (1 - p^{-2j-2})$, which equals the reciprocal zeta value in the same half plane of absolute convergence. Thus no limiting interchange near the critical line is involved.

For example,

$$c_1 = \frac{1}{\zeta(2)} - \frac{1}{\zeta(4)}, \qquad c_2 = \frac{1}{\zeta(2)} - \frac{2}{\zeta(4)} + \frac{1}{\zeta(6)}.$$

These formulas are exact. At large K, their alternating binomial terms can be individually much larger than their sum, so straightforward floating point evaluation may lose accuracy. A finite expression is not automatically a stable evaluation scheme. For a direct truncation of the original series at $n = N$, the elementary certified remainder bound is instead

$$\left| c_K - \sum_{n \le N} \mu(n) w_K(n) \right| \le \sum_{n > N} n^{-2} \le \frac{1}{N}.$$

This bound is useful for checking a finite computation, but it is much too coarse to prove the asymptotic criterion by itself.

The active scale follows from the weight. For a real variable $x > 1$, write $y = x^{-2}$. Then $w_K(x) = y(1-y)^K$, and differentiation in y puts its maximum at $y = 1/(K+1)$. The maximum of the continuously extended weight is therefore at $x = \sqrt{K+1}$, close to $M = \sqrt{K}$. For fixed $r > 0$, as M tends to infinity through square roots of integers,

$$M^2 w_{M^2}(rM) = r^{-2} \left(1 - \frac{1}{r^2 M^2}\right)^{M^2} \longrightarrow r^{-2} e^{-1/r^2}.$$

This is a scaling description of the positive weight. It says nothing about the sign distribution of $\mu(n)$ across its active window.

The size difference behind the criterion is now visible. There are on the order of M integers in a window of length proportional to M, and a typical weight there has order M^{-2}. An unsigned mass therefore naturally has order M^{-1}. The desired signed amplitude has the smaller scale $M^{-3/2+\varepsilon}$. This count is only orientation. The barrier proof below supplies a rigorous lower bound for the actual full absolute majorant, with its squarefree support and cube dependence included.

Let Q be a squarefree positive integer, and let $d = \omega(Q)$ be its number of distinct prime divisors. Its alternating prime cube acts on the weight by

$$\Delta_Q w_K(m) = \sum_{b \mid Q} \mu(b) w_K(bm). \tag{6.4}$$

Here b ranges over all positive divisors of Q. The result is a sum of 2^d terms, one for each subset of the prime coordinates.

Proposition 6.2 (Exact cube decomposition)**.** *For every $K \ge 1$ and squarefree Q,*

$$c_K = \sum_{\substack{m \ge 1 \\ (m,Q)=1}} \mu(m) \Delta_Q w_K(m). \tag{6.5}$$

The condition $(m, Q) = 1$ means that m and Q have greatest common divisor one.

Proof. Every squarefree integer n has a unique expression $n = bm$, where $b \mid Q$ and $(m, Q) = 1$. For these factors, $\mu(n) = \mu(b)\mu(m)$. Integers which are not squarefree contribute zero. Splitting the absolutely convergent series according to this expression gives (6.5). □

The cube creates an exact alternating difference, but the remaining bases still carry $\mu(m)$. The next section measures what happens if those signs are removed.

6.2 Two complete small cubes

For $Q = 6$, the divisors correspond to the subsets of $\{2, 3\}$, and the exact packet is

$$\Delta_6 w_K(m) = w_K(m) - w_K(2m) - w_K(3m) + w_K(6m).$$

For $Q = 30$, one adds a prime coordinate:

$$\begin{aligned}\Delta_{30} w_K(m) = {} & w_K(m) - w_K(2m) - w_K(3m) - w_K(5m) \\ & + w_K(6m) + w_K(10m) + w_K(15m) - w_K(30m).\end{aligned}$$

The complete divisor signs are shown here to make the grouping unambiguous.

b	1	2	3	5	6	10	15	30
$\mu(b)$	1	-1	-1	-1	1	1	1	-1

The $Q = 6$ row uses only columns $1, 2, 3, 6$. The condition on the remaining base matters: m must be coprime to the chosen Q. For example, the squarefree integer 42 becomes $6 \cdot 7$ for both cubes. The integer 35 becomes $1 \cdot 35$ for $Q = 6$, but $5 \cdot 7$ for $Q = 30$. In each case the coefficient factors correctly. A nonsquarefree integer such as 49 contributes zero in the original sum and remains zero in the grouped sum.

At $K = 1$, the weights are rational, since $w_1(n) = n^{-2} - n^{-4}$. Factoring the divisor sums gives

$$\Delta_6 w_1(m) = \frac{2}{3m^2} - \frac{25}{27m^4}, \qquad \Delta_{30} w_1(m) = \frac{16}{25m^2} - \frac{208}{225m^4}. \tag{6.6}$$

The following entries are exact. Every listed base is coprime to the Q of its row. The final column multiplies the cube value by the remaining Möbius sign.

Q	m	$\mu(m)$	$\Delta_Q w_1(m)$	Contribution
6	1	1	$-7/27$	$-7/27$
6	5	-1	$17/675$	$-17/675$
6	7	-1	$857/64827$	$-857/64827$
6	25	0	$449/421875$	0
6	35	1	$881/1620675$	$881/1620675$
30	1	1	$-64/225$	$-64/225$
30	7	-1	$6848/540225$	$-6848/540225$
30	11	-1	$17216/3294225$	$-17216/3294225$
30	49	0	$345536/1297080225$	0

These are finite arithmetic examples, not evidence for a universal decay estimate. They show three features which remain at larger scales. Cube values can change sign. Bases with nonzero cube values can have zero Möbius coefficient. Positive and negative contributions still survive after the cube has been formed.

The cube's shape is also scale dependent. At fixed K, large m makes every factor $(1 - 1/(bm)^2)^K$ close to one, so its leading coefficient is $\prod_{p|Q}(1 - p^{-2})$. At the active scale $m \asymp \sqrt{K}$, the different divisors sample different parts of the weight profile. This difference can create local cancellation within a cube. It does not remove the remaining sign $\mu(m)$, which is the quantity the signed criterion still asks us to control.

7 A barrier for discarding the remaining signs

Define the full triangle majorant

$$\mathcal{A}_Q(M) = \sum_{\substack{m\geq 1\\(m,Q)=1}} |\mu(m)|\,|\Delta_Q w_{M^2}(m)|. \tag{7.1}$$

It is finite and satisfies $|c_{M^2}| \leq \mathcal{A}_Q(M)$. We retain the squarefree support through $|\mu(m)|$. A lower bound on this majorant proves that this particular upper bound is insufficient. It gives no lower bound on $|c_{M^2}|$.

7.1 Which quantity is being bounded

For any finite interval I of remaining bases, distinguish its signed contribution from its absolute mass:

$$S_I = \sum_{\substack{m\in I\\(m,Q)=1}} \mu(m)\Delta_Q w_{M^2}(m), \qquad V_I = \sum_{\substack{m\in I\\(m,Q)=1}} |\mu(m)|\,|\Delta_Q w_{M^2}(m)|.$$

The valid inequalities are $|S_I| \leq V_I \leq \mathcal{A}_Q(M)$ and $|c_{M^2}| \leq \mathcal{A}_Q(M)$. There is no general inequality $|c_{M^2}| \leq V_I$. Terms outside the chosen interval remain in the full signed sum. Even a two term toy sum with values 3 inside the interval and 10 outside it shows why its interval mass cannot bound the full sum.

The opposite confusion is equally serious. A large absolute mass does not force a large signed value. The finite sum $3 - 2 - 1$ is zero while its absolute mass is six. The barrier will establish a lower bound on an absolute quantity which is available as an upper bound for the desired signed quantity. Such a lower bound prevents this upper bound from proving the desired estimate. It does not contradict the estimate itself.

Finally, dropping $|\mu(m)|$ creates a different mass. Its terms include nonsquarefree bases which the signed sum never sees. The exact cube table at $m = 25$ or $m = 49$ gives explicit examples: the cube value is positive, but the contribution to both the signed sum and its true triangle majorant is zero. This is why the support count below is part of the proof rather than an optional refinement.

Lemma 7.1 (Uniform positive cube tail)**.** *For a fixed $R > 0$, put*

$$\kappa_R = \frac{1}{\zeta(2)} - \frac{\zeta(4)}{R^2}. \tag{7.2}$$

For every squarefree Q and every integer $m \geq RM$,

$$\Delta_Q w_{M^2}(m) \geq \frac{\kappa_R}{m^2}. \tag{7.3}$$

In particular, $R = 3$ gives $\kappa_R > 0$.

Proof. The limiting divisor sum is

$$\sum_{b|Q} \frac{\mu(b)}{b^2} = \prod_{p|Q}(1 - p^{-2}) \geq \zeta(2)^{-1}.$$

For $0 \le x \le 1$, Bernoulli's inequality gives $0 \le 1 - (1-x)^K \le Kx$. With $K = M^2$, the error in replacing each power by one is at most

$$\frac{M^2}{m^2} \sum_{b \mid Q} b^{-4} \le \frac{\zeta(4)}{R^2}.$$

Multiplying the resulting lower bound by m^{-2} proves the lemma. □

The bound is uniform in the chosen primes. The finite product cannot fall below its product over all primes, while the error sum cannot exceed its sum over all positive integers. Taking $R = 3$ gives

$$\kappa_3 = \frac{6}{\pi^2} - \frac{\pi^4}{810} > 0.$$

The first term is approximately 0.6079 and the subtracted term approximately 0.1203. Thus the selected tail lies beyond the central region where the cube's own signs may fluctuate. Its positivity is a useful rigorous fact, but it still leaves the external signs $\mu(m)$.

To use this positive tail for the exact majorant, we must count squarefree coprime bases. Counting all coprime integers would include bases with zero Möbius coefficient and would not suffice.

Lemma 7.2 (Squarefree coprime bases)**.** *For fixed $R > 0$, the number $N_{Q,R}(M)$ of squarefree integers m in $[RM, 2RM]$ with $(m, Q) = 1$ satisfies, uniformly in squarefree Q,*

$$N_{Q,R}(M) = RM \frac{\varphi(Q)}{Q} \prod_{p \nmid Q} (1 - p^{-2}) + O_R(2^d \sqrt{M}). \tag{7.4}$$

Here $\varphi(Q)$ is Euler's totient function, which counts the positive residues coprime to Q.

Proof. Use the exact identity $\mu(m)^2 = \sum_{\ell^2 \mid m} \mu(\ell)$. In the coprime sector only $(\ell, Q) = 1$ contributes. Inclusion and exclusion over the prime factors of Q counts the possible multiples of ℓ^2 in the interval with main term $RM\ell^{-2}\varphi(Q)/Q$ and error $O(2^d)$. There are $O_R(\sqrt{M})$ values of ℓ, so the summed error is $O_R(2^d\sqrt{M})$. Extending the main sum from $\ell \le \sqrt{2RM}$ to all positive integers contributes only $O_R(\sqrt{M})$. Absolute convergence identifies the extended sum as $\prod_{p \nmid Q}(1 - p^{-2})$, giving (7.4). □

7.2 A worked support count

The squarefree identity in the proof can be understood prime by prime. If no prime square divides m, the only possible square divisor in its Möbius sum is built from no primes, and the sum is one. If some primes do have squares dividing m, summing over their squarefree subsets gives $(1-1)^r = 0$, where $r \ge 1$ is their number. Thus the identity is exactly the indicator of squarefree support.

For each permitted ℓ, write $m = \ell^2 v$. Coprimality of m with Q forces both $(\ell, Q) = 1$ and $(v, Q) = 1$. Inclusion and exclusion gives

$$\#\{v : A < v \le B,\ (v, Q) = 1\} = \sum_{b \mid Q} \mu(b) \left(\left\lfloor \frac{B}{b} \right\rfloor - \left\lfloor \frac{A}{b} \right\rfloor \right).$$

Replacing each floor by its real argument creates an error bounded by a constant times the number of divisors, namely 2^d. This explains exactly where cube size enters the density proof.

Summing over ℓ up to order $\sqrt{M}$ creates the error in Lemma 7.2. No distribution theorem for primes or Möbius signs is used.

For $Q = 6$, the density constant simplifies to

$$\frac{\varphi(6)}{6}\prod_{p\nmid 6}(1-p^{-2}) = \frac{1}{2\zeta(2)}.$$

For $Q = 30$, it is $5/(12\zeta(2))$. As a finite illustration, take $M = 10$ and $R = 3$. The interval is $[30, 60]$. Its $Q = 6$ squarefree coprime bases are

$$31, 35, 37, 41, 43, 47, 53, 55, 59.$$

For $Q = 30$, remove 35 and 55, leaving seven bases. In both coprime sectors, 49 would be counted if squarefree support were forgotten. The main term predicts approximately 9.12 bases for $Q = 6$ and 7.60 for $Q = 30$. These small examples illustrate the meaning of the density constant. The asymptotic lemma includes its error and does not claim that a main term rounded to the nearest integer always gives the exact finite count.

Theorem 7.3 (Absolute cube barrier)**.** *Fix R with $\kappa_R > 0$. Suppose squarefree $Q = Q(M)$ has $d(M) = o(\log M)$ distinct prime divisors. For all sufficiently large M,*

$$\mathcal{A}_Q(M) \geq \frac{\kappa_R}{8R\zeta(2)}\frac{\varphi(Q)}{QM} \geq \frac{\kappa_R}{8R\zeta(2)M(d+1)}. \tag{7.5}$$

Consequently $\mathcal{A}_Q(M) \geq M^{-1-o(1)}$. It cannot itself satisfy the target bound in Theorem 6.1.

Proof. Order the prime divisors as $p_1 < \cdots < p_d$. Since $p_i \geq i+1$,

$$\frac{\varphi(Q)}{Q} = \prod_{i=1}^{d}(1-p_i^{-1}) \geq \prod_{i=1}^{d}\frac{i}{i+1} = \frac{1}{d+1}.$$

Also $\prod_{p\nmid Q}(1-p^{-2}) \geq 1/\zeta(2)$. The error in (7.4) is $M^{1/2+o(1)}$, while its main term is at least $RM/[\zeta(2)(d+1)]$. Therefore eventually

$$N_{Q,R}(M) \geq \frac{RM}{2\zeta(2)}\frac{\varphi(Q)}{Q}.$$

On every counted base, $|\mu(m)| = 1$, and Lemma 7.1 gives $|\Delta_Q w_{M^2}(m)| \geq \kappa_R/(4R^2M^2)$. Their contribution is a lower bound for the full majorant. Multiplication gives (7.5). □

The dimension condition means that the cube size 2^d is $M^{o(1)}$. In this regime the target amplitude is $M^{-3/2+o(1)}$, whereas the triangle majorant is at least $M^{-1-o(1)}$. The missing factor is $M^{-1/2+o(1)}$. It must arise from cancellation or another operation which retains and controls the signs.

The theorem is specific to the exact majorant (7.1). It does not prohibit every possible use of a prime cube. A proof which estimates signed correlations between its remaining bases would be outside the barrier's hypothesis.

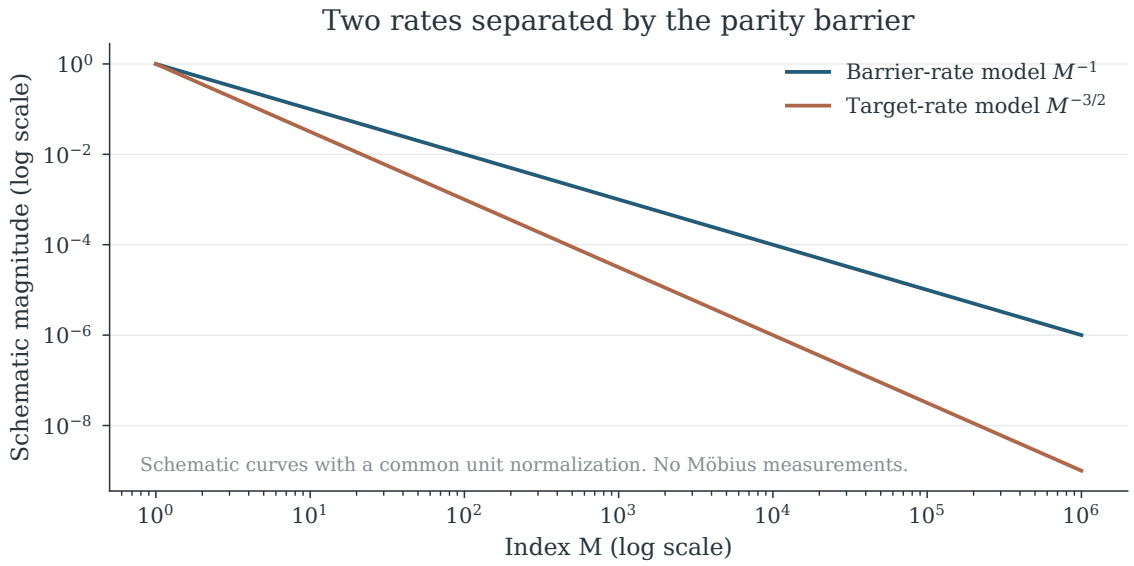


Figure 4: The two exponents in the repaired cube barrier. The curves show theorem scales, rather than measured Möbius amplitudes. The absolute majorant has a lower bound of order $M^{-1-o(1)}$, while the Báez Duarte criterion asks for the signed amplitude to be at most $M^{-3/2+o(1)}$.

7.3 Tracking the proof's scale losses

There are three separate estimates in the theorem. First, the totient factor loses at most $d+1$. This follows from ordering the primes, not from assuming that the chosen primes are the first d primes. Larger prime divisors make the totient density larger, so the lower bound is uniform over every squarefree Q with the stated dimension.

Second, the support error is negligible compared with its main term. The ratio is bounded by a constant times

$$\frac{2^d(d+1)}{\sqrt{M}}.$$

If $d = o(\log M)$, then $2^d = M^{o(1)}$ and $d+1 = M^{o(1)}$. The ratio therefore tends to zero. This is why the statement can retain an increasing cube, while still concluding that a fixed positive fraction of its predicted support is present. If the dimension condition is removed, this particular proof need not give a negligible error. We do not silently extend the theorem beyond that range.

Third, each selected base lies below $2RM$, so its positive cube tail is at least $\kappa_R/(4R^2M^2)$. Multiplying this by the proven support count gives the constant in (7.5). The argument uses only one interval as a lower bound on the full absolute sum, which is the direction an interval restriction actually supports.

To see the incompatibility with the target, choose any fixed amplitude exponent loss $\varepsilon < 1/2$ in the M formulation of the criterion. A quantity bounded below by $M^{-1-o(1)}$ cannot also be bounded above by a constant times $M^{-3/2+\varepsilon}$. Their ratio grows as $M^{1/2-\varepsilon-o(1)}$. The criterion requires every positive ε, so a failure for one such exponent already shows that this majorant cannot supply it. The signed coefficient can still be much smaller than the majorant.

8 Other mechanisms which preserve the obstruction

The saved program investigated several ways of creating contraction after the analytic reductions. Their failures are informative because each identifies a quantity that is preserved [9, 10].

First, an exact prime high pass followed by exact coprime reconstruction preserves a power mode. If a mode has local multiplier $1 - p^{-\rho}$, reconstruction contributes the reciprocal multiplier whenever it is defined. Their product is one. The high pass alone can look smaller, but exact reconstruction restores the removed factor. The historical Lean module `HighpassReconstructionNoGo.lean` certifies this algebraic cancellation. It does not estimate a prime sum.

Second, several diagonal filters of the same logarithmic generator do not create a mixing mechanism on a shared neutral vector. Every filter acts on that vector through its own scalar value, and a common preserved mode remains preserved under composition. The reported modules `DiagonalHypocoercivityNoGo.lean` and `NaturalScaleModeNoGo.lean` formalize this boundary in their declared algebraic settings. Their conclusion concerns those settings, rather than arbitrary arithmetic operators.

Third, global Hardy norms on a prime torus are blind to the squarefree Möbius sign. For a finite prime set, a squarefree monomial is $\prod_{p\in S} z_p$. Replacing every coordinate by $-z_p$ multiplies it by $(-1)^{|S|}$, exactly its Möbius coefficient. The substitution preserves Haar measure on the torus. Every global L^q norm formed with that measure is consequently unchanged by this sign operation. A bound in that invariant norm cannot distinguish the signed packet from the corresponding unsigned squarefree polynomial. This statement does not say that local restrictions or correlations with a separate observable are invariant.

Fourth, a raw positive fourth moment cannot lie below the square of its quadratic energy in a normalized probability space. If ν is its probability measure and $F \in L^4(\nu)$ is the packet, Cauchy Schwarz gives

$$\int |F|^4 \, d\nu \geq \left(\int |F|^2 \, d\nu \right)^2 .$$

This is the precise obstruction to a strict saving below that scale for the raw moment. It does not assert positivity of every separately grouped collision term for arbitrary signed coefficients. A viable replacement would require a signed expression, a justified quotient, or an independently proved correlation estimate.

The stable Hermite construction addresses another failure, the nearly dependent integer Gaussian frame. Its ground identity (5.7) then identifies what survives the repair. Taken together, these results reject particular insufficient operations. They do not establish that no approach to the conjecture can work.

9 The exact open conjecture in arithmetic coordinates

The prime cube identity permits an explicit formulation of the remaining signed estimate. It is independent of any unstated transport constant or choice of physical compression.

Conjecture 9.1 (Signed rough base bound)**.** *For every $\varepsilon > 0$, there is a constant C_ε such that, for*

all integers $K \geq 1$ and every squarefree Q,

$$\left| \sum_{\substack{m \geq 1 \\ (m,Q)=1}} \mu(m) \Delta_Q w_K(m) \right| \leq C_\varepsilon K^{-3/4+\varepsilon}. \tag{9.1}$$

Proposition 9.2 (Exact conditional implication)**.** *Conjecture 9.1 is equivalent to the Riemann hypothesis. The same is true if (9.1) is required only along one declared choice $Q = Q(K)$, including an admissible growing prime cube.*

Proof. By Proposition 6.2, the sum on the left is exactly c_K, regardless of Q. Theorem 6.1 supplies both implications. Its bound is independent of Q because its value is independent of Q. □

This equivalence does not make the estimate easier by itself. It states precisely what must be proved in this coordinate system. A choice of Q is useful only if it exposes a provable signed estimate. The barrier theorem shows that an admissible cube followed by its absolute majorant does not do so.

The heat formulation is equally precise: prove subexponential growth of $\widetilde{E}_J$ in (4.8). The Hermite formulation places that task in stable physical coordinates. The saved program proposes a lower bound on prime detail energy, combined with a historical upper bound and a record argument, as one possible route. That lower bound is currently unproved. We do not promote it to an established estimate by writing it as a spectral gap.

In particular, the expressions "first prime channel" and "rough base channel" describe where the arithmetic signs remain after the declared decompositions. They are not a claim that the Möbius and von Mangoldt packets are connected by a newly proved direct transfer identity. The two rigorous interfaces used here are the heat criterion and the classical Báez Duarte criterion. Each independently identifies a full Riemann hypothesis level bound.

A future result should therefore answer three concrete questions. Does it retain the actual arithmetic signs? Does it control the principal observable, rather than only an average over auxiliary characters? Does its loss remain small enough at the same scale where the conjecture is tested? A numerical profile floor, a sign invariant norm, or an exact transport identity without a cancellation estimate answers none of these by itself.

9.1 What a closing estimate would have to say

The heat criterion accepts a very specific type of conclusion. For every $\varepsilon > 0$, one needs a constant C_ε, depending on the fixed heat strip but not on J, such that

$$\int_{t_0}^{T} \int_{J}^{J+1} |\widetilde{\mathfrak{A}}_t(u) - \widetilde{\mathfrak{P}}_t(u)|^2 \, du \, dt \leq C_\varepsilon e^{\varepsilon J}.$$

The explicit archimedean estimate allows this to be replaced by the corresponding subexponential estimate for the transformed prime packet. An estimate after removing its ground component would require an independent argument recovering that component. An estimate averaged over an additional variable would require an argument bringing it back to the actual packet. These are mathematical obligations of those proposed routes, rather than changes to the criterion.

The arithmetic criterion is equally demanding about quantifiers. It asks for a bound at all positive integers K, with a constant uniform in that scale. A computed range, a favorable subsequence, or a choice of cube which changes at each scale does not replace the required bound unless a theorem covers the missing scales. Since every exact cube gives the same c_K, uniformity over Q in Conjecture 9.1 is automatic once the true signed bound is known. It cannot be supplied by numerical optimization of an unsigned cube profile.

The strongest completed boundary in the present arithmetic route is therefore constructive in a limited sense. It identifies which operation cannot provide the closing saving: the declared full absolute majorant for a cube with $d = o(\log M)$. A new attempt can keep the cube while seeking signed correlations, change the observable while proving an exact recovery estimate, or work directly with the heat packet. In each case the final output must reach one of the established criteria with its constants and quantifiers intact.

10 Reading map and notation

The argument can be read in two passes. A first pass can follow the objects: the centered zero coordinate, the heat packet, the exact prime formula, the stable Gaussian resolution, and the Möbius cube. The final open conjecture then states the desired conclusion without any hidden analytic removal. A second pass can follow the proof obligations: normal convergence, nonzero residues, unit band energy, physical metric control, exact divisor grouping, positive tails, and squarefree support. Each obligation closes a particular step of the representation changes.

Readers mainly interested in arithmetic can begin with Theorem 6.1 and Proposition 6.2, then read the complete barrier in Theorem 7.3. Readers interested in spectral structure can begin with Theorem 3.2, then follow Theorem 4.2 and the ground pairing (5.7). These are separate rigorous interfaces to the same conjecture. The manuscript does not assert that one concrete packet has been directly transformed into the other by a newly proved identity.

Symbol	Meaning
$\rho = \beta + i\gamma$	A nontrivial zero, counted with its multiplicity.
z_ρ	The centered coordinate $\rho - 1/2$.
t, u	Positive heat time and translation, respectively.
Θ_t	The Gaussian sum over centered zero modes.
L	The operator $1/4 - \partial_u^2$ removing the explicit pole channel.
$\widetilde{E}_J$	Squared packet energy on one unit translation band and one fixed heat strip.
$g_{t,a}$	A Gaussian of norm one in the physical measure du.
λ, R	Squared coherent displacement and Hermite cutoff. The cutoff R in this section is separate from the fixed tail multiplier in the barrier.
K, M	Integer sequence index and its natural scale $M = \sqrt{K}$.
Q, d	A squarefree product of primes and its number of prime coordinates.
$\Delta_Q w_K$	The alternating divisor cube acting on the positive weight.
$c_K, \mathcal{A}_Q$	The signed coefficient and its full absolute majorant.

The words exact, bounded, and open refer to different sorts of evidence here. An exact identity preserves a quantity for every object satisfying its hypotheses. A bounded compu-

tation tests a declared finite range or precision. An open estimate is not supplied by either description unless the missing universal argument is written and checked. Keeping those meanings stable lets the reader assess the completed work without mistaking the remaining Riemann hypothesis estimate for a routine final calculation.

11 Evidence, certification, and continuation

The recovered v0.13 report records four focused Lean modules checked in a pinned offline Lean 4.34.0 release candidate environment: physical metric bookkeeping, diagonal preservation, exact high pass reconstruction, and common normalized modes. It also records ordinary Mathlib logical foundations in the axiom audit and the rejection of two false controls. One rejected control exposed a missing nonnegativity hypothesis in a squared energy estimate, which was repaired in that checkpoint [9].

The master archive explicitly marks v0.10 verification as failed or incomplete and v0.11 and v0.12 Lean drafts as unverified. It reports 377 cumulative checkpoint entries and 104 distinct Lean source contents. These are source counts, not a count of independent theorems and not a cumulative certification result [12]. This monograph preserves those distinctions.

The Hermite report records 144 coherent expansion checks and 27 ground overlap checks with residuals at ordinary floating point precision. The prime cube report records finite orbit decompositions, interval masses, and high precision continuum profile optimization [7, 11]. Those experiments audit finite identities and implementation. They do not estimate the asymptotic cancellation in (9.1).

The revised barrier proof in this article makes two matters explicit. Its lower bound concerns the full absolute majorant, rather than treating one interval as an upper bound for the entire signed sum. It also counts squarefree coprime bases, so zero Möbius coefficients cannot inflate the claimed obstruction. The proofs of Lemma 7.2 and Theorem 7.3 supply that support correction. This is an ordinary proof added during consolidation, not a new historical certificate.

The continuation target is therefore well defined. In arithmetic coordinates it is Conjecture 9.1. In the physical zero packet it is the subexponential energy estimate of Theorem 4.2. The existing work supplies exact representations, explicit analytic removals, a stable compression, and several boundaries on proposed closure mechanisms. The signed estimate remains open, and the Riemann hypothesis remains unproved.

References

[1] NIST Digital Library of Mathematical Functions. *Riemann zeta function: reflection formulas*, Section 25.4, equations 25.4.3 and 25.4.4. Version 1.2.8, released 15 September 2026. https://dlmf.nist.gov/25.4. Accessed 8 October 2026.

[2] NIST Digital Library of Mathematical Functions. *Riemann zeta function: zeros*, Section 25.10(i). Version 1.2.8, released 15 September 2026. https://dlmf.nist.gov/25.10. Accessed 8 October 2026.

[3] NIST Digital Library of Mathematical Functions. *Gamma function: series expansions*, Section 5.7, equation 5.7.6. Version 1.2.8, released 15 September 2026. https://dlmf.nist.gov/5.7.E6. Accessed 8 October 2026.

[4] L. Báez Duarte. *A sequential Riesz-like criterion for the Riemann hypothesis*. International Journal of Mathematics and Mathematical Sciences **2005:21** (2005), 3527–3537. https://doi.org/10.1155/IJMMS.2005.3527. Theorem 1.1 and equation 3.6. Open journal mirror: https://www.maths.tcd.ie/EMIS/journals/IJMMS/Volume2005_21/141565.pdf.

[5] A. Bondarenko, D. Radchenko, and K. Seip. *Fourier Interpolation with Zeros of Zeta and L-Functions*. Constructive Approximation **57** (2023), 405–461. Equation 1.1, the unconditional Guinand and Weil explicit formula, and Section 2.4, equation 2.11, the zero counting asymptotic. https://doi.org/10.1007/s00365-022-09599-w.

[6] Komorebi P. *Prime packet heat kernels and a two point criterion for the Riemann hypothesis*. Saved TeX manuscript, 15 August 2026. `Prime_Packet_RH_Heat_Kernel_Core.tex`. Analytic manuscript, recovered for this release.

[7] Komorebi P. *Stable Hermite compression construction*. Saved research report, 12 September 2026. `STABLE_HERMITE_COMPRESSION_REPORT.md`. Analytic formulas and finite numerical audits, recovered for this release.

[8] Komorebi P. *Ground Hermite channel and preservation of the RH front*. Saved analytic note, 12 September 2026. `GROUND_MODE_FRONT_IDENTITY.md`. Manuscript derivation, recovered for this release.

[9] Komorebi P. *Komorebi RH final gap attack*, version 0.13.0. Saved report, 12 September 2026. `FINAL_GAP_ATTACK_REPORT.md`. Reports focused historical Lean verification and analytic conclusions.

[10] Komorebi P. *Exact claim boundary*, version 0.13.0. Saved scope report, 12 September 2026. `CLAIM_BOUNDARY(20260912-220305).md`.

[11] Komorebi P. *Prime cube parity barrier*, version 0.14.0. Saved research report, 13 September 2026. `PRIME_CUBE_PARITY_BARRIER_REPORT.md`. Original report preserved as evidence. This monograph supplies the full majorant and squarefree support corrections in its barrier proof.

[12] Komorebi P. *Komorebi RH master archive*, version 0.14.1, count correction. Saved archive scope record, 13 September 2026. `Komorebi_RH_Master_v0.14.1_README_START_HERE.md`. Historical source counts and checkpoint verification boundaries.