hexagon:2610.00688v1

THE RIEMANN HYPOTHESIS AS A MARGIN: A TWO-SIDED CRITERION VERIFIED TO $2 \cdot 10^{29}$, THE FOURTH-MOMENT ANATOMY OF MÖBIUS CANCELLATION, AND ONE CONSTANT SEEN FROM EVERY SIDE

PEDRO CACERES

Abstract. This paper collects, in one place and with uniform notation, a month of new results around a single way of looking at the Riemann hypothesis: as the statement that certain explicitly computable *margins* stay on one side of explicit constants for ever. Part I studies the Cesàro margin $M(N)$ between a smooth curve $8\sqrt{N} + O(\log N)$ and the weighted prime sum $2\sum_{q<N} \Lambda(q) q^{-1/2} \log(N/q)$: RH is equivalent to $M(N) \geq 0$ for all large N (Suzuki), and we prove, with correctly rounded and ball arithmetic, that $0 < M(N) < 0.18484$ for every real $1 < N \leq 2.15 \cdot 10^{29}$, with explicit unconditional two-sided bounds beyond. The margin is analyzed as a random walk against a curve, as an energy, and as a one-sided mean: RH holds if and only if a weighted mean of the prime-counting error stays below $-\widehat{B}$, $\widehat{B} = 1 + \frac{\gamma}{2} - \frac{1}{2}\log 4\pi$, for ever. Part II dissects the Möbius side of the same coin: the long-run variance of the truncated sawtooth sum $\sum_{d \leq D} \mu(d)(\{x/d\} - \frac{1}{2})$ obeys an exact positive-definite identity, its distribution is empirically Gaussian, and the Gaussian defect of the fourth moment is taken apart exactly — the conductor-1 orbit decouples, the cross term moves only at composite squarefree conductors, two of its four sectors are $O(D \log^4 D)$ unconditionally, a parity sieve annihilates every sector term using an odd number of frequency slots of even conductor, a third sector is proved $O(D^{3/2+o(1)})$, and an empirical jump law, first refuted, is proved exactly in its sector; the four sectors are then computed to $D = 800$, where the remaining conjectured object never exceeds 0.44 against the conjectured $O(D)$. Part III shows that one constant crosses everything: $2\widehat{B}$ is the Báez-Duarte constant of the Nyman–Beurling criterion, and we recompute the Nyman–Beurling distances from scratch through Vasyunin's cotangent formula, measuring $d_N^2 \log N \approx 0.046$ against $2\widehat{B} = 0.04619$ with no zero of ζ used anywhere. Part IV distils the author's twenty-two–part study of the cross-energy $C_2(n, s)$ of consecutive Dirichlet partial sums, largely consolidated with complete proofs in a companion monograph (the window studies of Part V are proved in their own installments): the trichotomy that singles out the critical line as the trace of the pole, the energy law — the mean energy of the prime-counting error is $\geq 2\widehat{B}$ unconditionally, with equality if and only if RH holds with simple zeros: a fourth mirror of the same constant — the energy slope $\partial_\sigma \log|\xi|$ as a truncated Weil functional exact in the primes, the functional reflection C_2^{FR} that reproduces Hardy's function from the integers below N and sees the real part of every zero, and the Davenport–Heilbronn and Beurling control methodology. Part V states precisely what none of this can do: a no-go theorem shows that weights indistinguishable from the primes' to $q^{\delta - 1/2}$ satisfy every averaged statement here and still carry an off-line zero, so the exact identity $\Lambda = \mu * \log$ must enter beyond its first moments. We do not prove the Riemann hypothesis, and no statement in this paper should be read as progress towards it beyond what it literally says.

Contents

Introduction: the programme and the new results 3

Part I. The Cesàro margin of the primes 7

Date: October 2026.

2020 *Mathematics Subject Classification.* Primary 11M26; Secondary 11N05, 11N37, 11M06, 11K65.

Key words and phrases. Riemann hypothesis, Cesàro means, explicit formula, Möbius function, Franel–Landau, variance, Nyman–Beurling, Báez-Duarte, Vasyunin sums, Weil positivity.

1. The margin: overview 7
2. The margin 8
2.1. RH in one variable 9
2.2. What is known unconditionally 10
2.3. Verification 11
3. The local minima and the slope 14
4. The margin as an energy 18
4.1. The energy programme: dyadic blocks and averages in L 21
5. The margin at large N: a probabilistic projection 23
6. The walk of the primes against the curve 25
6.1. The exact step 25
6.2. What it takes to cross 26
6.3. Anatomy of the close approaches 27
7. The margin as a one-sided mean of the prime-counting error 27

Part II. How random is the Möbius function? 30
8. The Möbius side: the exact anchor and its square-root refinement 30
9. The smooth part is tame, and the mass sits at the large divisors 31
10. The variance: an exact, positive-definite identity 32
11. The distribution is empirically Gaussian 34
12. The fourth moment: anatomy of the Gaussian defect 35
12.1. The resonance expansion 35
12.2. The intra-orbit part is exact 35
12.3. The measurement, split exactly 36
12.4. The prime-conductor cancellation 36
12.5. The four sectors of the cross term 38
12.6. The four sectors at scale 43
12.7. What remains 44

Part III. One constant, three mirrors 46
13. One constant 46
14. The mirror of the zeros 46
15. The mirror of the primes: the Cesàro margin 47
16. The mirror of the Möbius sawtooths: Nyman–Beurling 47
17. The computation 47
18. The triangle, and what it cannot do 49

Part IV. The cross-energy of Dirichlet partial sums 50
18.1. The trichotomy and the crossing law 50
18.2. Reflected pairs and the functional reflection 51
18.3. The energy law: a fourth mirror of the same constant 52
18.4. The energy slope 53
18.5. The Euler product inside the strip 54
18.6. The controls 54

Part V. The walls 56
19. What the arguments cannot see, and what they must use 56
19.1. What the energy programme cannot see: a planted zero 56
19.2. Using $\Lambda = \mu * \log$ exactly 57
20. Windows and Weil positivity: a summary 61

21. The wall, stated precisely 62
22. Discussion 63
Appendix A. A cognate positivity problem: the Laguerre kernel of ξ and the Gaussian-mixture route 63
Data and code 64
Use of AI assistance 64
References 65

Introduction: the programme and the new results

Everything in this paper is organized around one picture. The Riemann hypothesis, in each of its classical disguises, asserts that an error term is as small as square-root cancellation allows; in each of the three languages of the convolution identity $\Lambda = \mu * \log$ — primes, Möbius, zeros — that assertion can be renormalized into a *margin*: a computable quantity that must stay on one side of an explicit constant, for ever. The programme of this paper is to build those margins exactly, verify them as far as computation reaches, dissect their anatomy with theorems where theorems are available, and state with precision where the theorems stop. All results below are new (the paper is self-contained: it absorbs two unpublished preprints and a research note of the author, cross-checked across six rounds of independent review, and distils the consolidated cross-energy monograph, whose complete proofs it cites rather than repeats); classical inputs are cited where they enter.

Original contributions at a glance. For the reader who wants the value of the paper in one list, these are the results we believe are new to this work (the detailed statements, with their proved / measured / conditional status, follow by part):

(1) the first verification of the two-sided margin criterion, $0 < M(N) < 0.18484$ for all real $1 < N \le 2.15 \cdot 10^{29}$, by correctly rounded arithmetic on the primes and ball-certified zeros, with explicit unconditional envelopes $E(N)$ beyond (Section 2);
(2) the exact description of the local minima of the margin and their recurrence, through the slope $\alpha(x)$ (Section 3);
(3) the reading of the margin as the variogram of the slope under RH, and the equivalence RH $\iff \int_1^X \alpha^2\, dx/x = X^{o(1)}$ (Section 3);
(4) the unconditional energy identity for the slope, and the mean-energy lower bound attained exactly under RH with simple zeros (Section 4);
(5) the planted-zero no-go theorem: positive weights within $q^{\delta-1/2}$ of the primes' satisfy every averaged ingredient of the programme and still carry an off-line zero (Section 19);
(6) the Möbius representation of the slope, the seed of Part II (Section 19);
(7) the mean identity $M(N) = -2\int_1^N E^* K_N\, dt - 2\widehat{B}\log N$ and the one-sided mean form RH $\iff \mu_K(N) \le -\widehat{B}$, verified with gap $0.0171/\log N$ on the whole range (Section 7);
(8) the exact, positive-definite variance identity for $\widetilde{\Phi}_D$, diagonalized by Jordan's totient, identifying every constant in the data and the 46% gcd cancellation (Section 10);
(9) the primitive-sawtooth decomposition and its orthogonality (Proposition 10.5);
(10) the unit-interval lemma and the complete decoupling of the conductor-1 orbit from the fourth moment (Theorem 12.4);
(11) $X(p) = X(p-1)$ at every prime, by a real-space proof with no regularization (Theorem 12.6);
(12) the sector decomposition of the cross term, with the $(3,1)$ and $(2,2)$ sectors proved $O(D\log^4 D)$ unconditionally (Theorem 12.13);
(13) the closed forms on the $\{2, p, 2p\}$ clusters — $\mathrm{Cov}(\sigma_p^2, \sigma_{2p}^2) = -(p^2-1)/(48p^4)$ among them — and the double-ladder orthogonality (Proposition 12.16, Lemma 12.17);

(14) the exact $2p$-jump law $-(p^2-1)/(32p^4)$, proved at every prime in its sector, with its empirical failure at $p=11$ explained term by exact term (Theorem 12.18);
(15) the parity sieve: every sector term using an odd number of frequency slots of even conductor vanishes identically, because a resonance must solve a congruence modulo 2 with only odd units available (Lemma 12.10, Corollary 12.11);
(16) the pair-coefficient cotangent bound and the first unconditional estimate on a sector with three distinct conductors, $X_{211} \ll D^{3/2+o(1)}$, together with the computation of all four sectors to $D=800$, which locates the sign cancellation that the remaining conjecture needs (Lemma 12.14, Proposition 12.15, Section 12.6);
(17) the halving: the Báez-Duarte constant of the Nyman–Beurling criterion is $2\widehat{B}$, and the third mirror is an equivalence — each off-line quadruple adds the excess $2(2\beta-1)^2/(|\rho|^2|1-\rho|^2)$ (Section 14);
(18) the from-scratch recomputation of the Nyman–Beurling distances through Vasyunin's formula, $d_N^2 \log N \approx 0.046$ against $2\widehat{B}=0.04619$, with no zero of ζ used anywhere (Section 17);
(19) in the cross-energy series (by the companion's own accounting [CEn]): the energy-law formulation and its local Gaussian-window version, the equality characterization at $C_B = 2\widehat{B}$ — the fourth mirror — and the second-order constant $K=2.5125\ldots$ (Section 18.3);
(20) the exceptional-set bound and the Poisson-kernel identification and truncation of the energy slope (Theorem 18.5);
(21) the reading of Landau's formula on the circles of the Euler factors, twisted circles included (Section 18.1 and around);
(22) the functional reflection C_2^{FR}, the restricted-Weil equivalence on its own test functions, and the Hankel criterion in the power sums of $1/(\rho(1-\rho))$;
(23) the systematic control methodology: every criterion tested against the Davenport–Heilbronn function, its Euler-product twin $L(s,\chi_5)$, and a zero-sharing Beurling system;
(24) the refutation of the Gaussian-mixture route for the Laguerre kernel W of ξ, and the log-concavity conjecture left in its place (Appendix A).

What is due to others. The one-variable criterion of Part I — the function $C(N)$, the equivalence RH $\iff$ $M(N)\geq 0$ eventually, and the window $4(\log\xi)''(\frac{1}{2})$ — is a theorem of Suzuki [Su1, Su2], rooted in Kreĭn's screw functions; our proof was written before we found his papers, and priority is his. Weil's explicit formula and criterion, the Franel–Montgomery–Vaughan kernel, the Nyman–Beurling–Báez-Duarte framework with Vasyunin's formula [Vas], Burnol's lower bound [Burn] and the Bettin–Conrey–Farmer theorem [BCF], the zero-density and verification inputs [KLN, PT], and the classical layers of the cross-energy series (Cramér's theorem, Turing's method, Lagarias's pairing, Landau's formula, the hybrid products), itemized in [CEn, §1], are all prior art and cited where they enter. Everything in the list above is built on top of that floor.

In more detail, by part:

Part I (the Cesàro margin).

(1) The two-sided verification: $0 < M(N) < 0.18484 = 4(\log\xi)''(\frac{1}{2})$ for all real $1 < N \leq 2.15\cdot 10^{29}$ — from the primes to 10^{10} with correctly rounded arithmetic (extended to 10^{12} in the ancillary runs), and beyond from the first 4520 zeros certified by ball arithmetic, the verification of RH to $3\cdot 10^{12}$ [PT], and an explicit zero-density bound [KLN]; unconditionally, $-E(N) \leq M(N) \leq 0.18484 + E(N)$ with $E(N)$ explicit and $E(10^{20}) < 3\cdot 10^{-5}$ (Section 2).
(2) The anatomy: the local minima of M described exactly by the slope $\alpha(x) = \sum_{n\leq x}\Lambda(n)n^{-1/2} - 2\sqrt{x} + \frac{c_0}{2}$; the margin as the variogram of α under RH; an unconditional energy identity, with a mean-energy lower bound attained exactly under RH with simple zeros (Sections 3–4); the walk of the primes against the curve, step by step (Section 6); a probabilistic projection of the margin to large N (Section 5).

(3) The mean identity: $M(N) = -2\int_1^N E^* K_N\, dt - 2\widehat{B}\log N$, so RH holds if and only if the K_N-weighted mean $\mu_K(N)$ of the prime-counting error satisfies $\mu_K(N) \le -\widehat{B}$ for all large N; verified with margin $0.0171/\log N$ throughout $10^3 \le N \le 2.15\cdot 10^{29}$ (Section 7).

Part II (the Möbius sawtooths).

(4) The exact variance identity for $\widetilde{\Phi}_D = \sum_{d\le D}\mu(d)(\{x/d\} - \frac{1}{2})$, diagonalized by Jordan's totient into a positive combination of squared restricted Mertens sums; it identifies every constant observed in the data and shows the gcd structure of μ cancels 46% of the diagonal variance, unconditionally (Section 10); the decomposition into orthogonal primitive sawtooths that becomes the main tool (Proposition 10.5).

(5) The fourth-moment anatomy (Section 12): the conductor-1 orbit decouples completely by a one-line unit-interval lemma (Theorem 12.4); $X(p) = X(p-1)$ at every prime, by a real-space proof with no regularization (Theorem 12.6); the four sectors of the cross term, with the $(3,1)$ and $(2,2)$ sectors $O(D\log^4 D)$ unconditionally (Theorem 12.13); closed forms on the $\{2,p,2p\}$ clusters, including $\mathrm{Cov}(\sigma_p^2,\sigma_{2p}^2) = -(p^2-1)/(48p^4)$, and the double-ladder orthogonality (Proposition 12.16, Lemma 12.17); the exact $2p$-jump law $-(p^2-1)/(32p^4)$, proved at every prime in its sector, with its empirical failure at $p=11$ explained term by exact term (Theorem 12.18); the parity sieve, which annihilates half of each remaining sector (Lemma 12.10); the bound $X_{211} \ll D^{3/2+o(1)}$ (Proposition 12.15); and the computation of all four sectors to $D = 800$ (Section 12.6). What remains open is one counting problem, Conjecture 12.19, whose remaining quadruple sector the data say needs sign cancellation, not absolute values.

Part III (one constant, three mirrors).

(6) The halving $2\widehat{B} = \sum_\rho 1/(\rho(1-\rho)) = 2 + \gamma - \log 4\pi$: the Báez-Duarte constant of the Nyman–Beurling criterion is exactly twice the margin constant of Part I — a bridge we have not found stated in the literature; and a from-scratch recomputation of the Nyman–Beurling distances through Vasyunin's cotangent formula, with $d_N^2 \log N$ reaching 0.0461 at $N = 500$ against $2\widehat{B} = 0.04619$, using no zero of ζ anywhere (Sections 16–17).

Part IV (the cross-energy C_2 and its functional reflection).

(7) A distillation of the consolidated cross-energy monograph [CEn]: the trichotomy $\lim_n C_2(n,s) \in \{1/(\frac{1}{4}+t^2), 0, +\infty\}$ that characterizes the critical line; the crossing law $n^* = \gamma^2 + \frac{1}{3} + \cdots$, universal in $\lambda = s(1-s)$; the reflected-pair product, negative off the line on an explicit window with depth governed by $(2\beta-1)^2$; the energy law — the mean energy of the prime-counting error equals $\sum_\rho \lim_n C_2(n,\rho)$ when all zeros are simple, is finite iff RH holds, and is $\ge 2\widehat{B}$ unconditionally, with equality iff RH holds with simple zeros: the fourth mirror — with the second-order constant $K = 2.5125\ldots$; the energy slope $S = \partial_\sigma \log|\xi|$ with its disc, density and truncated-Weil theorems; the functional reflection C_2^{FR} and the Hankel-matrix criterion in the power sums of $1/(\rho(1-\rho))$; and the Davenport–Heilbronn / Dirichlet-twin / Beurling control methodology (Section 18.3 and around).

Part V (the walls).

(8) The no-go theorem: positive weights within $q^{\delta-1/2}$ of $\Lambda(q)q^{-1/2}$ satisfy the energy identity, the prime number theorem and Chebyshev's identity up to $x^{1/2+\delta}$, and yet carry an off-line zero; so no argument built from size estimates and fixed-order averages alone can prove RH — the identity $\Lambda = \mu * \log$ must be used exactly (Section 19); the computational study of Weil positivity in windows (Section 20); the precise statement of the two walls — uniformity

on the moving diagonal, and the exceptional set — on the Möbius side (Section 21); and, as an appendix, the refutation of a natural positivity route for the Laguerre kernel of ξ, with a new structural conjecture in its place (Appendix A).

Two remarks on reading. First, the parts are mutually reinforcing but independently readable; each keeps the notation of its subject, and the one constant that recurs across parts is $\widehat{B} = \sum_\rho \rho^{-1} = 1 + \frac{\gamma}{2} - \frac{1}{2}\log 4\pi$. Three notational warnings: $M(N)$ always denotes the Cesàro margin, while the Mertens function, which appears only in Part II, is always written $M(x)$ with argument x; $V(D)$ is the variance numerator of Part II, while $V_N(s)$ is the Bettin–Conrey–Farmer polynomial of Part III; the cross term $X(D)$ of Part II is unrelated to the random variable X of the Rice computations in Part I; and $E(N)$ (the explicit envelope of Part I), $E^*(t)$ (the prime-counting error of Part I), $E(u)$ (its normalization in Part IV) and $e_4(D)$ (the Gaussian defect of Part II) are four different objects. Second — and this sentence governs everything that follows — we do not prove the Riemann hypothesis; every theorem here either averages where RH is pointwise, or verifies where RH is perpetual, and Part V makes that distinction into mathematics.

Part I
The Cesàro margin of the primes

1. The margin: overview

For $N > 1$ put

$$C(N) = 2\sum_{n<N} \frac{\Lambda(n)}{\sqrt{n}} \log\frac{N}{n},$$

the Cesàro (Riesz) mean of order one of $\sum \Lambda(n)n^{-1/2}$, and

$$R(N) = 8\sqrt{N} - \Big(\log\pi + \gamma_E + \frac{\pi}{2} + 3\log 2\Big)\log N - 16 + \frac{\pi^2}{2} + 4G - 2\sum_{k\geq 1} \frac{N^{-(2k+1/2)}}{(2k+\frac{1}{2})^2}, \tag{1}$$

with γ_E Euler's constant and G Catalan's constant. The three pieces of R are the pole of ζ, the conductor with the Γ factor at $s = \frac{1}{2}$, and the trivial zeros. The object of this paper is the *margin* $M(N) = R(N) - C(N)$. Both R and C are of size $8\sqrt{N}$. By Suzuki's theorem [Su1, Su2] (Theorem 2.3), RH holds if and only if $M(N) \geq 0$ for all large N, and then $0 \leq M(N) \leq 0.18484$ for every $N > 1$. So RH is equivalent to confining a sum of size $8\sqrt{N}$ over the prime powers below N to a window of fixed width, from both sides. This is as hard as RH, and we do not prove it.

What the paper does. It studies this one-variable inequality from three sides and records exactly where each side stops.

(1) *The margin* (Section 2). The margin is a Fejér sum over the zeros (Proposition 2.1). The window width is $4(\log\xi)''(\frac{1}{2})$, computable without the zeros (Corollary 2.4). The two-sided inequality holds for all real $1 < N \leq 10^{12}$, by computations from the primes with rigorous error bounds, and, from the certified first 4520 zeros, up to $2.15\cdot 10^{29}$; without any hypothesis it holds up to an explicit error $E(N)$ (Proposition 2.8, Theorem 2.6, Corollary 2.7).

(2) *The minima and the slope* (Section 3). M is convex between prime powers, so its local minima form an explicit countable set, and RH is equivalent to their nonnegativity (Corollary 3.2). Under RH they accumulate at 0 (Proposition 3.3). The slope of the margin is the fluctuation α of $\sum \Lambda(n)n^{-1/2}$ (Proposition 3.4). The minima are the downward zeros of α, and they follow from one another by a recurrence (Proposition 3.6).

(3) *Energy* (Section 4). Under RH, M is the variogram of α, and the window is the inequality $0 \leq \langle|\Delta f|^2\rangle \leq 4\langle|f|^2\rangle$ (Proposition 4.1). RH is equivalent to finite energy of α (Proposition 4.2), and to a one-sided bound for the L-average of M (Proposition 4.7). Without hypothesis we prove an energy identity (Theorem 4.4) and a lower bound for the energy (Theorem 4.6). Block energies to 10^{10} agree with the prediction.

(4) *Limits* (Section 19). Theorem 19.1 plants an off-line zero in positive weights that satisfy every structural input of the energy arguments, so these arguments cannot prove RH. Proposition 19.3 carries the identity $\Lambda = \mu * \log$, which excludes the planted weights, through the slope exactly. The result is a type I sum of μ against the sawtooth (Remark 19.6). Section 20 summarizes the companion results on Weil's form in windows [CE20, CE21, CE22].

What is due to others. Weil's explicit formula and criterion [We52, Bo, IK]; the explicit formula for Riesz means [In]; the one-variable criterion, its bounds and its Laplace transform are Suzuki's [Su1, Su2] (we found his papers after our proofs were written); Landau's theorem on Laplace transforms [Wi, In]; the evaluation of $\sum_\rho(\rho - \frac{1}{2})^{-2}$ [Vo]; the limiting distributions of prime error terms [RuSa, ANS]; Beurling systems with large oscillation [DMV]; the distribution of the Mertens function [Ng]. What we believe is new: the verification to 10^{12}, the exact description of the minima

and their recurrence, the variogram reading of the margin, the energy identity and the energy lower bound, the planted-zero theorem, and the Möbius representation of the slope.

Controls. A statement that holds for the Davenport–Heilbronn function [DH, T], which has the functional equation of ζ but no Euler product and zeros off the line, cannot be the reason for RH. We use it, a Beurling system [DMV] and planted weights (Section 19.1) as controls throughout.

Notation. The nontrivial zeros are $\rho = \frac{1}{2} + i\gamma_\rho$, with γ_ρ real exactly when ρ is on the line; m_γ is the multiplicity. We write $T_0 = 3 \cdot 10^{12}$ for the height to which RH is verified [PT]. For a test function g, $\widehat{g}(z) = \int g(x) e^{izx} dx$. The explicit formula, for G even, continuous, of bounded variation and compact support [Bo, IK], reads

$$\sum_\rho \widehat{G}(\gamma_\rho) = \widehat{G}(\tfrac{i}{2}) + \widehat{G}(-\tfrac{i}{2}) - G(0)\log\pi + \frac{1}{2\pi}\int_{-\infty}^{\infty} \widehat{G}(r)\,\operatorname{Re}\psi\big(\tfrac{1}{4} + \tfrac{ir}{2}\big)dr - \sum_{n\geq 2} \frac{\Lambda(n)}{\sqrt{n}}\big(G(\log n) + G(-\log n)\big), \tag{2}$$

the sum over the zeros taken symmetrically. Weil's criterion says that RH holds if and only if the right side is nonnegative for every $G = g * \tilde{g}$, where $\tilde{g}(x) = \overline{g(-x)}$. The right side is *Weil's quadratic form* $Q(g)$.

2. The margin

The margin is the explicit formula for one test function, the autocorrelation of the indicator of $[0, \log N]$.

Proposition 2.1 (the margin over the zeros)**.** *Let $N > 1$ and $L = \log N$. Without any hypothesis,*

$$M(N) = \sum_\rho F_L(\gamma_\rho), \qquad F_L(z) = \frac{4\sin^2(Lz/2)}{z^2} = \Big|\int_0^L e^{izu} du\Big|^2 \ (z \in \mathbb{R}),$$

the sum over all nontrivial zeros with multiplicity, absolutely convergent. More generally, for a window $\varphi \geq 0$ on $[0, L]$ with autocorrelation $A(u) = \int_0^{L-u} \varphi(v)\varphi(v+u)dv$ and $t \in \mathbb{R}$, let $C_N^\varphi(t) = 2\sum_{n<N} \Lambda(n) n^{-1/2} A(\log n)\cos(t\log n)$ and let $R_N^\varphi(t)$ be the right side of (2) without the prime term for $G(u) = A(|u|)\cos(tu)$. Then $R_N^\varphi(t) - C_N^\varphi(t) = \sum_\rho \frac{1}{2}\big(\widehat{\varphi}(\gamma_\rho - t)\overline{\widehat{\varphi}(\bar{\gamma}_\rho - t)} + \widehat{\varphi}(\gamma_\rho + t)\overline{\widehat{\varphi}(\bar{\gamma}_\rho + t)}\big)$, a two-parameter family [CE20, Proposition 10]*; the case $\varphi = 1$, $t = 0$ is the first statement.*

Proof. Take $G(u) = (L - |u|)_+$, which is $g * \tilde{g}$ for $g = 1_{[0,L]}$, so $\widehat{G} = |\widehat{g}|^2 = F_L$ on $\mathbb{R}$; G is even, continuous, of bounded variation, supported in $[-L, L]$, and $\widehat{G}(z) = O(|z|^{-2})$ on horizontal lines, so (2) applies. Its prime side is $\sum_n \Lambda(n) n^{-1/2} \cdot 2(L - \log n)_+ = C(N)$. The pole terms are $\int G(u)(e^{u/2} + e^{-u/2})du = 8(\sqrt{N} + N^{-1/2} - 2)$. For the archimedean term, Gauss's integral $\psi(z) = \int_0^\infty (\frac{e^{-x}}{x} - \frac{e^{-zx}}{1-e^{-x}})dx$ and $\frac{1}{2\pi}\int \widehat{G}(r) e^{-irx/2} dr = G(\frac{x}{2})$ give, after $x = 2u$,

$$-L\log\pi + \int_0^\infty \Big[\frac{Le^{-2u}}{u} - \frac{2e^{-u/2}G(u)}{1-e^{-2u}}\Big]du,$$

and Proposition 2.2 evaluates the total as $R(N)$. The general window is the same computation with $G(u) = A(|u|)\cos(tu) = \frac{1}{2}\sum_\pm (k_\pm * \tilde{k}_\pm)(u)$, $k_\pm(u) = \varphi(u)e^{\pm itu}$. □

Proposition 2.2 (the remainder in closed form)**.** *The right side of (2) for $G(u) = (L - |u|)_+$, without the prime term, is $R(N)$ given by (1).*

Proof. The pole term is $8(\sqrt{N} + N^{-1/2} - 2)$. Extending the archimedean integral to $(0, \infty)$,

$$L\int_0^\infty \Big[\frac{e^{-2u}}{u} - \frac{2e^{-u/2}}{1-e^{-2u}}\Big]du + \int_0^\infty \frac{2u\,e^{-u/2}}{1-e^{-2u}}du - \int_L^\infty \frac{2(u-L)e^{-u/2}}{1-e^{-2u}}du.$$

The first integral is $\psi(\frac14) = -\gamma_E - \frac{\pi}{2} - 3\log 2$. Expanding $(1-e^{-2u})^{-1} = \sum_k e^{-2ku}$, the second is $2\sum_{k\ge0}(2k+\frac12)^{-2} = \frac12\psi'(\frac14) = \frac{\pi^2}{2} + 4G$, and the third is $2\sum_{k\ge0} N^{-(2k+1/2)}(2k+\frac12)^{-2}$, whose term $k=0$ cancels the $8N^{-1/2}$ of the pole. □

A second derivation, independent of Proposition 2.1, integrates von Mangoldt's explicit formula for $\sum_{n<N}\Lambda(n)n^{-1/2}$ twice in $\log N$ [In, MV]: it gives $C(N) = R(N) - M(N)$ with

$$R(N) = 8(\sqrt{N}-1) - 2\frac{\zeta'}{\zeta}(\tfrac12)\log N + 2\sum_{m\ge1}\frac{1-N^{-(2m+1/2)}}{(2m+\frac12)^2}, \qquad M(N) = 2\sum_\rho \frac{1-N^{i\gamma_\rho}}{\gamma_\rho^2},$$

and comparing with (1) gives $\log\pi + \gamma_E + \frac{\pi}{2} + 3\log 2 = 2\frac{\zeta'}{\zeta}(\frac12)$, which is $\xi'(\frac12) = 0$. So the remainder is the smoothed sum of the smooth part of the partial sum (the pole, the constant $\frac{\zeta'}{\zeta}(\frac12)$ and the trivial zeros), and the margin is the smoothed sum of its oscillating part.

2.1. **RH in one variable.**

Theorem 2.3 (Suzuki [Su1, Su2])**.** *RH holds if and only if $M(N) \ge 0$ for all sufficiently large N. Under RH,*

$$M(N) = 4\sum_{\gamma>0}\frac{1-\cos(\gamma\log N)}{\gamma^2} \in \Big[0,\ 8\sum_{\gamma>0}\gamma^{-2}\Big], \tag{3}$$

so then $M(N) \ge 0$ for every $N > 1$.

Proof. By Proposition 2.1, $M(e^L) = \sum_\rho F_L(\gamma_\rho)$. Under RH every term is ≥ 0, and pairing $\pm\gamma$ gives (3). Conversely, suppose $M(e^L) \ge 0$ for $L \ge L_0$. For $\operatorname{Re} s > \frac12$, since $|F_L(\gamma)| \le (2+2e^{L/2})|\gamma|^{-2}$ and $\sum|\gamma_\rho|^{-2} < \infty$, Fubini gives

$$\int_0^\infty M(e^L)e^{-sL}dL = \sum_\rho \frac{2}{s(s^2+\gamma_\rho^2)} =: \mathcal{G}(s),$$

meromorphic on $\mathbb{C}$ with poles at 0 and at $s = \pm i\gamma_\rho$. A pole s fixes $\gamma^2 = -s^2$, so only $\pm\gamma$ meet there, with equal nonzero residues, and no cancellation occurs. An off-line zero $\frac12 + \delta + i\gamma_0$, $\delta > 0$, has $\gamma_0 \ne 0$ because ζ has no real zeros in $(0,1)$, and gives a pole at $\delta + i\gamma_0$ off the real axis; the zeros on the line give poles on the imaginary axis. So $\mathcal{G}$ is analytic on $(0,\infty)$. By Landau's theorem [In, MV], the Laplace integral of a function that is ≥ 0 on $[L_0,\infty)$ has a singularity at the real point of its abscissa of convergence; it differs from $\mathcal{G}$ by an entire function, so the abscissa is ≤ 0, and $\mathcal{G}$ has no pole with $\operatorname{Re} s > 0$. □

Status: proved. It is Landau's oscillation method for a Riesz mean of $\psi(x)-x$ weighted by $n^{-1/2}$, and the theorem is Suzuki's: with $M = 2\Psi(\log N)$, RH holds if and only if $\Psi \ge 0$, and if and only if $\Psi = O(1)$ [Su1, Theorems 1.6, 1.7]; the criterion with the main term alone is [Su2, Theorem 1]. Our proof is independent. The constants in (3) are values of ξ at the centre: Hadamard's product gives $(\log\xi)''(s) = -\sum_\rho(s-\rho)^{-2}$, and at $s = \frac12$

$$(\log\xi)''(\tfrac12) = \sum_\rho\frac{1}{\gamma_\rho^2} = -8 + \frac{\pi^2}{4} + 2G + \Big(\frac{\zeta''}{\zeta} - \Big(\frac{\zeta'}{\zeta}\Big)^2\Big)(\tfrac12) = 0.0462100, \tag{4}$$

unconditionally; under RH it is $2\sum_{\gamma>0}\gamma^{-2}$, and with the 4520 zeros below 5000 and the tail the two agree to $6\cdot10^{-10}$.

Corollary 2.4 (the sandwich)**.** *The following are equivalent: (1) RH; (2) $M(N) \ge 0$ for all large N; (3) $M(N) \le B$ for some B and all large N; (4) for every $N > 1$,*

$$R(N) - 4(\log\xi)''(\tfrac12) \ \le\ C(N) \ \le\ R(N), \qquad 4(\log\xi)''(\tfrac12) = 0.18484. \tag{5}$$

Proof. (1)⇒(4) is (3) with (4); (4)⇒(2),(3) are immediate; (2)⇒(1) is Theorem 2.3; for (3)⇒(1) apply the same argument to $B - M(e^L) \geq 0$, whose Laplace integral has the same poles off the real axis. □

Status: proved; the equivalences (1)–(3) and the bounds are in [Su1], and the evaluation (4) of $\sum_\rho \gamma_\rho^{-2}$ is classical [Vo]. Writing the window this way makes (5) checkable without the zeros, which is what the verification below uses. The two sides are not redundant in general: a real zero $\beta \in (\frac{1}{2}, 1)$ would push M upwards only, and only the upper side sees it. For ζ, which has no such zeros, either side alone is equivalent to RH; for the Beurling system and the fake curve of [CE20] the lower side never fails and the upper side fails for all large N.

2.2. **What is known unconditionally.**

Proposition 2.5 (unconditional bound, crude form)**.** *For every $N > 1$, $M(N) \geq -8(3+\sqrt{N})S \geq -2.29 \cdot 10^{-11}(3+\sqrt{N})$, with $S = \sum_{\operatorname{Im}\rho > T_0} (\operatorname{Im}\rho)^{-2} \leq 2.854 \cdot 10^{-12}$.*

Proof. In Proposition 2.1 the zeros on the line contribute ≥ 0. By Platt and Trudgian [PT] every zero with $0 < \operatorname{Im}\rho \leq T_0$ is on the line, so an off-line zero has $\gamma_\rho = \gamma_0 + id$ with $|\gamma_0| > T_0$, $|d| < \frac{1}{2}$. With $F_L(z) = (2 - e^{iLz} - e^{-iLz})/z^2$, $|F_L(z)| \leq (2 + N^{|d|} + N^{-|d|})/|z|^2 \leq (3+\sqrt{N})/\gamma_0^2$, and summing over the off-line ordinates, which come in quadruples $\pm\gamma_0 \pm id$, gives at most $8(3+\sqrt{N})S$ (a conservative constant). Stieltjes integration and Trudgian's bound for $N(u)$ [Tr] give $S \leq \int_{T_0}^\infty 2N(u)u^{-3}du \leq 2.854 \cdot 10^{-12}$. □

The factor $\sqrt{N}$ is the price of a hypothetical zero near $\operatorname{Re} s = 1$. A zero-density estimate says that such zeros are rare, and it improves the bound by many orders of magnitude. Let $N(\sigma, T)$ be the number of zeros $\beta + i\gamma$ with $\beta > \sigma$, $0 < \gamma < T$. Kadiri, Lumley and Ng [KLN] prove, for $T \geq H_0 = 3.0610046 \cdot 10^{10}$ and each σ in their Table 1 ($\sigma = 0.60, 0.65, \ldots, 0.85, 0.86, \ldots, 0.99$), with constants A_σ, B_σ and $k_\sigma \leq 1$ listed there,

(6)
$$N(\sigma,T) \leq A_\sigma (\log k_\sigma T)^{2\sigma} (\log T)^{5-4\sigma} T^{\frac{8}{3}(1-\sigma)} + B_\sigma (\log T)^2 \leq A_\sigma (\log T)^{5-2\sigma} T^{\frac{8}{3}(1-\sigma)} + B_\sigma (\log T)^2.$$

Theorem 2.6 (explicit unconditional lower bound)**.** *Let $T_0 = 3 \cdot 10^{12}$, $\widetilde{N}(t) = \frac{t}{2\pi}\log\frac{t}{2\pi}$, and let $\frac{1}{2} = \sigma_0 < \sigma_1 < \cdots < \sigma_J < \sigma_{J+1} = 1$ be $\frac{1}{2}$, the twenty values of σ in* [KLN, Table 1]*, and 1. Put*

$$S_{1/2} = \int_{T_0}^\infty \widetilde{N}(t)\, t^{-3}\, dt, \qquad S_j = 2\int_{T_0}^\infty \min\Big(N_{\mathrm{KLN}}(\sigma_j, t), \tfrac{1}{2}\widetilde{N}(t)\Big) t^{-3}\, dt \quad (1 \leq j \leq J),$$

with N_{KLN} the right side of (6). Then for every $N > 1$

$$M(N) \geq -E(N), \qquad E(N) = 4\Big[(3 + N^{\sigma_1 - \frac{1}{2}})S_{1/2} + \sum_{j=1}^{J} \big(N^{\sigma_{j+1} - \frac{1}{2}} - N^{\sigma_j - \frac{1}{2}}\big) S_j\Big]. \tag{7}$$

One may take $S_{1/2} = 1.4822 \cdot 10^{-12}$, $S_j = S_{1/2}$ to four digits for $\sigma_j \leq 0.75$, $S(0.80) = 4.71 \cdot 10^{-13}$, $S(0.85) = 8.63 \cdot 10^{-15}$, $S(0.90) = 1.57 \cdot 10^{-16}$, $S(0.95) = 2.85 \cdot 10^{-18}$, $S(0.99) = 1.14 \cdot 10^{-19}$, and $E(N)$ takes the values of Table 1. In particular $|\min(M(N), 0)| < 10^{-3}$ for $N < 5 \cdot 10^{24}$ and $E(N) < 0.043$ for $N < 2.3 \cdot 10^{29}$.

Proof. As in Proposition 2.5, $M(N) \geq -\sum' |F_L(\gamma_\rho)|$, the sum over the off-line zeros, all of which have $|\operatorname{Im}\rho| > T_0$, and $|F_L(\gamma_0 + id)| \leq (3 + N^{|d|})/\gamma_0^2$. A zero $\rho = \beta + i\gamma$ with $\beta > \frac{1}{2}$, $\gamma > T_0$, brings the quadruple $\rho, \bar\rho, 1-\rho, 1-\bar\rho$, so

$$\sum{}' |F_L| \leq 4 \sum_{\beta > \frac{1}{2},\ \gamma > T_0} \frac{3 + N^{\beta - 1/2}}{\gamma^2}.$$

Write $\Sigma(\sigma)=\sum_{\beta>\sigma,\gamma>T_0}\gamma^{-2}$. Since $N^{\beta-1/2}\le N^{\sigma_{j+1}-1/2}$ when $\sigma_j<\beta\le\sigma_{j+1}$, Abel summation over the bins gives

$$\sum_{\beta>\frac12,\gamma>T_0}\frac{N^{\beta-1/2}}{\gamma^2}\le N^{\sigma_1-\frac12}\Sigma(\tfrac12)+\sum_{j=1}^{J}\big(N^{\sigma_{j+1}-\frac12}-N^{\sigma_j-\frac12}\big)\Sigma(\sigma_j).$$

It remains to show $\Sigma(\frac12)\le S_{1/2}$ and $\Sigma(\sigma_j)\le S_j$. Since $N(\sigma,T_0)=0$, Stieltjes integration gives $\Sigma(\sigma)=2\int_{T_0}^{\infty}N(\sigma,t)t^{-3}dt$. By the symmetry $\beta\leftrightarrow1-\beta$, at most half of the zeros up to height t lie to the right of the line, so $N(\sigma,t)\le\frac12 N(t)$; and $N(t)\le\widetilde{N}(t)$ for $t\ge T_0$, since $N(t)=\frac{t}{2\pi}\log\frac{t}{2\pi e}+\frac78+\theta(0.112\log t+0.278\log\log t+2.51+0.2/t)$ with $|\theta|\le1$ [Tr], and the term $-\frac{t}{2\pi}$ dominates. This gives $\Sigma(\frac12)\le S_{1/2}$. For $\sigma_j\ge0.6$, (6) applies since $T_0>H_0$, and $\Sigma(\sigma_j)\le S_j$. The values used come from adaptive quadrature in $u=\log t$ (`ebound.py`); they are rigorous upper bounds, because each integrand $g(u)$ is positive and decreasing for $t\ge T_0$ (both $N_{\mathrm{KLN}}(\sigma,t)t^{-2}$ and $\widetilde{N}(t)t^{-2}$ decrease), so a left Riemann sum of step 10^{-4} on $[\log T_0,\log T_0+120]$ plus an explicit tail bounds the integral from above, and every quadrature value exceeds that bound (`ebound_rig.py`; the ratio is at most 0.9984). For instance $S_{1/2}=(\log\frac{T_0}{2\pi}+1)/(2\pi T_0)=1.47970\cdot10^{-12}$ exactly, the Riemann sum gives $1.47977\cdot10^{-12}$ and the value used is $1.48222\cdot10^{-12}$. □

Corollary 2.7 (unconditional two-sided bound)**.** *For every $N>1$,*

$$-E(N)\le M(N)\le4(\log\xi)''(\tfrac12)+E(N)=0.18484+E(N).$$

Proof. The lower side is Theorem 2.6. For the upper side split $M=\sum_\rho F_L(\gamma_\rho)$ into zeros on and off the line. A pair $\pm\gamma$ on the line contributes $4(1-\cos\gamma L)/\gamma^2\le8/\gamma^2$, and the off-line zeros contribute at most $E(N)$ in absolute value, as in the proof of Theorem 2.6. Identity (4) holds without RH, with $\gamma_\rho=(\rho-\frac12)/i$ complex for an off-line zero. A quadruple $\pm\gamma_0\pm id$ contributes to it $4(\gamma_0^2-d^2)/(\gamma_0^2+d^2)^2>0$, since $|d|<\frac12<\gamma_0$. Hence $\sum_{\text{on the line}}\gamma^{-2}\le(\log\xi)''(\frac12)$, and the zeros on the line contribute at most $4(\log\xi)''(\frac12)$. □

So without any hypothesis (5) holds up to $E(N)$ on both sides; RH removes $E(N)$.

TABLE 1. The off-line contribution: crude bound of Proposition 2.5 against $E(N)$ of Theorem 2.6.

N	10^{10}	10^{15}	10^{20}	10^{25}	10^{27}	10^{30}
$8(3+\sqrt{N})S$	$2.3\cdot10^{-6}$	$7.2\cdot10^{-4}$	0.23	72	720	$2.3\cdot10^{4}$
$E(N)$	$1.0\cdot10^{-8}$	$4.7\cdot10^{-7}$	$2.4\cdot10^{-5}$	$1.3\cdot10^{-3}$	$6.3\cdot10^{-3}$	0.072

The crude bound is below the least margin observed (0.0427, Proposition 2.8) as long as $N<3.5\cdot10^{18}$; the explicit bound (7) stays below it up to $2.3\cdot10^{29}$. Beyond that it grows like $N^{0.35}$ (the bin $0.80<\beta\le0.85$ dominates) and says nothing. This is the familiar fact that verifying RH to a height H, together with zero-density estimates, gives bounds of RH quality only in a range polynomial in H [Jn]; it comes from the zeros, not from the primes. Theorem 2.6 is not a step towards RH: it holds whether or not RH is true, and it is silent for $N>10^{30}$.

2.3. **Verification.**

Proposition 2.8 (the sandwich verified)**.** (1) *For every real N with $1<N\le10^{10}$, $0<M(N)<0.14$; more precisely $M(N)\ge0.0148\log N$ for $1<N\le e^{1/15}$, $M(N)\ge0.0276$ for $e^{1/30}\le N\le10^{10}$, $M(N)\ge0.0429$ for $10^3\le N\le10^{10}$, and $M(N)\le0.1379$ throughout.*

(2) *For every real N with $10^{10}\le N\le10^{12}$, $0.04268\le M(N)\le0.13122$. The least value of M on $(10^3,10^{12}]$ is attained near $N=1.056\cdot10^{11}$.*

(3) $M(N) \geq 0.0415$ *for every real* N *with* $10^8 \leq N \leq 10^{15}$, $M(N) \geq 0.0481$ *for* $10^{15} \leq N \leq 10^{27}$, *and* $M(N) > 0$ *for* $10^{15} \leq N \leq 2.15 \cdot 10^{29}$; $M(N) \leq 0.1377$ *for* $10^8 \leq N \leq 10^{27}$ *and* $M(N) < 0.18484$ *for* $10^8 \leq N \leq 4.28 \cdot 10^{29}$; *so* (5) *holds on* $[10^8, 2.15 \cdot 10^{29}]$. *Here the* 4520 *zeros below* 5000 *are certified: with ball arithmetic,* $N(5000) = 4520$ *and each ordinate used lies within* $9.1 \cdot 10^{-12}$ *of the true one.*

Since $0.14 < 0.18484$, *parts (1) and (3) give the two-sided inequality* (5) *for every real* N *with* $1 < N \leq 2.15 \cdot 10^{29}$; *part (2) checks the range* $[10^{10}, 10^{12}]$ *independently of the zeros.*

Proof. (1) Write $m(L) = M(e^L)$ and $r(L) = R(e^L)$. *Near* $N = 1$, no prime power lies in $(1, 2)$ and $M = R$; elementary bounds on the series in (1) give $r(L) \geq 0.01488L$ for $L \leq \frac{1}{15}$ and $r(L) \leq 0.1016$ for $L \leq \frac{1}{30}$, and on $[\frac{1}{30}, \log 2]$ monotonicity of each piece on 3300 intervals gives $0.0670 \leq M \leq 0.1282$. *From* 2 *to* 10^{10}: between consecutive prime powers $q < q'$, $m(L) = r(L) - 2AL + 2B$ with $A = \sum_{n \leq q} \Lambda(n) n^{-1/2}$ and $B = \sum_{n \leq q} \Lambda(n) n^{-1/2} \log n$ constant, and $m'' = r'' = 2e^{L/2} - 2e^{-L/2}/(e^{2L} - 1) > 0$ for $N \geq 2$, so m is convex there and its minimum on $[\log q, \log q']$ is bounded below by the tangent lines at the two ends; its maximum is at an end. *Arithmetic*: Python's `decimal` at 34 digits, whose $+, -, \times, /, \sqrt{\ }, \log, \exp$ are correctly rounded; π, γ_E, G entered to 48 digits; a margin of 10^{-15} is subtracted from every lower bound and added to every upper bound. The computation ran over all 455 062 595 prime powers below 10^{10}, from a segmented sieve; 8 intervals below $N = 100$ needed subdivision. The least lower bound for $N \geq 10^3$ is 0.042997, on $[34\,186\,367, 34\,186\,499]$, where the minimum of M below 10^{10} is attained; the largest upper bound is 0.137816 at $N = 319\,439$.

(2) Same convexity argument, over all 37 607 992 088 prime powers below 10^{12} (the count equals $\pi(10^{12}) = 37\,607\,912\,018$ plus the 80 070 powers p^k, $k \geq 2$, which checks the sieve). On $[\log q, \log q']$, $\min m \geq \max\big(m(a) - (b-a)\max(0, -m'(a)),\ m(b) - (b-a)\max(0, m'(b))\big)$. The arithmetic is IEEE double precision. The operations $+, -, \times, /, \sqrt{\ }$ are correctly rounded; the library logarithm was checked at run time to be within 0.5 ulp of the correctly rounded value on $2 \cdot 10^4$ test points, and an error of 2 ulp is assumed in the bounds. The sums A and B are accumulated in segments of length $5 \cdot 10^7$, each total exactly rounded (`math.fsum`) and the totals added exactly in rational arithmetic. Every lower bound is lowered, and every upper bound raised, by an explicit bound for all rounding errors (the error in $A(10^{12}) = 1\,999\,997.3576058$ is at most $2.1 \cdot 10^{-9}$, and the error of one evaluation of m is of order 10^{-7}). The least lower bound is 0.042686, near $N = 105\,606\,345\,001$; the largest upper bound on $[10^{10}, 10^{12}]$ is 0.131211, near $N = 742\,864\,907\,663$. The same program reproduces part (1) on $[10^3, 10^{10}]$ (0.042997, 0.137816, 455 062 595 prime powers). The minimum near $1.056 \cdot 10^{11}$ is also visible from the zeros: there the sum over the 4520 zeros below 5000 in part (3) is 0.041711, and the difference 0.00098 is the mean contribution $4 \sum_{\gamma > 5000} \gamma^{-2} = 0.00098$ of the omitted zeros.

(3) Unconditionally $M = \sum_\rho F_L(\gamma_\rho)$; the zeros up to $3 \cdot 10^{12}$ are on the line and contribute ≥ 0, so dropping those not in the list and bounding the rest as in Proposition 2.5, $M(N) \geq 4 \sum_{\gamma \in \text{list}} (1 - \cos \gamma L)/\gamma^2 - 8(3 + \sqrt{N})S$. The sum is Lipschitz with constant $4 \sum 1/\gamma = 14.132$; on a grid of 805 906 points covering $[\log 10^8, \log 10^{15}]$, the least value of the sum on a cell, minus the Lipschitz loss and minus $8(3 + \sqrt{N})S$ at the right end of the cell (at most $7.2 \cdot 10^{-4}$), is 0.04156; it is attained near $N = 1.06 \cdot 10^{11}$, where the sum is 0.041711, the Lipschitz loss $1.41 \cdot 10^{-4}$ and the tail term $7 \cdot 10^{-6}$. For $10^{15} \leq N \leq 10^{30}$ the same computation with $E(N)$ of Theorem 2.6 in place of $8(3 + \sqrt{N})S$, on a grid of step $2 \cdot 10^{-5}$ in L (1 726 941 points), gives the least lower bound 0.048104 up to 10^{27}, 0.047549 up to 10^{28} and 0.034188 up to 10^{29}; the first cell where $W - E \leq 0$ is at $N = 2.159 \cdot 10^{29}$, where $E = 0.0418$ and the sum is near its least value on the range, 0.038908 at $N = 2.18 \cdot 10^{29}$ (Figure 1). For the upper side, the zeros on the line above 5000 contribute at most $8 \sum_{\gamma > 5000} \gamma^{-2}$, and by Stieltjes integration with $N(5000) = 4520$ and Trudgian's bound for $N(t)$ this is at most $8 \cdot 2.446 \cdot 10^{-4} = 1.957 \cdot 10^{-3}$. (As a check, $2(\sum_{\text{list}} 4\gamma^{-2} + 4 \cdot 2.446 \cdot 10^{-4}) = 0.184841$,

against $4(\log\xi)''(\frac{1}{2}) = 0.184840$.) So $M \le W + 1.957\cdot 10^{-3} + E(N)$. The largest upper bound on the cells, the maximum of W plus the Lipschitz gain, is 0.137298 up to 10^{20}, 0.137626 up to 10^{27} and 0.161553 up to 10^{29}, and it first reaches 0.18484 at $N = 4.29\cdot 10^{29}$; the maximum of W on $[10^8, 10^{30}]$ is 0.135632. An error δ in each ordinate changes the sum by at most $4\delta L\sum\gamma^{-2} \le 3.2\,\delta$ for $N \le 10^{30}$. □

Status: (1) is computer-assisted, with correctly rounded arithmetic, and was checked in three independent ways (an extended-precision scan of A and B to 10^{10}, 45-digit values below $2\cdot 10^6$, and an independent sieve on three windows). (2) is computer-assisted in double precision with rigorous error bounds; it relies on the stated accuracy of the library logarithm, which was tested but not proved. (3) uses the zeros below 5000, certified with Arb [Arb] through python-flint: `acb.zeta_zeros` returns balls guaranteed to contain the n-th zeros, isolated by sign changes of Hardy's Z with the count checked by Turing's method. The 4520-th zero lies in $[4999.32968137297, 4999.32968137298]$ and the 4521-st above 5000.2343169313, so the list is complete, and it is accurate to $9.1\cdot 10^{-12}$, which changes W by less than $3\cdot 10^{-11}$. The grid evaluation of W in double precision has error below 10^{-10} (each term $4(1-\cos\gamma L)/\gamma^2$ to a few units in the last place of $\gamma L \le 3.5\cdot 10^5$). Neither reaches the statement for all large N: the primes cost one evaluation per prime power, and the zeros stop near $2\cdot 10^{29}$ by Theorem 2.6.

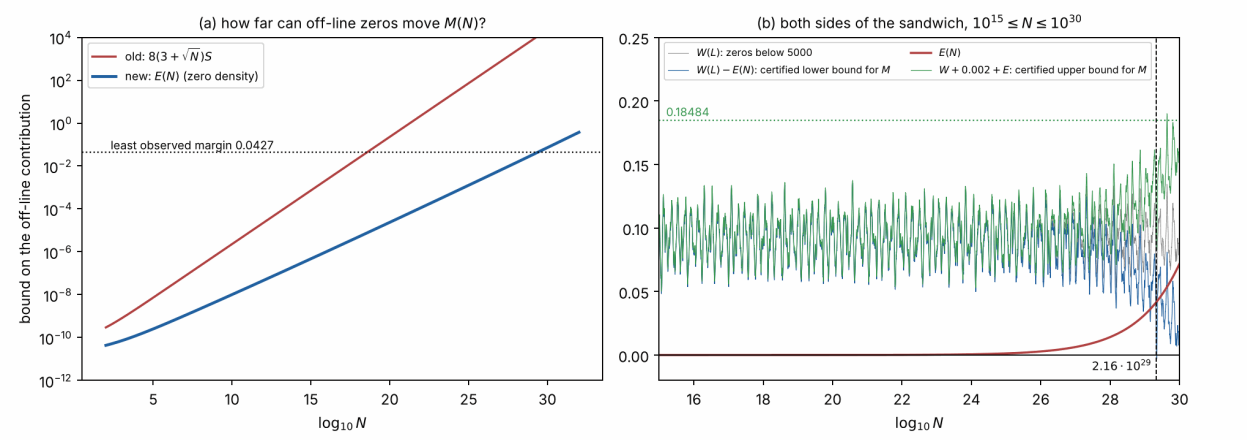

FIGURE 1. (a) The bound on how far off-line zeros could move $M(N)$: the crude bound of Proposition 2.5 and $E(N)$ of Theorem 2.6, which uses the zero density of [KLN]. (b) For $10^{15} \le N \le 10^{30}$: the contribution W of the 4520 zeros below 5000 (grey), the bound $E(N)$ (red) the certified lower bound $W - E$ for M (blue), positive up to $2.16\cdot 10^{29}$ (dashed), and the upper bound $W + 1.96\cdot 10^{-3} + E$ (green), below 0.18484 up to $4.29\cdot 10^{29}$.

The wave. Formula (3) says what the primes must produce: an almost periodic function of $\log N$ with mean $4\sum_{\gamma>0}\gamma^{-2} = 0.0924$, frequencies at the ordinates and amplitudes $4/\gamma^2$. From the 5 762 859 prime powers below 10^8, at 6000 values of $\log N$, M agrees with the right side of (3) truncated at the zeros below 5000 to $5.4\cdot 10^{-5}$; its spectrum in $\log N$, from the primes alone, has peaks at 14.132, 21.023, 25.014, 30.425, ... with heights agreeing with $4/\gamma^2$ to 10^{-5} (Figure 2). Over $4\cdot 10^5$ values of $\log N$ in $[\log 2, \log 10^8]$ the mean is 0.092367 (predicted $2(\log\xi)''(\frac{1}{2}) = 0.092420$) and the standard deviation 0.017219 (predicted $\sqrt{8\sum\gamma^{-4}} = 0.017245$ for independent phases).

A least-squares fit of $m(L) = M(e^L)$ on the first 40 ordinates, from the prime powers below 10^7, gives cosine amplitudes equal to $-4/\gamma_k^2$ to within 0.4% and sine amplitudes below 10^{-5}, as (3) predicts under RH. The decay $4/\gamma^2$ is the regularity of m: it is continuous, and its slope jumps by $-2\Lambda(q)/\sqrt{q}$ at $L = \log q$ (Proposition 3.1).

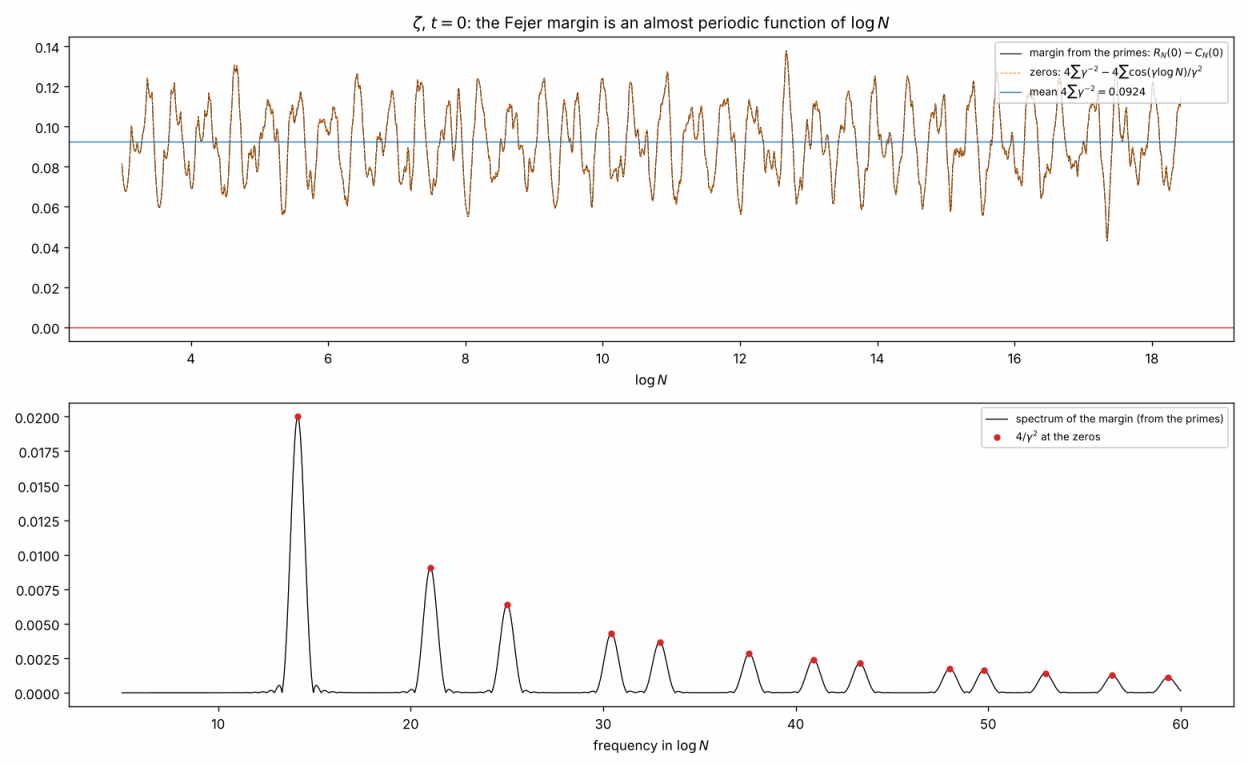

FIGURE 2. Top: the margin $M(N)$ computed from the primes below 10^8 (black) and the right side of (3) from the 4520 zeros below 5000 (dashed); mean 0.0924, least value 0.043. Bottom: its spectrum in the variable $\log N$, from the primes, against $4/\gamma^2$ at the zeros.

Where the difficulty sits. By blocks at $N = 0.999 \cdot 10^8$: $R(N) = 79853.63$, the primes give 79665.63, the prime squares 130.68, the higher powers 57.21, and $M = 0.114$. The squares are controlled without RH; the whole difficulty is in the sum over primes, which has to be known to $O(1)$ out of $8\sqrt{N}$. Without RH the best bound for that sum is of size $\sqrt{N}\exp(-c(\log N)^{3/5}(\log\log N)^{-1/5})$, and closing the gap to $O(1)$ is RH. For the Davenport–Heilbronn function the same construction gives its own window $w_f = 0.78456$; its margin computed from its coefficients up to $2 \cdot 10^7$ stays inside, and the model with its zeros predicts the first exit near $N \approx 2.3 \cdot 10^8$ [CE20]: in one variable the criterion is the simplest, but the slowest to detect an off-line zero.

3. The local minima and the slope

The proof of Proposition 2.8 uses that the margin is convex between consecutive prime powers. This has a sharper consequence: the minima of M are a discrete set, one at most between two prime powers, with an explicit location and value. The criterion of Theorem 2.3, a statement about every real N, becomes a statement about countably many points given by sums over primes. We keep the notation of the proof of Proposition 2.8: $m(L) = M(e^L)$, $r(L) = R(e^L)$, and for consecutive prime powers $q < q'$

$$A_q = \sum_{n \le q} \frac{\Lambda(n)}{\sqrt{n}}, \qquad B_q = \sum_{n \le q} \frac{\Lambda(n)\log n}{\sqrt{n}}, \qquad m(L) = r(L) - 2A_q L + 2B_q \quad (\log q \le L \le \log q').$$

Write $c_0 = \log\pi + \gamma_E + \frac{\pi}{2} + 3\log 2 = 5.3721834$ and $K_0 = -16 + \frac{\pi^2}{2} + 4G = -7.4013354$, so that by (1)

$$r(L) = 8e^{L/2} - c_0 L + K_0 - \tau(L), \qquad \tau(L) = 2\sum_{k\ge 1} \frac{e^{-(2k+\frac12)L}}{(2k+\frac12)^2}. \tag{8}$$

Proposition 3.1 (the local minima of the margin)**.** (1) *On $L > L_c = 0.2824$ ($N > 1.326$), $r''(L) = 2e^{L/2} - 2e^{-L/2}/(e^{2L}-1) > 0$; so m is strictly convex on every interval between consecutive prime powers.*

(2) *At a prime power $q = p^k$ the derivative of m jumps down: $m'(\log q^+) - m'(\log q^-) = -2\Lambda(q)/\sqrt{q} < 0$. Hence no prime power is a local minimum of m.*

(3) *For $N \ge 2$, the local minima of M are exactly the points $N_q^* = e^{L_q^*}$ with $r'(L_q^*) = 2A_q$ and $q < N_q^* < q'$, one at most between two consecutive prime powers; at such a point*

$$M(N_q^*) = 2B_q + K_0 - 8u_q\big(\log u_q - 1\big) + \epsilon_q, \qquad u_q = \frac{2A_q + c_0}{4}, \tag{9}$$

with $N_q^* = u_q^2(1 + O(u_q^{-5}))$ and ϵ_q coming from τ: $|\epsilon_q| \le 0.035$, and $|\epsilon_q| \le 3\cdot 10^{-6}$ for $N_q^* > 100$.

(4) *On $1 < N < 2$, where $M = R$, the only local minimum is at $N = 1.5904$, with $M = 0.0793$, and $M(N) \to 0$ as $N \to 1^+$.*

Proof. (1) Differentiating (8), $r''(L) = 2e^{L/2} - 2\sum_{k\ge1} e^{-(2k+\frac12)L} = 2e^{L/2} - 2e^{-L/2}/(e^{2L}-1)$, which is positive exactly when $e^L(e^{2L}-1) > 1$, that is for $L > 0.2824$. Between prime powers $m'' = r''$. (2) $C(e^L) = 2\sum_n \Lambda(n) n^{-1/2}(L - \log n)_+$ gains the slope $2\Lambda(q)/\sqrt{q}$ at $L = \log q$, and r is smooth. A local minimum at $\log q$ would need $m'(\log q^-) \le 0 \le m'(\log q^+)$, which the jump forbids. (3) By (1) and (2), for $L \ge \log 2$ every local minimum is an interior critical point of a strictly convex piece, so it is the unique solution of $m' = r' - 2A_q = 0$ in that piece, if it lies there. With $r'(L) = 4e^{L/2} - c_0 + \tau_1(L)$, $\tau_1 = -\tau'$, the equation is $4e^{L/2} = 2A_q + c_0 - \tau_1(L)$, so $e^{L^*/2} = u_q - \tau_1(L^*)/4$. Inserting $e^{L^*/2} = u_q$ and $L^* = 2\log u_q$ in $m = r - 2A_qL + 2B_q$ gives $8u_q - 2\log u_q\,(c_0 + 2A_q) + K_0 + 2B_q = 2B_q + K_0 - 8u_q(\log u_q - 1)$; the difference with the exact value comes from τ and τ_1, of order $N^{-5/2}$, and its size is the computed one. (4) On $(1,2)$ the sum over primes is empty, $r'' < 0$ for $N < 1.326$ and $r'' > 0$ after, r' vanishes once on $(1.326, 2)$, at $N = 1.5904$, and $r(0) = 0$ with $r(L) \ge 0.0148L$ near 0 (Proposition 2.8). □

Status: proved; the bounds for ϵ_q and the values in (4) are numerical, from the exact solution of $r'(L) = 2A_q$ by Newton's method. Part (1) is the convexity already used in Proposition 2.8; parts (2) and (3) are new to this paper as far as we know.

Corollary 3.2 (a discrete criterion)**.** *RH holds if and only if $M(N_q^*) \ge 0$ for every prime power q with $q < N_q^* < q'$, that is, if and only if*

$$2B_q \;\ge\; 8u_q\big(\log u_q - 1\big) - K_0 - \epsilon_q \qquad \text{whenever } q < N_q^* < q'. \tag{10}$$

Proof. If RH holds, $M \ge 0$ everywhere by Theorem 2.3. Conversely, assume (10) and suppose $M(N_0) < 0$ for some $N_0 \ge 2$ ($M > 0$ on $(1,2)$ by Proposition 2.8). Let $L_0 = \log N_0$. If $m'(L_0^-) > 0$, follow m to the left: on the convex piece containing L_0 it decreases towards the left end unless it meets a critical point, which would be a local minimum with value $< m(L_0) < 0$; at a prime power the left derivative is larger than the right one by (2), so m keeps decreasing to the left. Either we meet a local minimum with negative value, contradicting (10), or we reach $L = \log 2$ with $m(\log 2) < 0$, contradicting Proposition 2.8. If $m'(L_0^+) < 0$, follow m to the right in the same way: either a local minimum with negative value appears, or m decreases for all $L > L_0$; in the second case $M \le M(N_0)$ for all $N \ge N_0$, so M is bounded above, and Corollary 2.4(3) gives RH, hence

$M \geq 0$, a contradiction. One of the two cases occurs, since $m'(L_0^-) \geq m'(L_0^+)$ and they cannot be both ≤ 0 and ≥ 0 unless L_0 is itself a critical point, which is then a local minimum with negative value. □

So RH is equivalent to countably many inequalities between two explicit sums over the prime powers below q, at the prime powers q for which the stationary point u_q^2 falls before the next prime power. The reduction removes the need to look at every real N; it does not make the inequality easier. Both sides of (10) are of size $4\sqrt{q}\log q$, and their difference is $M(N_q^*) \in [0, 0.185]$ under RH.

Proposition 3.3 (no margin)**.** *If RH holds, then* $\liminf_{N\to\infty} M(N) = 0$*, and the local minima of Proposition 3.1 accumulate at* 0*:* $\liminf M(N_q^*) = 0$ *over the prime powers* q *with* $q < N_q^* < q'$*.*

Proof. Under RH, $m(L) = 4\sum_{\gamma>0}(1-\cos\gamma L)/\gamma^2$ by (3), an absolutely convergent trigonometric series, so m is uniformly almost periodic with mean $2(\log\xi)''(\frac12) = 0.0924$. Given $\varepsilon \in (0, 0.09)$, choose T with $8\sum_{\gamma>T}\gamma^{-2} < \varepsilon/2$. By Dirichlet's simultaneous approximation theorem there are arbitrarily large L_0 with $\|\gamma L_0/2\pi\|$ so small for every $0 < \gamma \leq T$ that $4\sum_{\gamma\leq T}(1-\cos\gamma L_0)/\gamma^2 < \varepsilon/2$; then $m(L_0) < \varepsilon$. Since m is almost periodic with mean $0.0924 > \varepsilon$, the set $\{m > \varepsilon\}$ is relatively dense, so there are $a < L_0 < b$ with $m(a), m(b) > \varepsilon$. The minimum of m on $[a, b]$ is $< \varepsilon$, is not at a or b, and is not at a prime power by Proposition 3.1(2); so it is a local minimum N_q^* with $M(N_q^*) < \varepsilon$, and $N_q^* \geq e^a$ is arbitrarily large. □

Status: proved, conditional on RH, and needs no linear independence of the ordinates. So no proof of $M \geq 0$ can rest on a uniform positive lower bound: under RH the margin is positive without margin, as the least eigenvalues of the windows are in Section 20.

The minima from the primes. Below 10^7 there are 665 134 prime powers and 10 341 of the intervals between them, 1.6%, contain a local minimum. The least is $M(N_q^*) = 0.055041$ at $N_q^* = 3099.71$, between $q = 3089$ and $q' = 3109$; the largest is 0.137816 at $N_q^* = 319\,433.67$, next to the maximum of M found in Proposition 2.8; the mean is 0.091059, against the mean 0.0924 of M. The first ones are $M = 0.0718, 0.0754, 0.0862, 0.0652$ at $N_q^* = 2.448, 3.598, 4.298, 5.934$. The minimum 0.042997 of M below 10^{10}, near $N = 3.4\cdot 10^7$, is necessarily one of these points.

Both sides of (10) are of size $4\sqrt{q}\log q$, and their difference is at most 0.185. In this subsection we remove the common smooth part of the two sums A and B and look at what remains. The criterion becomes a condition on the fluctuation of one weighted Chebyshev sum, B, at the points where the fluctuation of the other, A, vanishes. The identities below are exact and elementary, and they contain no information beyond Corollary 3.2. What they do is separate the two roles that the primes play at a minimum, where it lies and how deep it is, and show the size of the problem.

For real $x \geq 1$ let $A(x) = \sum_{n\leq x}\Lambda(n)n^{-1/2}$ and $B(x) = \sum_{n\leq x}\Lambda(n)n^{-1/2}\log n$, so that $A_q = A(q)$ and $B_q = B(q)$, and define their fluctuations

$$\alpha(x) = A(x) - \Big(2\sqrt{x} - \frac{c_0}{2}\Big), \qquad \beta(x) = B(x) - \big(2\sqrt{x}\log x - 4\sqrt{x}\big). \tag{11}$$

The smooth parts are the integrals $\int_0^x t^{-1/2}\,dt = 2\sqrt{x}$ and $\int_0^x t^{-1/2}\log t\,dt = 2\sqrt{x}\log x - 4\sqrt{x}$, that is, the two sums with $d\psi(t)$ replaced by dt; the constant $-c_0/2$ in α is the one that the closed form (8) imposes. By the prime number theorem and partial summation, $\alpha(x) = o(\sqrt{x})$ and $\beta(x) = o(\sqrt{x}\log x)$; under RH, $\psi(x) - x = O(\sqrt{x}\log^2 x)$ gives $\alpha(x) = O(\log^2 x)$ and $\beta(x) = O(\log^3 x)$. Between two consecutive prime powers A and B are constant, so α decreases, with $d\alpha/d\log x = -\sqrt{x}$, and at a prime power q it jumps up by $\Lambda(q)/\sqrt{q}$. It is a sawtooth (Figure 3(a)).

Proposition 3.4 (the two sums give the slope and the value of the margin)**.** *Let $N \geq 2$ not be a prime power, $L = \log N$, and $\tau_1 = -\tau'$, that is, $\tau_1(L) = 2\sum_{k\geq 1} e^{-(2k+\frac{1}{2})L}/(2k+\frac{1}{2})$. Then*

$$\alpha(N) = \frac{\tau_1(L) - m'(L)}{2}, \tag{12}$$

$$\beta(N) - L\,\alpha(N) = \frac{m(L) - K_0 + \tau(L)}{2}. \tag{13}$$

Equivalently, $M(N) = K_0 + 2\big(\beta(N) - \alpha(N)\log N\big) - \tau(\log N)$ and $\beta(N) = \frac{1}{2}\big(m - Lm' - K_0 + \tau + L\tau_1\big)$.

Proof. Let $q < N < q'$ be the consecutive prime powers around N. On $(\log q, \log q')$ we have $A(e^L) = A_q$, $B(e^L) = B_q$ and $m(L) = r(L) - 2A_qL + 2B_q$. Differentiating (8), $r'(L) = 4e^{L/2} - c_0 + \tau_1(L)$, so

$$m'(L) = 4e^{L/2} - c_0 + \tau_1(L) - 2A_q = \tau_1(L) - 2\alpha(N),$$

which is (12). Write $A_q = 2e^{L/2} - \frac{c_0}{2} + \alpha$ and $B_q = 2Le^{L/2} - 4e^{L/2} + \beta$, with $\alpha = \alpha(N)$, $\beta = \beta(N)$. Inserting these and (8) in $m = r - 2A_qL + 2B_q$ gives

$$m = 8e^{L/2} - c_0L + K_0 - \tau - 2L\Big(2e^{L/2} - \frac{c_0}{2} + \alpha\Big) + 2\big(2Le^{L/2} - 4e^{L/2} + \beta\big) = K_0 - \tau + 2(\beta - L\alpha).$$

The terms in $e^{L/2}$, $Le^{L/2}$ and L cancel, and what remains is (13). □

Status: proved, elementary. So α is, up to the factor $-\frac{1}{2}$ and the trivial-zero term τ_1, the slope of the margin in the variable $\log N$. The combination $\beta - \alpha\log N$ is its value, up to the constant K_0 and the factor 2. The trivial-zero terms are small: $\tau(\log 2) = 0.0615$ and $\tau_1(\log 2) = 0.165$, and both are $O(N^{-5/2})$.

Corollary 3.5 (the minima as zeros of α)**.** *For $N \geq 2$, the local minima of M are the points N_q^* of an interval (q, q') between consecutive prime powers at which the sawtooth α crosses the level $\frac{1}{2}\tau_1(\log N)$, a quantity below 0.083 that tends to 0. At such a point*

$$M(N_q^*) = K_0 + 2\,\beta(N_q^*) - \tau(L_q^*) - L_q^*\,\tau_1(L_q^*), \qquad L_q^* = \log N_q^*. \tag{14}$$

Consequently RH holds if and only if

$$\beta(N_q^*) \;\geq\; -\frac{K_0}{2} + \frac{\tau(L_q^*) + L_q^*\tau_1(L_q^*)}{2} \qquad \text{at every local minimum } N_q^*, \tag{15}$$

where $-K_0/2 = 8 - \pi^2/4 - 2G = 3.7006677$ and the correction is below $1.8\cdot 10^{-5}$ for $N_q^* > 100$.

Proof. By Proposition 3.1(3) the local minima are the solutions of $m'(L) = 0$ inside the intervals $(\log q, \log q')$, and by (12) these are the points where $\alpha = \frac{1}{2}\tau_1$. On $L \geq \log 2$ the function $\alpha - \frac{1}{2}\tau_1$ is strictly decreasing between prime powers, since its derivative in L is $-e^{L/2} + \sum_{k\geq 1} e^{-(2k+\frac{1}{2})L} < 0$; so there is at most one such point in each interval, and α crosses the level from above. At that point (13) gives $m = K_0 - \tau + 2\beta - 2L\alpha = K_0 - \tau + 2\beta - L\tau_1$, which is (14). The equivalence with RH is Corollary 3.2. The bound for the correction is the computed maximum of $\frac{1}{2}(\tau + L\tau_1)$ over the minima with $N_q^* > 100$. □

In words: a minimum of the margin lies where the staircase A comes down to its smooth curve $2\sqrt{x} - \frac{c_0}{2}$, inside an interval free of prime powers. Its depth is twice the excess of the staircase B over its smooth curve at that point, minus 7.40. The large term $\alpha\log N$ in (13) is switched off exactly at the minima, which is why the value of a minimum depends on β alone. One function of the primes locates the minima and the other measures them.

The numbers. On the 10 341 local minima below 10^7 the identities hold to the precision of the computation: $|\alpha(N_q^*) - \frac{1}{2}\tau_1| \leq 5.5\cdot 10^{-12}$, and (14) holds to $2.2\cdot 10^{-10}$. The values of β at the minima lie in $[3.72819, 3.79680]$, with mean 3.74621, against the threshold 3.70067 of (15) (Figure 3(c)). The least value is at $N_q^* = 3099.71$, the least minimum of M below 10^7, and its

excess over the threshold, 0.02752, is half of $M(N_q^*) = 0.05504$. The least value of M below 10^{10}, 0.042997, corresponds to $\beta = 3.72217$.

The minima one after another. Let $N_1^* < N_2^* < \cdots$ be the local minima of M on $[2, \infty)$, and write $L_j = \log N_j^*$, $u_j = \sqrt{N_j^*}$ and $\beta_j = \beta(N_j^*)$.

Proposition 3.6 (the recurrence of the minima)**.** *The first minimum is $N_1^* = 2.44777$, with $\beta_1 = (\log 2)^2/\sqrt{2} - 4u_1(\log u_1 - 1) = 3.796801$ and $M(N_1^*) = 0.071795$. For $j \geq 1$, N_{j+1}^* is the stationary point u^2, $4u = 2A_q + c_0 - \tau_1(2\log u)$, of the first interval (q, q') after N_j^* that contains it. With dA the measure of mass $\Lambda(n)/\sqrt{n}$ at each prime power n,*

$$\beta_{j+1} - \beta_j = \int_{N_j^*}^{N_{j+1}^*} \log \frac{x}{N_j^*} \Big(dA(x) - \frac{dx}{\sqrt{x}} \Big) + \frac{L_j}{2} \big(\tau_1(L_{j+1}) - \tau_1(L_j) \big), \tag{16}$$

and the two measures have equal mass on (N_j^*, N_{j+1}^*) up to $\frac{1}{2}(\tau_1(L_{j+1}) - \tau_1(L_j))$. RH holds if and only if the running sum β_J never falls below the threshold of (15).

Proof. The position follows from Proposition 3.1(3) and Corollary 3.5. Subtracting $\alpha(N_j^*) = \frac{1}{2}\tau_1(L_j)$ from $\alpha(N_{j+1}^*) = \frac{1}{2}\tau_1(L_{j+1})$ gives $A(N_{j+1}^*) - A(N_j^*) - 2(u_{j+1} - u_j) = \frac{1}{2}(\tau_1(L_{j+1}) - \tau_1(L_j))$, which is the statement about the masses. With $S(u) = 4u(\log u - 1)$ we have $\beta_{j+1} - \beta_j = (B(N_{j+1}^*) - B(N_j^*)) - (S(u_{j+1}) - S(u_j))$. Write $\log n = \log(n/N_j^*) + L_j$ in the first difference and $\int x^{-1/2} \log x \, dx = \int x^{-1/2} \log(x/N_j^*)\, dx + L_j \int x^{-1/2} dx$ in the second. The terms in L_j combine to L_j times the mass difference, which gives (16). The last statement is Corollary 3.5. □

Status: proved, elementary. Between two minima, the primes and the smooth density $x^{-1/2}dx$ contribute the same mass, and β changes by the difference of their first moments in $\log(x/N_j^*)$. RH says that, starting from $\beta_1 = 3.796801$, these mass-balanced moments never lose more than $\beta_1 + K_0/2 = 0.0961$. Run forward with the prime powers alone, the recurrence finds the same 10 341 minima below 10^7, with β correct to $1.7 \cdot 10^{-11}$; 5 258 steps go up and 5 082 go down. Like Corollary 3.2, it is exact but contains no information beyond Landau's form of the criterion [In].

4. The margin as an energy

By Proposition 3.4, α is, up to the factor $-\frac{1}{2}$ and the term τ_1, the slope of the margin. Under RH, α is almost periodic in $\log N$, and the margin turns out to be the mean square of its increments. This makes both sides of the sandwich of Corollary 2.4 visible: $M \geq 0$ because it is a mean of squares, and $M \leq 4(\log \xi)''(\frac{1}{2})$ because the variance of α is $(\log \xi)''(\frac{1}{2})$. The same identity turns RH into a statement about energy: α has finite mean square in $\log x$.

Write $\langle f \rangle = \lim_{U\to\infty} \frac{1}{U} \int_0^U f(u)\, du$ when the limit exists. B^2 is the Besicovitch space of almost periodic functions, the closure of trigonometric polynomials in the seminorm $\langle |f|^2 \rangle^{1/2}$. For $f \in B^2$ with Fourier series $\sum_\omega a_\omega e^{i\omega u}$, Parseval's identity $\langle |f|^2 \rangle = \sum |a_\omega|^2$ holds [Bes].

Proposition 4.1 (the margin is the variogram of α)**.** *Assume RH. Then $u \mapsto \alpha(e^u)$ belongs to B^2, has mean 0, and has the Fourier series*

$$\alpha(e^u) \sim \sum_\gamma \frac{i\, m_\gamma}{\gamma}\, e^{i\gamma u} = -2 \sum_{\gamma>0} \frac{m_\gamma}{\gamma} \sin(\gamma u),$$

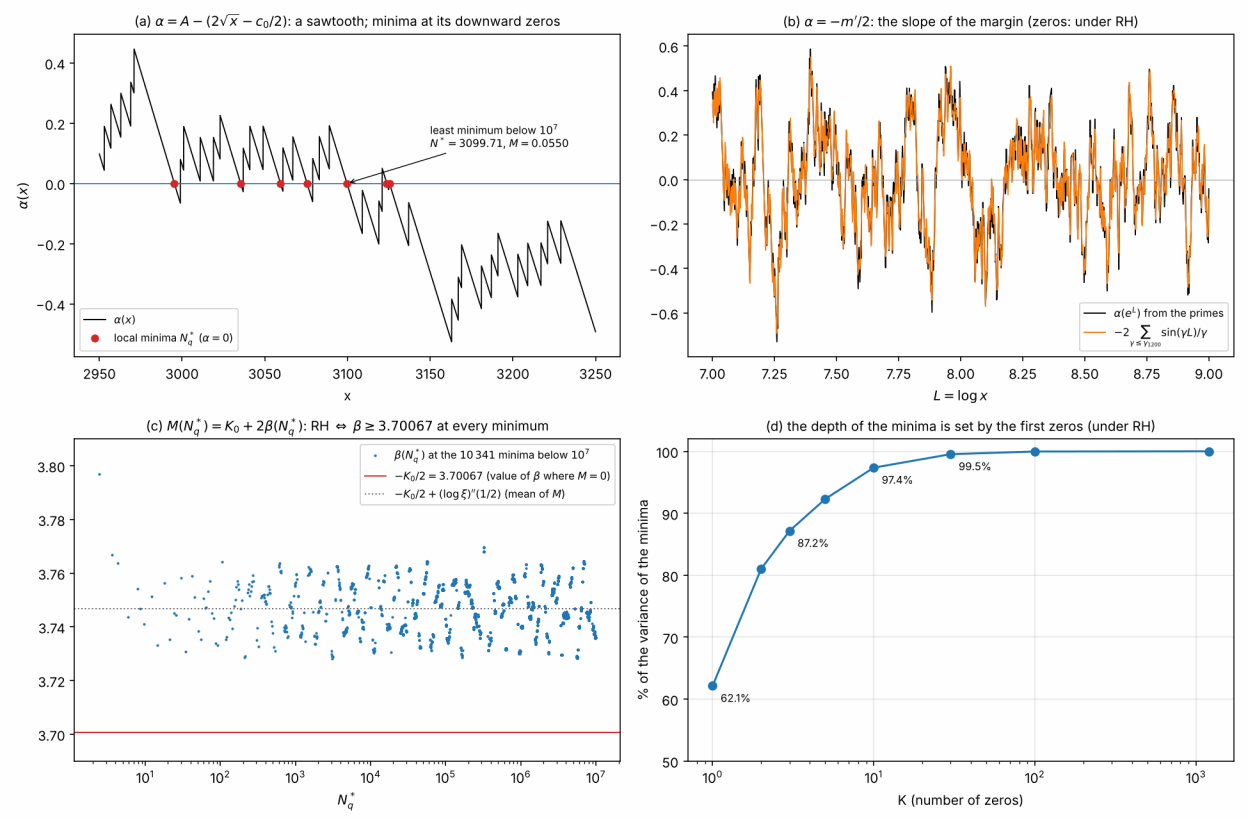

FIGURE 3. The minima read from the slope. (a) The sawtooth $\alpha(x)$ on $[2950, 3250]$; its downward zeros (red) are the local minima N_q^*, among them the least minimum below 10^7. (b) $\alpha(e^L)$ from the primes and $-2\sum_{\gamma\le\gamma_{1200}} \sin(\gamma L)/\gamma$ from the zeros. (c) $\beta(N_q^*)$ at the 10 341 minima below 10^7; the red line is $M = 0$ by (14). (d) Share of the variance of the minima reproduced by the first K zeros.

where γ runs over the distinct ordinates of the zeros (both signs in the first sum) and m_γ is the multiplicity. Consequently, for every $L > 0$,

$$V(L) := \big\langle |\alpha(e^{u+L}) - \alpha(e^u)|^2 \big\rangle = 4 \sum_{\gamma>0} \frac{m_\gamma^2\,(1-\cos\gamma L)}{\gamma^2} \;\ge\; M(e^L), \tag{17}$$

$$\langle \alpha^2 \rangle = 2 \sum_{\gamma>0} \frac{m_\gamma^2}{\gamma^2} \;\ge\; (\log\xi)''(\tfrac12) = 0.046210\ldots, \tag{18}$$

with equality in both, for all L, if and only if all zeros are simple. In that case $M(e^L) = V(L)$ is a mean of squares, and the sandwich $0 \le M \le 4(\log\xi)''(\frac12)$ is the elementary inequality

$$0 \le \langle |f(u+L) - f(u)|^2 \rangle \le 2\langle |f(u+L)|^2 \rangle + 2\langle |f(u)|^2 \rangle = 4\langle |f|^2 \rangle.$$

Proof. Let $\phi(u) = e^{-u/2}(\psi(e^u) - e^u)$. Under RH, $\phi \in B^2$ with Fourier series $-\sum_\rho e^{i\gamma_\rho u}/\rho$ [RuSa, ANS]. This is the explicit formula $\psi(x) = x - \sum_\rho x^\rho/\rho - \log 2\pi - \frac12 \log(1 - x^{-2})$, with the series converging in B^2 because $\sum_\rho |\rho|^{-2} < \infty$.

Partial summation with $A(x) = \sum_{n<x} \Lambda(n) n^{-1/2}$ gives, for $x = e^u$ not a prime power,

$$A(x) = \frac{\psi(x)}{\sqrt{x}} + \frac12 \int_1^x \frac{\psi(t)}{t^{3/2}}\,dt, \qquad \alpha(e^u) = \phi(u) + \frac12 \int_0^u \phi(v)\,dv + \frac{c_0}{2} - 1.$$

Integrating the explicit formula termwise, which is justified as in [In, Ch. IV], gives

$$\frac{1}{2}\int_0^u \phi = -\frac{1}{2}\sum_\rho \frac{e^{i\gamma_\rho u}-1}{i\gamma_\rho\,\rho} + \kappa + o(1)$$

for a constant κ. Here the series converges absolutely and uniformly, since $\sum|\gamma\rho|^{-1}<\infty$. Hence $\alpha(e^u)\in B^2$, and its coefficient at $e^{i\gamma u}$ is

$$-m_\gamma\Big(\frac{1}{\rho}+\frac{1}{2i\gamma\rho}\Big) = -m_\gamma\,\frac{2i\gamma+1}{2i\gamma\rho} = -\frac{m_\gamma}{i\gamma} = \frac{i\,m_\gamma}{\gamma}.$$

The mean is 0. Indeed, for $\Re s>\frac12$, the Laplace transform of $\alpha(e^u)$ is

$$\mathcal{L}[\alpha](s) = -\frac{\zeta'}{\zeta}\big(s+\tfrac12\big)\frac{1}{s} - \frac{2}{s-\frac12} + \frac{c_0}{2s}.$$

By the Abelian theorem [Wi], the mean equals $\lim_{\sigma\to0^+}\sigma\mathcal{L}[\alpha](\sigma) = -\frac{\zeta'}{\zeta}(\frac12)+\frac{c_0}{2}$, which is 0 by the identity $c_0 = 2\frac{\zeta'}{\zeta}(\frac12)$ of Section 2.

Translation and difference preserve B^2. So $\alpha(e^{u+L})-\alpha(e^u)$ has coefficients $\frac{im_\gamma}{\gamma}(e^{i\gamma L}-1)$, and Parseval gives (17), because $|e^{i\gamma L}-1|^2 = 2(1-\cos\gamma L)$ and each $\gamma>0$ appears with its negative. Parseval for α itself gives (18).

For the comparisons, (3) counts zeros with multiplicity, so $M(e^L) = 4\sum_{\gamma>0} m_\gamma(1-\cos\gamma L)/\gamma^2$. Similarly $(\log\xi)''(\frac12) = -\sum_\rho(\frac12-\rho)^{-2} = 2\sum_{\gamma>0}m_\gamma/\gamma^2$. Since $m^2\ge m$, with equality iff $m=1$, both inequalities and the equality cases follow, because $1-\cos\gamma L$ does not vanish identically in L. □

Proposition 4.2 (RH as finite energy)**.** *The following are equivalent:*

(i) *the Riemann hypothesis;*
(ii) *α has bounded mean energy in $\log x$:* $\frac{1}{U}\int_0^U \alpha(e^u)^2\,du = O(1)$*;*
(iii) $\int_0^U \alpha(e^u)^2\,du = e^{o(U)}$*, that is,* $\int_1^X \alpha(x)^2\frac{dx}{x} = X^{o(1)}$.

Proof. (i)⇒(ii) is Proposition 4.1. In substance it is Cramér's theorem that, under RH, $\int_2^X(\psi(x)-x)^2x^{-2}dx\sim c\log X$ [In].

(ii)⇒(iii) is trivial.

(iii)⇒(i): let $I(U)=\int_0^U\alpha(e^u)^2du = e^{o(U)}$ and $\sigma>0$. Integration by parts against dI shows that $\int_0^\infty\alpha(e^u)^2e^{-2\sigma' u}du<\infty$ for every $\sigma'>0$. By Cauchy–Schwarz with $0<\sigma'<\sigma$,

$$\int_0^\infty|\alpha(e^u)|e^{-\sigma u}du \le \Big(\int_0^\infty\alpha^2e^{-2\sigma' u}du\Big)^{1/2}\Big(\int_0^\infty e^{-2(\sigma-\sigma')u}du\Big)^{1/2}<\infty.$$

So $\mathcal{L}[\alpha](s)$ converges absolutely in $\Re s>0$ and is holomorphic there. For $\Re s>\frac12$ it equals $-\frac{\zeta'}{\zeta}(s+\frac12)\frac1s-\frac{2}{s-\frac12}+\frac{c_0}{2s}$, by partial summation from $\sum_n\Lambda(n)n^{-1/2-s} = -\frac{\zeta'}{\zeta}(s+\frac12)$.

By analytic continuation, $-\frac{\zeta'}{\zeta}(w)/(w-\frac12)-2/(w-1)$ is holomorphic in $\Re w>\frac12$. The pole of $-\frac{\zeta'}{\zeta}$ at $w=1$ is cancelled, since its residue 1 divided by $w-\frac12=\frac12$ gives $\frac{2}{w-1}$. So ζ has no zeros in $\Re w>\frac12$. □

Remark 4.3 (what Proposition 4.2 changes, and what it does not). Proposition 4.2 replaces the positivity $M\ge0$ by an energy bound for the slope, $\int_1^X\alpha(x)^2\frac{dx}{x} = X^{o(1)}$. Positivity is then automatic: under (ii), Proposition 4.1 applies and M is a mean of squares. The bound in (iii) is weak: it allows sub-exponential growth in $\log X$. Yet it is equivalent to RH.

Unconditionally, $\alpha(x) \ll x^{\Theta-1/2+\varepsilon}$, where Θ is the supremum of the real parts of the zeros. This gives only $\int_1^X \alpha^2 \frac{dx}{x} \ll X^{2\Theta-1+\varepsilon}$, and the gap between $X^{2\Theta-1}$ and $X^{o(1)}$ is exactly the gap between the known zero-free regions and RH.

By the Paley–Wiener theorem, (ii) says that $\mathcal{L}[\alpha]$ belongs to the Hardy space H^2 of every half-plane $\Re s > \sigma > 0$. Its holomorphy there is the content of RH. Its L^2 size on vertical lines is a mean-square bound for $\frac{\zeta'}{\zeta}$ to the right of the critical line, which is known under RH [MV]. So an unconditional proof of (iii) would have to control $\frac{\zeta'}{\zeta}$ in mean square on every line $\Re w = \frac{1}{2} + \sigma$ and show that it extends holomorphically. No method we know does the first without the second. We do not claim progress on it.

Numerically, with the prime powers below 10^8, $M(e^L)$ and the variogram of α agree to three digits at $L = 0.2, 0.6, \dots, 3$, and the empirical variance of $\alpha(e^u)$ is 0.04625, against $(\log \xi)''(\frac{1}{2}) = 0.04621$.

4.1. **The energy programme: dyadic blocks and averages in L.** Proposition 4.2 reduces RH to an energy bound for α. This subsection follows the two natural ways of attacking that bound: in dyadic blocks, using the recurrence of the minima, and by averaging the margin in L. Each route gives exact statements and a clear account of where it stops. Neither route proves RH.

An exact identity: the energy of the slope is a C-transform of the margin. Write $m(L) = M(e^L)$, so that $m' = \tau_1 - 2\alpha$ by Proposition 3.4.

Theorem 4.4 (energy identity)**.** *For $1 < X < Y$, neither of them a prime power, and without any hypothesis,*

$$\int_{\log X}^{\log Y} m'(L)^2 \, dL = 2 \sum_{X<q<Y} \frac{\Lambda(q)}{\sqrt{q}} \, M(q) \; - \; \int_{\log X}^{\log Y} m(L)\, r''(L)\, dL \; + \; \Big[m\, m' \Big]_{\log X}^{\log Y},$$

where q runs over prime powers and $r''(L) = 2e^{L/2} - 2\sum_{k\ge1} e^{-(2k+\frac{1}{2})L}$.

Proof. Between consecutive prime powers, $m = r - 2A_q L + 2B_q$ is smooth with $m'' = r''$. At a prime power q, m is continuous, and $m' = r' - 2A$ jumps by $-2\Lambda(q)/\sqrt{q}$. On each gap, $\int m'^2 = [mm'] - \int m\, r''$. Summing over the gaps, the interior boundary terms combine into $\sum_q m(\log q)\big(m'(\log q^-) - m'(\log q^+)\big) = 2\sum_q \frac{\Lambda(q)}{\sqrt{q}} M(q)$. □

The prime measure $\frac{\Lambda(q)}{\sqrt{q}}\delta_{\log q}$ and its smooth density $e^{L/2}dL = \frac{1}{2}r''dL + O(e^{-5L/2})$ both look at the margin. Theorem 4.4 says that the energy of the slope is twice the difference between what the primes see and what the continuum sees, which is the C-transform of M with the weight $\Lambda(n)/\sqrt{n}$.

Corollary 4.5 (the rate is the window width)**.** *Assume RH and that the zeros are simple. Then*

$$2\sum_{q<Y} \frac{\Lambda(q)}{\sqrt{q}} M(q) - \int_0^{\log Y} m\, r''\, dL = 4\langle \alpha^2 \rangle \log Y + o(\log Y) + O\big(1 + |\alpha(Y)|\big),$$

with $4\langle\alpha^2\rangle = 4(\log\xi)''(\frac{1}{2}) = 0.18484\ldots$ In particular, the left side divided by $\log Y$ tends to 0.18484 along every sequence with $\alpha(Y) = o(\log Y)$, and so on a set of Y of logarithmic density 1. In words: per unit of $\log N$, the primes see the margin more than the continuum does by exactly the width of the window of Corollary 2.4.

Proof. Under RH, $0 \le m \le 0.18484$ (Corollary 2.4), and $m' = \tau_1 - 2\alpha$ with $\tau_1(L) = O(e^{-5L/2})$. So the boundary term of Theorem 4.4 is $O(1 + |\alpha(Y)|)$, and $\int_0^{\log Y} m'^2 = 4\int_0^{\log Y} \alpha^2 + O(1) = 4\langle\alpha^2\rangle \log Y + o(\log Y)$ by Proposition 4.1. For simple zeros, (18) gives the value. Finally, $\alpha(e^u)$ has a limiting distribution under RH [ANS], so $|\alpha(Y)| \le K$ on a set of logarithmic density tending to 1 as $K \to \infty$. (The boundary term cannot be dropped altogether: α is unbounded, by Littlewood's theorem.) □

Theorem 4.6 (a lower bound for the energy, without hypothesis)**.** *Unconditionally,*

$$\limsup_{X\to\infty} \frac{1}{\log X}\int_1^X \alpha(x)^2\frac{dx}{x} \;\geq\; (\log\xi)''(\tfrac{1}{2}) = 0.046210\ldots,$$

and the limit exists and equals $(\log\xi)''(\frac{1}{2})$ *if and only if RH holds and all zeros are simple.*

Proof. If RH fails, Proposition 4.2 shows that $\int_1^X \alpha^2\frac{dx}{x}$ is not $X^{o(1)}$, so the lim sup is infinite. If RH holds, Proposition 4.1 gives the limit $\langle\alpha^2\rangle = 2\sum_{\gamma>0} m_\gamma^2/\gamma^2$. This is $\geq 2\sum_{\gamma>0} m_\gamma/\gamma^2 = (\log\xi)''(\frac{1}{2})$, with equality iff every $m_\gamma = 1$. □

So the energy of the slope is a computable test of "RH and simple zeros" together, with an excess $2\sum(m_\gamma^2 - m_\gamma)/\gamma^2$ for multiple zeros. The empirical value is 0.04625 up to 10^8, and 0.0465 from the block energies to 10^{10} below.

Dyadic blocks. Let $E_k = \int_{k\log 2}^{(k+1)\log 2} \alpha(e^u)^2 du$ be the energy in the block $[2^k, 2^{k+1})$; by Proposition 4.2, RH is equivalent to $E_k = 2^{o(k)}$. We computed E_k and $W_k = \int m'^2$ for every block below 10^{10}, exactly between consecutive prime powers, from all 455 062 595 prime powers (Figure 4(a)). The block energies do not grow: the mean of $E_k/\log 2$ over $6 \leq k \leq 32$ is 0.0465, against 0.0462, and $W_k/\log 2$ has mean 0.1858, against the rate 0.1848 of Corollary 4.5. Theorem 4.4 holds block by block to $4\cdot 10^{-11}$ for $k \geq 15$.

Why this does not give a proof: on a block, $\alpha(x) = \alpha(2^k) + D_k(x)$, where D_k depends only on the prime powers of the block. Its energy is the energy of the normalized error $\big(\psi(x) - \psi(2^k) - (x-2^k)\big)/2^{k/2}$ of the prime number theorem on an interval as long as its position. A bound $2^{o(k)}$ for it is already RH. The recurrence of Proposition 3.6 moves α exactly from one prime power to the next and reproduces D_k, so it gives nothing more. Finally, no computation can see an off-line zero. Such a zero has $\gamma > T_0$, so its term in α is at most $10^{-12}x^{\beta-1/2}$, which reaches the size 0.2 of α only beyond $x = e^{54}$ even for $\beta = 1$.

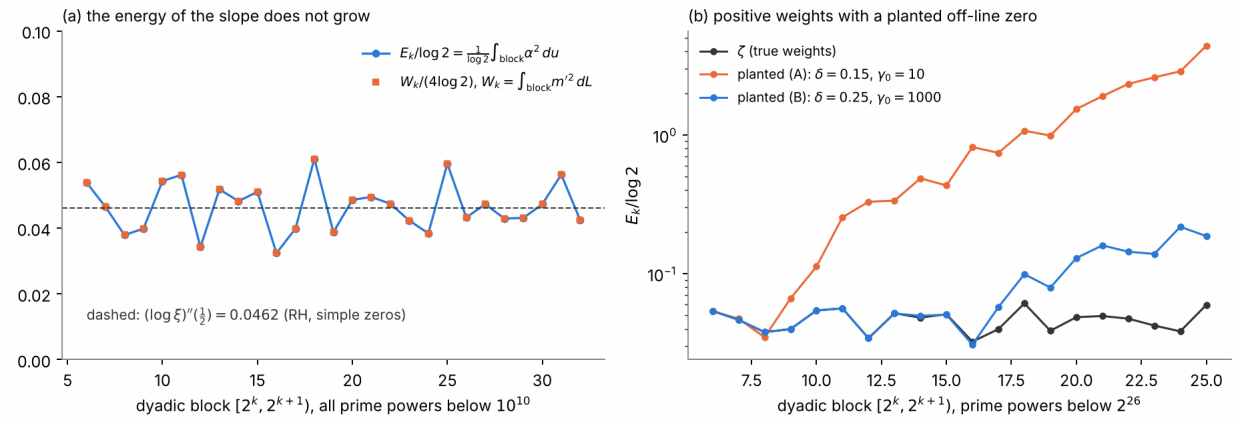

FIGURE 4. (a) Energy per dyadic block of the slope, $E_k/\log 2$, and of the margin's derivative, $W_k/(4\log 2)$, from all prime powers below 10^{10}, against $(\log\xi)''(\frac{1}{2})$. (b) The same block energy, below 2^{26}, for the true weights and for the weights (43) with a planted zero: (A) $\delta = 0.15$, $\gamma_0 = 10$; (B) $\delta = 0.25$, $\gamma_0 = 1000$ (Theorem 19.1).

Averages in L. Averaging the margin in L is a second-order Riesz mean:

$$M_2(N) := \int_0^{\log N} M(e^L)\,dL = R_2(N) - \sum_{n<N}\frac{\Lambda(n)}{\sqrt{n}}\log^2\frac{N}{n}, \qquad R_2(N) = \int_0^{\log N} R(e^L)\,dL.$$

Under RH and simple zeros, $M_2(N) = 4\sum_{\gamma>0}\gamma^{-2}\log N - 4\sum_{\gamma>0}\sin(\gamma\log N)/\gamma^3 = 2\langle\alpha^2\rangle\log N + O(1)$, which is the identity $\langle M\rangle = 2\langle\alpha^2\rangle$.

Proposition 4.7 (averaging in L)**.** *The following are equivalent:*

(i) *RH;*
(ii) *for every* $\varepsilon > 0$, $M_2(N) \geq -C_\varepsilon N^\varepsilon$*;*
(iii) $|M_2(N)| = N^{o(1)}$.

Proof. (i)⇒(ii),(iii): under RH, $0 \leq M \leq 0.18484$ by Corollary 2.4, so $0 \leq M_2(N) \leq 0.18484\log N$.

(iii)⇒(i): integrating by parts, $\int_0^\infty M(e^L)e^{-sL}dL = s\int_0^\infty M_2(e^L)e^{-sL}dL$ converges for $\Re s > 0$, so the function $\mathcal{G}$ of the proof of Theorem 2.3 is holomorphic in $\Re s > 0$. But each zero $\rho = \frac{1}{2}+\delta+it$ with $\delta > 0$ gives $\mathcal{G}$ a pole at $s = \delta + it$.

(ii)⇒(i): apply Landau's theorem [Wi] to $M_2(e^L)+C_\varepsilon e^{\varepsilon L} \geq 0$. Its Laplace transform is $\mathcal{G}(s)/s + C_\varepsilon/(s-\varepsilon)$, which is holomorphic at every real point $s > \varepsilon$, because zeros of ζ are not real and $\zeta(\sigma) \neq 0$ for $\frac{1}{2} < \sigma < 1$. So the integral converges in $\Re s > \varepsilon$, and $\mathcal{G}$ has no pole there. □

Proposition 4.7 is weaker than Theorem 2.3 in its hypothesis, since it needs only a one-sided bound on an average and allows N^ε. It is equivalent all the same, and it is the same statement as Proposition 4.2, because $M_2 = 2\langle\alpha^2\rangle\log N + O(1)$ under RH. The least eigenvalue of a window (Section 20) is an infimum of such energies over test functions, so it enters through the same door.

Remark 4.8 (where both routes stop). Every statement in this subsection that is equivalent to RH is equivalent through a Landau-type argument. The analytic content is always the same: the Laplace transform of m, of α, or of the kernel's remainder is $-\frac{\zeta'}{\zeta}$ shifted by $\frac{1}{2}$, divided by a polynomial in s, plus explicit terms. Each bound in the programme (energy, average, one-sided bound) is equivalent to the holomorphy of $\frac{\zeta'}{\zeta}$ in $\Re w > \frac{1}{2}$.

The two routes make this concrete. The block computation shows that "local" in $\log x$ means a prime number theorem on intervals as long as their position. Averaging shows that averaging changes the order of the Riesz mean but not the abscissa. Theorem 4.4 is the one exact link between the energy and the primes. To be useful it would need an independent lower bound for the continuum term or an upper bound for the prime term, and each of these is again a statement about $\frac{\zeta'}{\zeta}$ to the right of the line.

We record the programme as a reformulation with exact identities and verified numerics. It does not prove RH.

5. The margin at large N: a probabilistic projection

Proposition 3.3 says that under RH the margin comes arbitrarily close to 0. This section asks how fast. The answer is a projection, not a theorem: it assumes RH and the linear independence of the ordinates (LI), as in the work of Rubinstein and Sarnak [RuSa] and of Akbary, Ng and Shahabi [ANS] on limiting distributions.

The limiting distribution. Under RH, (3) writes $M(e^L) = \sum_{\gamma>0} a_\gamma(1-\cos\gamma L)$ with $a_\gamma = 4/\gamma^2$. Under LI, by the Kronecker–Weyl theorem the phases γL modulo 2π become equidistributed on the infinite torus as $L \to \infty$, so the proportion of $L \in [0,T]$ with $M(e^L) < x$ tends to $\mathbb{P}(X < x)$, where

$$X = \sum_{\gamma>0} a_\gamma(1-\cos\theta_\gamma), \qquad \mathbb{E}\,e^{itX} = \prod_{\gamma>0} e^{ita_\gamma}J_0(ta_\gamma), \tag{19}$$

with independent θ_γ uniform on $[0,2\pi)$. Its mean is $2(\log\xi)''(\frac{1}{2}) = 0.0924$ and its standard deviation $(8\sum\gamma^{-4})^{1/2} = 0.01724$. Fourier inversion of (19) with the 4520 certified zeros (the rest contribute their mean 0.00098 and a variance below $3\cdot10^{-11}$) gives its density, a smooth bump on $[0.04, 0.145]$

with two shallow maxima. For the lower tail we use the saddle point: with $K(s) = \log \mathbb{E}e^{sX} = \sum_\gamma \big(sa_\gamma + \log I_0(sa_\gamma)\big)$ and $K'(\hat{s}) = x$, the Lugannani–Rice formula gives

$$\mathbb{P}(X < 0.04) = 3.7 \cdot 10^{-6}, \quad \mathbb{P}(X < 0.03) = 1.0 \cdot 10^{-11},$$
$$\mathbb{P}(X < 0.02) = 4.2 \cdot 10^{-27}, \quad \mathbb{P}(X < 0.01) = 1.7 \cdot 10^{-97},$$

in agreement with Fourier inversion to 0.6% where both apply ($x \geq 0.03$). Numerically $\log \mathbb{P}(X < x) \approx -C/x^2$: halving x multiplies the exponent by about 3.7.

How often the margin dips. The rate of downcrossings of a level x, per unit of $L = \log N$, is given by Rice's formula $\nu(x) = f_X(x)\, \mathbb{E}\big[|X'| \,\big|\, X = x\big]/2$, with $X' = \sum_\gamma \gamma a_\gamma \sin\theta_\gamma$. Conditioning on $X = x$ by exponential tilting turns each θ_γ into a von Mises variable of concentration $\kappa_\gamma = |\hat{s}|a_\gamma$, for which $\mathbb{E}\sin^2\theta = I_1(\kappa)/(\kappa I_0(\kappa))$; taking X' conditionally Gaussian,

$$\nu(x) = \frac{f_X(x)}{\sqrt{2\pi}} \Big(\sum_{\gamma > 0} \frac{16}{\gamma^2} \frac{I_1(\kappa_\gamma)}{\kappa_\gamma I_0(\kappa_\gamma)} \Big)^{1/2}, \qquad f_X(x) \approx \frac{e^{K(\hat{s}) - \hat{s}x}}{\sqrt{2\pi K''(\hat{s})}}. \tag{20}$$

We tested the model against the function itself, computed from the zeros on $7 \leq L \leq 3000$: for $x = 0.060, 0.055, 0.050, 0.045$ the observed proportion of time below x is $2.12 \cdot 10^{-2}$, $6.1 \cdot 10^{-3}$, $1.17 \cdot 10^{-3}$, $1.2 \cdot 10^{-4}$ against $2.14 \cdot 10^{-2}$, $6.3 \cdot 10^{-3}$, $1.16 \cdot 10^{-3}$, $1.1 \cdot 10^{-4}$, and the observed number of excursions per unit of L on $[1000, 3000]$ is 0.686, 0.259, 0.068, 0.0085 against 0.683, 0.262, 0.065, 0.0083 from (20).

The record low. If dips below x arrive as a Poisson process of rate $\nu(x)$, then $\mathbb{P}\big(\min_{N' \leq N} M(N') \geq x\big) \approx e^{-\nu(x)\log N}$. The median value of $\log N$ at which M first drops below x is $\log 2/\nu(x)$:

x	0.040	0.035	0.030	0.025	0.020	0.015	0.010
$1/x$	25	28.6	33.3	40	50	66.7	100
median $\log_{10} N$	$7.5 \cdot 10^2$	$7.9 \cdot 10^4$	$1.3 \cdot 10^8$	$2.7 \cdot 10^{13}$	$1.2 \cdot 10^{23}$	$5.8 \cdot 10^{42}$	$6.4 \cdot 10^{92}$

Since $\log \mathbb{P}(X < x) \approx -C/x^2$, the record low decreases like $(\log\log N)^{-1/2}$. Fitting the median gives

$$\max_{10^3 \leq N' \leq N} \frac{1}{M(N')} \approx 10.4 \sqrt{\log_{10}\log N} \qquad (N \to \infty), \tag{21}$$

which reproduces the median to within 3% for $\log_{10}\log N \geq 13$; for smaller N the median lies above the law (Figure 5(b)).

The actual record, under RH only. Under RH alone, $M(e^L) = W(L) + T(L)$, where W is the sum (3) over the 4520 certified zeros and $0 \leq T \leq 2 \cdot 0.00098$, T having mean 0.00098. We followed the record low of $W + 0.00098$ for $7 \leq L \leq 10^5$: a coarse pass with the first 100 zeros, whose partial sum is a pointwise lower bound for W, selects the candidates, which are then refined with all 4520 zeros. The records are

$$0.0430 \ (N = 3.4 \cdot 10^7), \quad 0.04269 \ (1.06 \cdot 10^{11}), \quad 0.0399 \ (2.2 \cdot 10^{29}),$$
$$0.0385 \ (10^{1521}), \quad 0.0369 \ (10^{8274}), \quad 0.0362 \ (10^{24461}),$$

the first two agreeing with the primes (Proposition 2.8) to 10^{-5}. The last three fall inside the 10–90% band of the model; the first three come early (each had probability 2–6% of arriving so soon), so the region $N < 10^{30}$ is somewhat atypically low.

Remark 5.1 (status). Formula (19) for the limiting distribution is a theorem under RH and LI; the saddle-point values are numerical approximations checked against Fourier inversion; the rate (20), the Poisson approximation and the law (21) are heuristics, checked against the function itself where it can be computed. Nothing here bears on the truth of RH: the projection describes how the margin behaves if RH holds. Read through the reciprocal, it says that $1/M$ is positive and

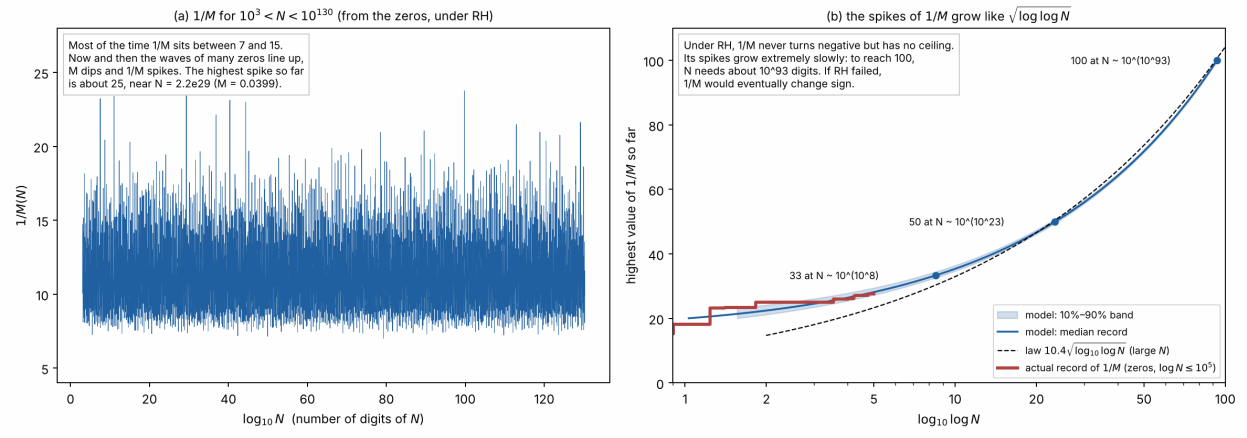

FIGURE 5. The reciprocal $1/M$. (a) $1/M(N)$ for $10^3 < N < 10^{130}$, from the certified zeros under RH: it stays between 7 and 15 most of the time and spikes when many waves are in phase. (b) The record of $1/M$ against $\log_{10}\log N$: actual records for $\log N \le 10^5$ (red), the median record and 10–90% band of the model (20) under RH and LI (blue), and the law (21) (dashed).

unbounded, and that its spikes grow like $\sqrt{\log\log N}$, so that $1/M$ reaches 100 only for N with about 10^{93} digits. If RH failed, $1/M$ would eventually change sign.

6. The walk of the primes against the curve

Proposition 3.1 can be read as a geometric reformulation of the whole problem. Write $P_q = (A_q, B_q)$ for the pair of prime sums at the prime power q, and

$$D_q = 2B_q + K_0 - \Phi(u_q), \qquad \Phi(u) = 8u(\log u - 1), \qquad u_q = \frac{2A_q + c_0}{4} = \sqrt{q} + \frac{\alpha(q)}{2}, \tag{22}$$

the last equality because $A_q = 2\sqrt{q} + \alpha(q) - \frac{c_0}{2}$. By (9), D_q is, up to $|\epsilon_q| \le 3\cdot 10^{-6}$, the value of the local minimum of M attached to the block $[q, q')$ when that minimum is interior. Every local minimum of M on $N \ge 2$ is of this form (Proposition 3.1(3)), and a sign change of the continuous function M, positive near $N = 1$, would happen at a negative local minimum. So RH is the statement that *the walk $q \mapsto P_q$, pushed forward one prime power at a time, never crosses the fixed curve $2B + K_0 = \Phi\big(\frac{2A+c_0}{4}\big)$* beyond some point; D_q is the height of the walk above the curve (Figure 6).

6.1. The exact step.

Proposition 6.1 (the step of the walk)**.** *Let $q < q'$ be consecutive prime powers, $\lambda = \Lambda(q')q'^{-1/2}$. Then, for some $\xi \in (u_q, u_{q'})$,*

$$D_{q'} - D_q = 2\lambda\Big(\log\frac{q'}{q} - 2\log\Big(1 + \frac{\alpha(q)}{2\sqrt{q}}\Big)\Big) - \frac{\lambda^2}{\xi}. \tag{23}$$

In particular the step is $2\lambda\log(q'/N_q^*) + O(\lambda^2/\sqrt{q})$: the prime power q' lowers *the floor when it arrives before the pending minimum N_q^*, and* raises *it when it arrives after; and no single prime power moves the floor appreciably: for $q \ge 10^6$,*

$$|D_{q'} - D_q| \le \frac{2\log q'}{\sqrt{q'}}\Big(\log\frac{q'}{q} + \frac{|\alpha(q)|+1}{\sqrt{q}}\Big) \le \frac{2\log q\,(|\alpha(q)|+2)}{q}.$$

Proof. $A_{q'} = A_q + \lambda$, $B_{q'} = B_q + \lambda \log q'$, $u_{q'} = u_q + \frac{\lambda}{2}$. So $D_{q'} - D_q = 2\lambda \log q' - \big(\Phi(u_q + \frac{\lambda}{2}) - \Phi(u_q)\big)$, and by Taylor with $\Phi'(u) = 8 \log u$, $\Phi''(u) = 8/u$, the bracket is $4\lambda \log u_q + \lambda^2/\xi$. Now $2 \log u_q = \log q + 2 \log(1 + \frac{\alpha(q)}{2\sqrt{q}})$, which gives (23); and $2 \log u_q = \log N_q^* + O(u_q^{-5})$ by Proposition 3.1(3). The bound follows from $\Lambda(q') \le \log q'$ and the Bertrand bound $q' < 2q$, and for $q \ge 10^6$ from the known prime gaps ($q' - q < 1500$ below $4 \cdot 10^{18}$). □

Identity (23) was checked against the exact D_q over 10^5-long stretches at $q \sim 10^6$, agreeing to $1.9 \cdot 10^{-11}$; the largest single step there is $1.6 \cdot 10^{-5}$, falling to $9 \cdot 10^{-8}$ at $q \sim 10^8$ and $3.3 \cdot 10^{-9}$ at $q \sim 10^9$. The two terms of the step are the two clocks of the problem: the gap $\log(q'/q)$, what the primes do, against $\alpha(q)/\sqrt{q}$, where the running account stands. Their accumulated difference is $-2 \int \alpha \, d \log x$ (Section 3), so no individual prime matters and only the *area* under the slope does.

6.2. **What it takes to cross.**

Proposition 6.2 (local crossing bound)**.** *Let $x < y$ and suppose M changes sign in $[x, y]$, with $M(x) > 0$. Then*

$$\max_{t \in [x,y]} \alpha(t) \;\ge\; \frac{M(x)}{2 \log(y/x)},$$

and the maximum is attained at a prime power of $[x, y]$. Moreover $\alpha(t) - \alpha(x)$ depends only on the prime powers in $(x, t]$: $\alpha(t) - \alpha(x) = \sum_{x < n \le t} \Lambda(n) n^{-1/2} - 2(\sqrt{t} - \sqrt{x})$.

Proof. With $m(L) = M(e^L)$, $m' = -2\alpha(e^L) + \tau_1(L)$ and $\tau_1 > 0$ (Section 3), so for $L \in [\log x, \log y]$, $m(L) \ge m(\log x) - 2(L - \log x) \max_{[x,y]} \alpha \ge M(x) - 2 \log(y/x) \max \alpha$. If M changes sign the left side is < 0 somewhere. Between prime powers $\alpha(t) = A - 2\sqrt{t} + \frac{c_0}{2}$ decreases, and it jumps up at each prime power, so its maximum on a closed block sits at a prime power. □

TABLE 2. What a first sign change in a dyadic block $[x, 2x]$ would require, from Proposition 6.2 with the verified lower bounds for $M(x)$: a surplus $\alpha \ge M(x)/(2 \log 2)$, i.e. about $\alpha \sqrt{x}/\log x$ prime powers above the smooth count, present at some prime power of the block.

x	10^{12}	10^{20}	10^{27}	10^{29}
$M(x) \ge$	0.0427	0.0481	0.0481	0.0342
needs $\alpha \ge$	0.031	0.035	0.035	0.025
extra prime powers	$1.1 \cdot 10^3$	$7.5 \cdot 10^6$	$1.8 \cdot 10^{10}$	$1.2 \cdot 10^{11}$

Status: both propositions are proved, and they answer a natural question – can one added prime, or one added zero, flip the sign of M? – precisely. A zero on the line only adds the nonnegative window $4(1 - \cos \gamma L)/\gamma^2$; an off-line zero flips the sign for large N, but that is RH itself; a single prime power beyond 10^6 moves the floor by less than $3 \cdot 10^{-5}$. What Proposition 6.2 does *not* give is RH: the required surplus, about 0.15 standard deviations of α (whose mean square is $(\log \xi)''(\frac{1}{2})$ by Theorem 4.6), occurs in a positive proportion of blocks. The content is the exact bookkeeping: a sign change cannot start in a block where the primes run behind their smooth count, and the surplus must be present at a prime power, locally auditable through the second identity of Proposition 6.2. The obstruction to converting this into a proof is the one of Theorem 19.1: the bound uses $|\alpha|$, not the identity $\Lambda = \mu * \log$.

6.3. **Anatomy of the close approaches.** The walk has come closest to the curve at the record minima of Section 5. Two complementary mechanisms show up, one in the primes and one in the zeros.

In the primes. At the record $N = 34\,186\,433$ ($M = 0.042997$), the surplus was assembled one octave earlier: on $(8.5 \cdot 10^6,\, 1.7 \cdot 10^7]$ the weighted count of prime powers exceeds $2\sqrt{x}$ by $+0.223$ in the slope – about 49 extra prime powers – while the neighbouring octaves run near balance ($+0.017$) or behind (-0.193, -0.203). The final descent is short: M crosses its mean 0.0924 for the last time at $N = 3.09 \cdot 10^7$ and reaches the record a factor 1.106 later, paying $2\int \alpha\, d\log x = 0.0504$ with α averaging $+0.25$ on the stretch; the bottom is reached exactly when the surplus is spent, $\alpha(N_1) = 0.0149$ (the minimum is the zero of α, Proposition 3.1). Figure 7(a) shows α from the primes against $\alpha = -2\sum_\gamma \sin(\gamma \log x)/\gamma$ from the 4520 certified zeros: the two are indistinguishable at this scale. At the next record, $N = 1.056 \cdot 10^{11}$, the primes (Proposition 2.8(2)) give 0.042686 and the zeros give 0.042689: the same event seen from both sides, to $3 \cdot 10^{-6}$.

In the zeros. At every record the low-lying zeros conspire: of the first 20 phases $\cos(\gamma_k \log N)$, between 13 and 16 exceed $\frac{1}{2}$ simultaneously (Figure 7(b)), the first 10 zeros carry 74–92% of the depth $0.0924 - M$, and the first 100 carry 99%. Independent phases would align like this with probability of order $3^{-14} \approx 10^{-7}$ per sample, which is the rarity that the rates of Section 5 quantify; the next deeper conspiracies are predicted near $N \sim 10^{10^5}$.

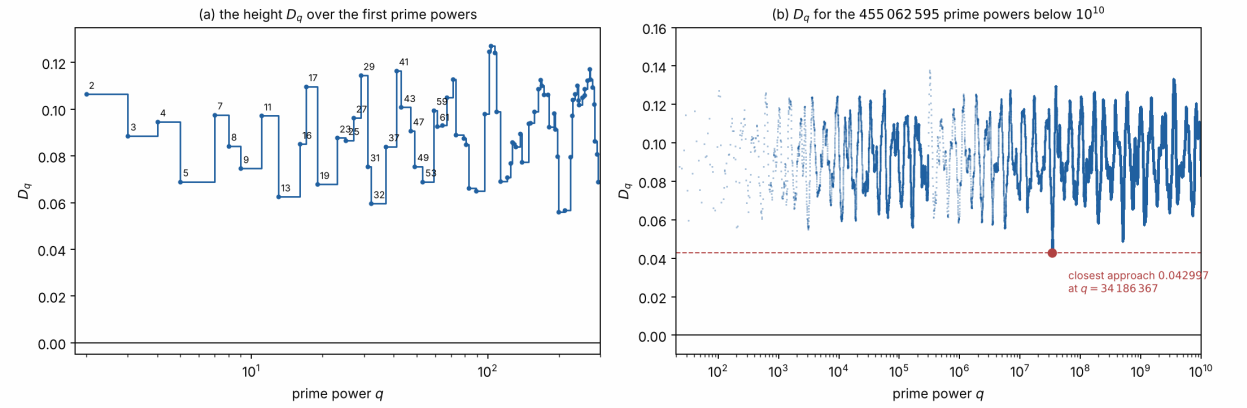

FIGURE 6. The walk against the curve. (a) The height D_q of (22) over the first prime powers: each prime power moves the point, down when it arrives before the pending minimum N_q^*, up when after (23). (b) D_q for all prime powers below 10^{10}; the closest approach 0.042997 at $q = 34\,186\,367$ is the global minimum of M beyond $N = 10^3$.

7. The margin as a one-sided mean of the prime-counting error

The attempt to push the verification beyond the computed range with explicit prime-counting bounds fails, and it fails for a reason worth recording exactly. Let

$$K_N(t) = t^{-3/2}\Big(1+\tfrac{1}{2}\log\tfrac{N}{t}\Big), \qquad E^*(t) = \psi(t) - t + \log 2\pi + \tfrac{1}{2}\log(1-t^{-2}), \qquad \widehat{B} = 1 + \frac{\gamma}{2} - \frac{1}{2}\log 4\pi,$$

so that E^* is the remainder of ψ against the full smooth model of the explicit formula, $\widehat{B} = \sum_\rho \rho^{-1} = 0.0230957\ldots$ is the Hadamard constant of ξ, and $\int_1^N K_N(t)\,dt = \log N$ exactly, since $-t^{-1/2}\log(N/t)$ is an antiderivative of K_N.

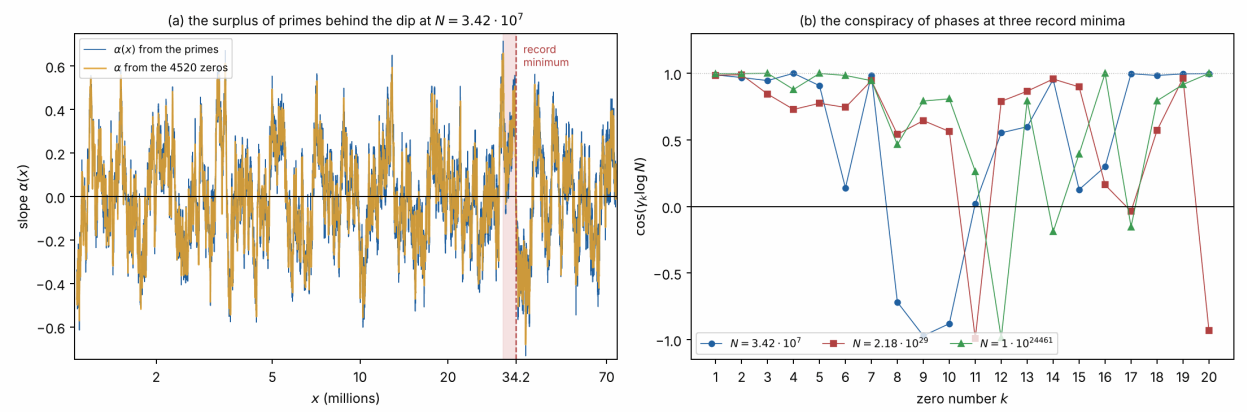

FIGURE 7. Anatomy of the record minima. (a) The slope α around the record at $N = 3.42 \cdot 10^7$, from the primes and from the zeros: a surplus of about 49 prime powers accumulates on $(8.5 \cdot 10^6, 1.7 \cdot 10^7]$, the margin pays for it on the shaded stretch, and the record sits where α returns to 0. (b) The phases $\cos(\gamma_k \log N)$ of the first 20 zeros at three records: the dips are conspiracies of the low zeros.

Proposition 7.1 (the margin as a weighted mean)**.** *For every $N > 1$, exactly and unconditionally,*

$$M(N) = -2 \int_1^N E^*(t)\, K_N(t)\, dt - 2\widehat{B} \log N. \tag{24}$$

Proof. All three steps are elementary. *First*, $C(N) = 2\int_1^N \psi(t) K_N(t)\, dt$: by the antiderivative above, $2\int_q^N K_N = 2q^{-1/2}\log(N/q)$, so the right side is $\sum_{q<N} 2\Lambda(q) q^{-1/2}\log(N/q) = C(N)$, with no boundary terms. *Second*, with $S(t) = t - \log 2\pi - \frac{1}{2}\log(1-t^{-2})$, termwise integration with $\int_1^N t^{a-1}(1 + \frac{1}{2}\log(N/t))\, dt = \frac{N^a - 1}{a} + \frac{N^a - 1 - a\log N}{2a^2}$, $a = -(2k+\frac{1}{2})$ for the k-th term of $-\frac{1}{2}\log(1 - t^{-2}) = \frac{1}{2}\sum_k t^{-2k}/k$, gives

$$2\int_1^N S(t) K_N(t)\, dt = 8\sqrt{N} + \Big(2 - 5\log 2 - 2\log\pi - \frac{\pi}{2}\Big)\log N + K_0 - \tau(\log N)\,:$$

the $N^{-2k-\frac{1}{2}}$ coefficients are $\frac{1}{k}\big(\frac{1}{a} + \frac{1}{2a^2}\big) = -\frac{2}{(2k+\frac{1}{2})^2}$, which is $-\tau$ term by term; the $\log N$ coefficient uses $\sum_{k\geq 1} \frac{1}{k(4k+1)} = \psi(\frac{5}{4}) + \gamma = 4 - 3\log 2 - \frac{\pi}{2}$; and the constant uses $\sum_{k\geq 1} \frac{1}{k(4k+1)^2} = 8 - 3\log 2 - \frac{\pi}{2} - \frac{\pi^2}{4} - 2G$, which follows from the trigamma value $\psi_1(\frac{1}{4}) = \pi^2 + 8G$, and comes out as exactly $K_0 = -16 + \frac{\pi^2}{2} + 4G$. *Third*, comparing with (1),

$$R(N) - 2\int_1^N S(t) K_N(t)\, dt = \big(-c_0 - 2 + 5\log 2 + 2\log\pi + \tfrac{\pi}{2}\big)\log N = (\log 4\pi - \gamma - 2)\log N = -2\widehat{B}\log N.$$

Since $\psi = S + E^*$, $M = R - C = -2\int E^* K_N - 2\widehat{B}\log N$. □

Identity (24) was checked to machine precision against (1) for $1.5 \leq N \leq 10^{15}$ (`ident.py`). The proof never mentions the zeros; the Hadamard constant $\widehat{B} = \sum_\rho \rho^{-1}$ materializes on the prime side out of the digamma values at $\frac{1}{4}$.

Corollary 7.2 (RH as one inequality for one mean)**.** *Let $\mu_K(N) = \frac{1}{\log N}\int_1^N E^*(t) K_N(t)\, dt$, the K_N-weighted mean of E^*, so that $M(N) = -2\log N\,(\mu_K(N) + \widehat{B})$. Then RH holds if and only if*

$$\mu_K(N) \leq -\widehat{B} \qquad \textit{for all sufficiently large } N,$$

and in that case $-\widehat{B} - \frac{0.18484}{2\log N} \le \mu_K(N) \le -\widehat{B}$ *for* $N > 1$, *so* $\mu_K(N) \to -\widehat{B}$ one-sidedly. *Unconditionally,* $\mu_K(N) \le -\widehat{B} + E(N)/(2\log N)$ *by Theorem 2.6, and the verification gives* $\mu_K(N) \le -\widehat{B} - 0.0171/\log N$ *for* $10^3 \le N \le 2.15 \cdot 10^{29}$ *(Figure 8).*

Three remarks on what (24) says, and what it cannot do.

Where the weight sits. The measure $K_N(t)\,dt$ has total mass $\log N$, of which the interval $[1,2]$ – before the first prime – carries the share $1 - 2^{-1/2} + O(1/\log N)$, about 30%, and $[1,4]$ about 51%; the oscillation at height t enters with density $\sim t^{-1/2}$ per unit. The mean is dominated by the first integers.

The head start and the payback. On $[1,2]$, $\psi = 0$ and the integrand is forced: that interval alone contributes $d\log N + 0.115$ to $M(N)$, with $d = \int_1^2 S(t)t^{-3/2}dt = 0.06344$. Since the total drift in (24) is $-2\widehat{B}\log N = -0.04619\log N$, the primes from 2 on must pay back at the exact asymptotic rate $d - 2\widehat{B} = 0.01725$ per unit of $\log N$ – they must run, in this weighted sense, slightly *ahead* of the smooth model forever – and RH is the statement that they never overpay: the accumulated account $\mu_K(N) + \widehat{B}$ reaches 0 from below at most.

Why explicit prime bounds are powerless here. The sharpest explicit estimates of the ψ-remainder, such as $|\psi(t) - t| \le 0.94\sqrt{t}$ for $11 < t \le 10^{19}$ [Bu], give only $|\mu_K(N)| \le 0.94\,(1 + \frac{1}{4}\log N) + O(1)$, which does not even keep μ_K bounded; the criterion needs the *signed* mean accurate to the absolute constant $\widehat{B} = 0.0231$. At $N = 10^{12}$ the gap is a factor of about 300. This is the quantitative form of the obstruction of Section 19: positivity of M is a statement about cancellation in E^* at the scale $\sqrt{t}$, which no absolute-value estimate can see, and Theorem 19.1 forbids any route that uses only $|E^*|$-information. We record (24) because it compresses RH into a single one-sided inequality for a single weighted mean, with an explicit classical constant on the other side.

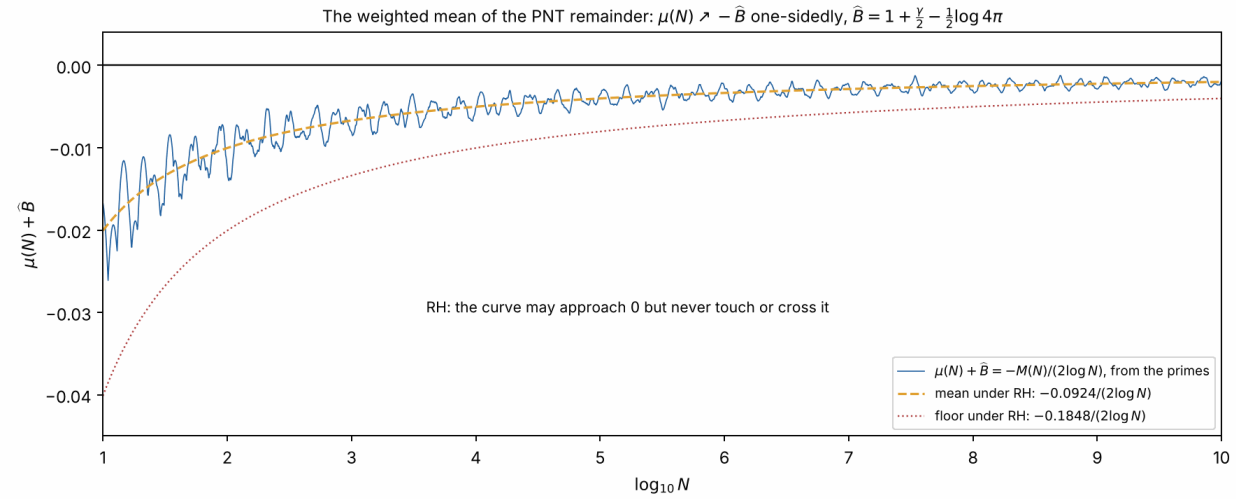

FIGURE 8. The one-sided convergence that is RH: $\mu_K(N) + \widehat{B} = -M(N)/(2\log N)$, computed from the primes up to 10^{10}, against the mean $-0.0924/(2\log N)$ and the extreme $-0.1848/(2\log N)$ allowed under RH. RH says the blue curve approaches the axis from below and never touches it.

Part II
How random is the Möbius function?

8. The Möbius side: the exact anchor and its square-root refinement

Two facts frame this paper. The first is an identity, valid for every real $x \geq 1$:

$$\sum_{d \leq x} \mu(d) \left\lfloor \frac{x}{d} \right\rfloor = \sum_{n \leq x} \sum_{d \mid n} \mu(d) = 1. \tag{25}$$

The summands have size x/d; the sum is always exactly 1. This is the Möbius orthogonality $\sum_{d \mid n} \mu(d) = [n = 1]$ in its rawest form, and it is the engine of Legendre's sieve. The second fact is that the Riemann hypothesis is the square-root refinement of (25). Writing $\{y\} = y - \lfloor y \rfloor$,

$$\Phi(x) = \sum_{d \leq x} \mu(d) \left\{ \frac{x}{d} \right\} = x \sum_{d \leq x} \frac{\mu(d)}{d} - 1, \tag{26}$$

and RH is equivalent to $\sum_{d \leq x} \mu(d)/d \ll x^{-1/2+\varepsilon}$, that is, to

$$\Phi(x) \ll x^{1/2+\varepsilon} \tag{27}$$

(see [T, §14.25–14.27]; the equivalent statements for $M(x) = \sum_{d \leq x} \mu(d)$ and for $\sum \mu(d)/d$ are classical, and Franel and Landau [FL] recast them as the equidistribution of the Farey fractions). The prime number theorem gives only $\Phi(x) \ll x \exp(-c(\log x)^{3/5-\varepsilon})$.

So the question "how far is RH?" has, in this language, a concrete form: *how is the cancellation in* (26) *organized, which parts of it can be proved, and exactly which part cannot?* This paper answers the first two questions with theorems and measurements, and the third with a precise statement of the obstruction. The contributions, in order of appearance:

(1) the z-smooth part of Φ is bounded, by an elementary complete-sieve argument, and the measurements locate the mass of the problem at the divisors $d \asymp x$ (Section 9);
(2) the long-run variance of the truncated sum obeys an exact, positive-definite identity, diagonalized by Jordan's totient (29); it identifies every constant observed in the data and shows that the gcd structure of μ cancels 46% of the diagonal variance (Section 10); a decomposition of $\widetilde{\Phi}_D$ into orthogonal *primitive sawtooths*, one per squarefree conductor, re-proves the identity and becomes the main tool of the paper (Proposition 10.5);
(3) the distribution of $\widetilde{\Phi}_D$ is empirically Gaussian, out to 4.5 standard deviations (Section 11);
(4) the Gaussian defect of the fourth moment has an exact anatomy (Section 12): it splits into an intra-orbit part, computed exactly — with the closed law $\varepsilon_4(p) = -(1 - p^{-4})/120$ at prime conductors — and a cross-orbit term $X(D)$. We prove that the conductor-1 orbit decouples completely (the classes carrying an integer frequency cancel identically in X) and that $X(p) = X(p-1)$ at every prime p, so the cross term moves only at composite squarefree conductors; the proofs are real-space identities needing no regularization, confirmed in exact rational arithmetic. Splitting X into four sectors by conductor multiplicities, we prove the $(3,1)$ and $(2,2)$ sectors are $O(D \log^4 D)$ unconditionally, prove the exact jump law of the $(2,2)$ sector at the conductors $2p$ (Theorem 12.18), prove a parity sieve that annihilates every sector term using an odd number of frequency slots of even conductor (Lemma 12.10), and prove $X_{211} \ll D^{3/2+o(1)}$ — the first unconditional bound on a sector with three distinct conductors (Proposition 12.15). All four sectors are then computed to $D = 800$ (Section 12.6): the conjectured object stays below 0.44 in absolute value, and the growth of its absolute-value majorant indicates that, if the trend persists, the

quadruple sector cannot be bounded by absolute values alone, locating exactly what kind of cancellation Conjecture 12.19 still needs;

(5) a precise statement of why none of this reaches the Riemann hypothesis (Section 21, placed with the other walls in Part V).

The data reach $x = 2 \cdot 10^7$ for the full sum, $D = 3 \cdot 10^4$ for distributions, 10^8 samples for the fourth moment, and exact computations over the full period $\operatorname{lcm}(1, \dots, D)$ for $D \le 18$. Everything is reproducible from the ancillary code. Part I carries the complementary, one-sided form of RH that motivates this part: there, RH is the statement that a weighted mean of the prime-counting error stays below the explicit constant $-\big(1 + \frac{\gamma}{2} - \frac{1}{2}\log 4\pi\big)$ for ever; here we study the Möbius side of the same coin. We do not prove the Riemann hypothesis, and no statement below should be read as progress towards it beyond what it literally says.

Throughout, $B_1(y) = \{y\} - \frac{1}{2}$ is the sawtooth, $\mu^2(k) = 1$ marks squarefree k, $J_2(k) = k^2 \prod_{p \mid k}(1 - p^{-2})$ is Jordan's totient, and

$$\widetilde{\Phi}_D(x) = \sum_{d \le D} \mu(d) B_1\Big(\frac{x}{d}\Big), \qquad m_k(Y) = \sum_{\substack{m \le Y \\ (m,k)=1}} \frac{\mu(m)}{m}$$

are the truncated sawtooth sum and the restricted Mertens sums; $\Phi_D(x) = \sum_{d \le D} \mu(d)\{x/d\}$ is the plain truncation of (26). "Long-run" averages, written $\langle \cdot \rangle = \lim_{T \to \infty} \frac{1}{T} \int_0^T (\cdot)\, dx$, exist because $\widetilde{\Phi}_D$ is periodic with period $\operatorname{lcm}(1, \dots, D)$. One convention is used silently throughout: every conditionally convergent sum over frequencies is a *symmetric* (principal-value) limit, $\lim_{U \to \infty} \sum_{0 < |n| \le U} \cdot$ This is the convergence in which $B_1(y) = \sum_{n \ne 0} e(ny)/(-2\pi i n)$ holds at every non-integer y, and all the moment computations below inherit it.

9. The smooth part is tame, and the mass sits at the large divisors

Lemma 9.1 (the z-smooth part is bounded)**.** *Fix $z \ge 2$ and let $P(z) = \prod_{p \le z} p$. The contribution to $\Phi(x)$ of the d free of prime factors $> z$ is*

$$S_z(x) = \sum_{d \mid P(z)} \mu(d)\Big\{\frac{x}{d}\Big\} = x \prod_{p \le z}\Big(1 - \frac{1}{p}\Big) - \#\{n \le x : (n, P(z)) = 1\} \qquad (x \ge P(z)),$$

a periodic function of x with period $P(z)$, mean 0, and $|S_z(x)| \le 2^{\pi(z)-1}$.

Proof. A squarefree z-smooth d divides $P(z)$, and for $x \ge P(z)$ all divisors of $P(z)$ are $\le x$, so the sum is finite and complete; the displayed identity is (25) relativized to $P(z)$ (Legendre's formula). Periodicity is that of each $\{x/d\}$; the mean of $\{x/d\}$ over a period is $\frac{1}{2}$ and $\sum_{d \mid P(z)} \mu(d) = 0$ for $z \ge 2$; the bound is $\frac{1}{2} \cdot \#\{d \mid P(z)\}$. □

Numerically $\max |S_z| = 1.37,\ 3.03,\ 5.46$ for $z = 7, 13, 19$ on $x \le 2 \cdot 10^7$ (Figure 9(d)), far below the trivial bound. Whatever is hard in (27), it is not the smooth divisors.

Lemma 9.2 (the pairing $d \leftrightarrow pd$)**.** *For any function g and any prime p,*

$$\sum_{d \le x} \mu(d) g(d) = \sum_{d \le x,\ p \nmid d} \mu(d)\big[g(d) - g(pd)\,\mathbf{1}_{pd \le x}\big].$$

Proof. A squarefree d is either coprime to p or equals pd' with $p \nmid d'$, and $\mu(pd') = -\mu(d')$. □

Iterating over primes turns the Möbius sum into difference operators $(\mathrm{I} - T_p)$ applied to the kernel: the Euler product in operator form. Its statistical signature is visible in the data. Splitting $\Phi(x)$ into octaves $d \in [2^j, 2^{j+1})$, over 120 samples of x near $2 \cdot 10^7$: *(i)* every octave shows square-root cancellation with the same constant, $\mathrm{sd}_j \approx 0.22\sqrt{2^j}$ for all $j \le 22$ (the constant is identified

in Remark 10.3 below as $1/\sqrt{2\pi^2} = 0.2251$); *(ii)* the octaves are not independent: $\mathrm{Var}(\Phi) = 0.40\sum_j \mathrm{Var}(\text{octave } j)$, with neighbouring octaves anticorrelated by as much as -0.56 — the pairing of Lemma 9.2 seen as statistics; *(iii)* truncations $\Phi_D(x)$ stay negligible until $D \approx x/30$: the mass sits at $d \asymp x$, the type II range, where $\{x/d\} = x/d - k$ on $(\frac{x}{k+1}, \frac{x}{k}]$ is piecewise smooth and each block reduces exactly to short Mertens sums,

$$\sum_{\frac{x}{k+1}<d\le\frac{x}{k}} \mu(d)\Big\{\frac{x}{d}\Big\} = x\sum_{\frac{x}{k+1}<d\le\frac{x}{k}} \frac{\mu(d)}{d} - k\Big[M\Big(\frac{x}{k}\Big) - M\Big(\frac{x}{k+1}\Big)\Big].$$

Figure 9 shows all four facts. The reduction localizes the problem — RH is square-root cancellation in smooth-weighted short Mertens sums at the top octaves — without solving any of it: unconditionally those blocks are only known to be $O(x\exp(-c(\log x)^{3/5-\varepsilon}))$.

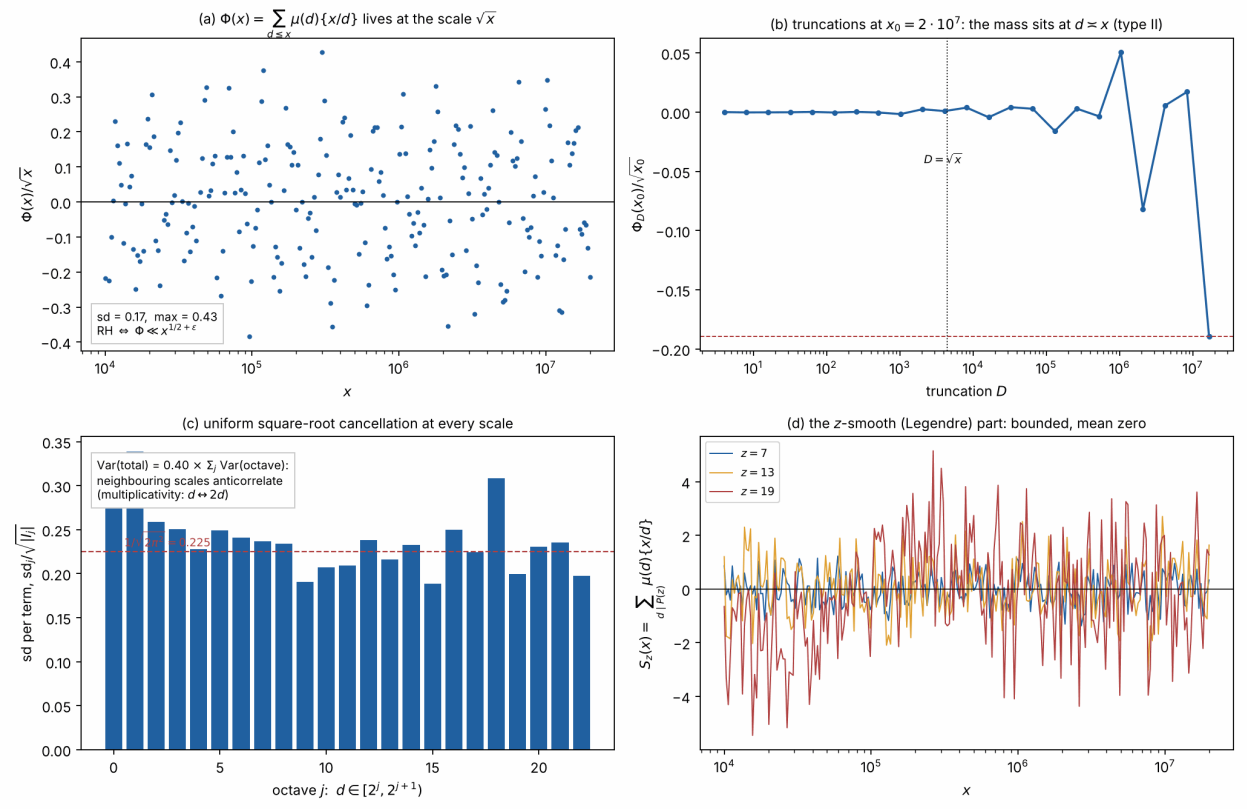

FIGURE 9. The scales of $\Phi(x)$ to $2\cdot 10^7$. (a) $\Phi/\sqrt{x}$ is bounded in the range. (b) Truncations: the mass sits at $d \asymp x$. (c) Square-root cancellation with the same constant at every octave. (d) The z-smooth part, bounded (Lemma 9.1).

10. The variance: an exact, positive-definite identity

Lemma 10.1 (the Franel–Montgomery–Vaughan kernel; cf. [FL, MV])**.** *For $d, e \ge 1$,*

$$\lim_{T\to\infty}\frac{1}{T}\int_0^T B_1\Big(\frac{x}{d}\Big)B_1\Big(\frac{x}{e}\Big)\,dx = \frac{(d,e)^2}{12\,de}.$$

Proof. Both factors are periodic with common period de/g, $g = (d,e)$, so the long-run average is the integral over one period, and Parseval applies with absolute convergence: writing $B_1(y) = \sum_{n\neq 0} e(ny)/(-2\pi i n)$ (valid off the integers, a null set), the average keeps only the resonant pairs

$n/d + m/e = 0$, i.e. $n = kd/g$, $m = -ke/g$ with $k \neq 0$; such a pair contributes $-1/(4\pi^2 nm) = g^2/(4\pi^2 k^2 de)$, and $\sum_{k\neq 0} k^{-2} = \pi^2/3$. □

Theorem 10.2 (the variance identity)**.** *For every $D \geq 1$,*

$$\lim_{T\to\infty} \frac{1}{T}\int_0^T \widetilde{\Phi}_D(x)^2\,dx = \frac{V(D)}{12}, \qquad V(D) = \sum_{d,e\leq D} \frac{\mu(d)\mu(e)\,(d,e)^2}{de}, \tag{28}$$

and V diagonalizes as a positive combination of squared restricted Mertens sums:

$$V(D) = \sum_{k\leq D} \mu^2(k)\,\frac{J_2(k)}{k^2}\; m_k\Big(\frac{D}{k}\Big)^2 \;\geq\; 0. \tag{29}$$

Proof. (28) is Lemma 10.1 summed over $d, e \leq D$ with weights $\mu(d)\mu(e)$. For (29), insert Jordan's identity $(d,e)^2 = \sum_{k|(d,e)} J_2(k)$ and exchange sums:

$$V(D) = \sum_{k\leq D} J_2(k)\Big(\sum_{\substack{d\leq D\\ k|d}} \frac{\mu(d)}{d}\Big)^2,$$

and non-squarefree k contribute nothing, since no squarefree d is divisible by such a k and the inner sum vanishes. For squarefree k, writing $d = km$ with $(m,k) = 1$ (any other d has $\mu(d) = 0$) gives $\sum_{k|d\leq D} \mu(d)/d = (\mu(k)/k)\, m_k(D/k)$ by multiplicativity. Squaring removes the sign $\mu(k)^2 = 1$ and leaves (29). □

Remark 10.3 (the constants of the data).

(1) The diagonal $d = e$ alone gives variance $(6/\pi^2)D/12 = D/(2\pi^2)$ as $D \to \infty$: per squarefree term, $1/12$; per integer, $1/(2\pi^2)$, i.e. standard deviation 0.22508 per term — the constant of Figure 9(c).

(2) The measured $c_V(D) = V(D)/D$ settles quickly: $c_V(3000) = 0.33043$, $c_V(10^4) = 0.33138$, $c_V(10^5) = 0.33111$, $c_V(2\cdot 10^6) = 0.33114$ — three stable digits from $D = 3000$ on. (Whether $\lim V(D)/D$ exists is not proved here; the identity (29) is exact at every D.) Taking $c_V = 0.33114$,

$$\operatorname{sd}\widetilde{\Phi}_D = \sqrt{\frac{c_V D}{12}} = 0.16612\,\sqrt{D},$$

against $0.2251\sqrt{D}$ for independent terms: *the gcd structure of μ cancels* 46% *of the diagonal variance*, unconditionally. The moving-truncation measurement ($D = x$, Figure 9(a)) gave $0.17\sqrt{x}$.

(3) In (29), the top octave $D/2 < k \leq D$, where $m_k(D/k) = 1$ trivially, carries 81% of $V(2\cdot 10^6)$; the top two octaves carry 96%. The small k, which carry the restricted Mertens sums at large argument, are negligible *unconditionally*, since $m_k(Y) \to 0$ by the prime number theorem. The second moment at fixed D is, in this sense, completely understood without any hypothesis.

The gcd kernel is the aggregate view. The sharper tool, used for the rest of the paper, is to group the Fourier frequencies of $\widetilde{\Phi}_D$ by their reduced denominator.

Lemma 10.4 (class weights)**.** *Group the lifts of a reduced fraction u/q, $(u,q) = 1$: for every $D \geq q$,*

$$\sum_{d\leq D,\ n/d=u/q} \frac{\mu(d)}{n} = \frac{\mu(q)}{u}\, m_q\Big(\frac{D}{q}\Big).$$

Proof. $d = qm$ with $(m,q) = 1$, $n = um$, and $\mu(qm)/(um)$ summed over $m \leq D/q$ is the statement, by multiplicativity. □

Proposition 10.5 (primitive decomposition)**.** *Group the frequencies of $\widetilde{\Phi}_D$ by their reduced denominator. Then, exactly,*

$$\widetilde{\Phi}_D(x) = \sum_{\substack{q \le D \\ q \text{ squarefree}}} \mu(q)\, m_q\Big(\frac{D}{q}\Big)\, \sigma_q\Big(\frac{x}{q}\Big), \qquad \sigma_q(y) = \sum_{t \mid q} \frac{\mu(t)}{t}\, B_1(ty),$$

where σ_q is the primitive sawtooth *of conductor q: its Fourier frequencies are exactly the u with $(u, q) = 1$. Distinct components have disjoint frequency sets, hence are mutually orthogonal, and*

$$\int_0^1 \sigma_q^2 = \frac{J_2(q)}{12\, q^2},$$

which gives a second proof of (29): the variance identity is the orthogonality of the primitive components.

Proof. By Lemma 10.4 the aggregated coefficient of the reduced frequency u/q in $\widetilde{\Phi}_D$ is $\mu(q)m_q(D/q)/(-2\pi i u)$; and σ_q as defined has coefficient $1/(-2\pi i u)$ at u coprime to q and 0 otherwise, by Möbius inclusion–exclusion over $t \mid q$. The L^2 norm is $\sum_{(u,q)=1}(2\pi u)^{-2} = \frac{1}{2\pi^2}\zeta(2)\prod_{p\mid q}(1-p^{-2})$. □

The identity was verified against sampling: at $D = 3000$, sampled variance 82.35 against $V(D)/12 = 82.61$; the theoretical line fits from $D = 300$ to $3 \cdot 10^4$ (Figure 10). Positivity of (29) has a one-sided consequence worth noting: any single k gives a free *lower* bound for the variance; upper bounds for individual m_k at large arguments are, of course, exactly the hard direction.

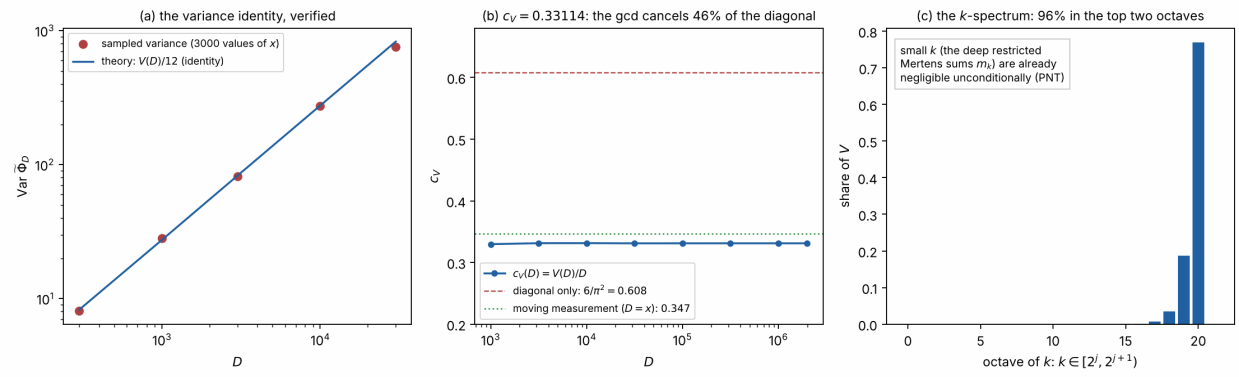

FIGURE 10. The variance identity. (a) Verified against sampling. (b) $c_V = 0.33114$, between the diagonal $6/\pi^2$ and the moving measurement. (c) The spectrum of (29) in k.

11. The distribution is empirically Gaussian

With the second moment exact, the next question is the shape; this section is measurement, not theorem. (For the distribution of the Mertens function $M(x)$ itself, see Ng [Ng].) Over 10^6 samples at $D = 300$, $3 \cdot 10^5$ at $D = 3000$ and $6 \cdot 10^4$ at $D = 3 \cdot 10^4$ (windows of x near 10^8), the normalized $\widetilde{\Phi}_D/\mathrm{sd}$ is Gaussian to remarkable accuracy: skewness -0.003, $+0.003$, $+0.006$; kurtosis 2.992, 3.009, 2.986, each within 1.5 standard errors of 3; and the empirical tails follow the normal law out to $\lambda = 4.5$, where the probability is $7 \cdot 10^{-6}$ (Figure 12(a,b)). The symmetry is worth a remark: the one-sidedness that the Riemann hypothesis needs in its margin form (Part I) is nowhere in the distribution of Φ; it lives in the Fejér positivity of the explicit formula, not here.

12. The fourth moment: anatomy of the Gaussian defect

12.1. **The resonance expansion.** The method of Lemma 10.1 computes every moment. For the fourth, define the *Gaussian defect*

$$e_4(D) = \lim_{T\to\infty} \frac{1}{T}\int_0^T \widetilde{\Phi}_D(x)^4\,dx - 3\Big(\frac{V(D)}{12}\Big)^2. \tag{30}$$

The subtracted term collects the Wick pairings of resonances — the count a Gaussian would give — so $e_4(D)$ is what the genuine four-term resonances leave behind: tuples (n_i, d_i), $n_i \neq 0$, $d_i \le D$, with $\sum_{i\le 4} n_i/d_i = 0$ and no vanishing two-term subsum, weighted by $\prod \mu(d_i)/(-2\pi i n_i)$, plus the lower-order corrections from coincident pairings. *The excess kurtosis measures* e_4. By Lemma 10.4, every four-term resonance class of reduced fractions (u_i/q_i) with $\sum u_i/q_i = 0$ enters e_4 with weight $\frac{1}{(2\pi i)^4}\prod_i \frac{\mu(q_i)}{u_i} m_{q_i}(D/q_i)$.

Regularization. Two conventions make this expansion precise. Since $\widetilde{\Phi}_D$ is bounded and the symmetric partial sums of the Fourier series of B_1 are uniformly bounded and converge off the integers, bounded convergence computes every joint moment as the limit of truncated frequency sums, the truncations tending to infinity independently; at this stage the value is order-independent. The expansion then groups the surviving quadruples into the classes of Lemma 10.4; within a class the numerator sum is only conditionally convergent, and it is always evaluated in the symmetric (principal-value) sense fixed in Section 8. The main theorems of this section turn out to need none of this: they are proved in real space, as identities between long-run averages of bounded periodic functions over a common period. The symmetric convention governs the class-by-class identities of Remark 12.7 and the heuristic discussion; we do not claim a general interchange theorem for the infinite sum over classes, and every assertion that ties class sums to the measurable quantities e_4, I, X is backed by the exact full-period computations, which involve no limits at all.

Two structural consequences. *(i) Classes containing an integer frequency cancel identically in* X: genuine such classes exist, with nonzero individual weights, at every fixed $D \ge 6$ — the conductor multiset $\{1,2,3,6\}$, with witness $\frac12+\frac13+\frac16-1=0$, is live from $D=6$ — but the conductor-1 orbit decouples completely from the fourth moment (Theorem 12.4 below); the factor $m_1(D)\to 0$ makes them small, the decoupling makes them irrelevant. *(ii) The live classes need all* $q_i \asymp D$, where $m_{q_i}(D/q_i) = 1 + O(q_i/D)$ gives no damping; clearing denominators then forces the multiplicative matching of Proposition 12.2 below. The sign $\mu(q_1)\cdots\mu(q_4)$ is $+1$ whenever every prime divides an even number of the conductors, though not always (a prime may divide exactly three: the multiset $\{2,6,6,6\}$, with witness $-\frac12+\frac16+\frac16+\frac16=0$, carries sign -1); the real damping is that the numerators live on a rank-3 lattice of covolume $\asymp D$, so the weights $1/(u_1u_2u_3u_4)$ suppress every class. This is the structural reason to expect the non-pairing resonances to total far less than the trivial $O(D^2)$ scale of the expansion (two of the four conductors free, each $\asymp D$).

12.2. **The intra-orbit part is exact.** The decomposition of Proposition 10.5 splits the Gaussian defect into the *intra-orbit* defects, which are exactly computable, and the *cross-orbit* resonances:

$$e_4(D) = \underbrace{\sum_{q\le D} \mu^2(q)\, m_q\Big(\frac{D}{q}\Big)^4 \varepsilon_4(q)}_{I(D)} + X(D), \qquad \varepsilon_4(q) = \Big\langle \sigma_q^4 \Big\rangle - 3\Big\langle \sigma_q^2 \Big\rangle^2, \tag{31}$$

where $X(D)$ is *defined* by this identity — so it is as rigorous as e_4 and I — and, in the resonance expansion, collects the quadruples of frequencies from at least two different conductors. Each $\varepsilon_4(q)$ is a finite rational computation: σ_q is piecewise linear with nodes on $\frac1q\mathbb{Z}$ and slope $\sum_{t\mid q}\mu(t) = [q=1]$, so its moments are finite sums. The first values are

$$\varepsilon_4(1) = -\tfrac{1}{120},\quad \varepsilon_4(2) = -\tfrac{1}{128},\quad \varepsilon_4(3) = -\tfrac{2}{243},\quad \varepsilon_4(5) = -\tfrac{26}{3125},\quad \varepsilon_4(6) = -\tfrac{1}{216},\quad \dots$$

all negative. For prime conductor there is a closed law with a pleasant proof.

Lemma 12.1 (prime conductors)**.** *For $q = p$ prime, σ_p is the* discrete *uniform variable on the p equally spaced values $\frac{2j+1-p}{2p}$, $0 \le j < p$, and*

$$\varepsilon_4(p) = -\frac{1}{120}\big(1 - p^{-4}\big).$$

Proof. Write $y = (j+u)/p$ with j uniform on $\{0, \dots, p-1\}$ and u uniform on $[0,1)$, independent. Then $\sigma_p(y) = B_1(y) - \frac{1}{p}B_1(py) = \frac{j+u}{p} - \frac{1}{2} - \frac{u-\frac{1}{2}}{p} = \frac{2j+1-p}{2p}$: the continuous part cancels. The moments of the symmetric discrete uniform give $\langle\sigma_p^2\rangle = (1-p^{-2})/12$ and, with $\sum_j (2j+1-p)^4 = p(p^2-1)(3p^2-7)/15$, the stated $\varepsilon_4(p)$. □

The law was also confirmed to 10^{-17} for all primes up to 1000 by the finite algorithm.

12.3. **The measurement, split exactly.** With $I(D)$ computed exactly and the fourth moment known exactly for $D \le 18$ (full-period computation) and to high precision at $D = 100$ (10^8 samples) and $D = 1000$:

D	6	10	12	15	18	100	1000
$e_4(D)$, eq. (30)	−0.0181	−0.0253	−0.0336	−0.0327	−0.0411	−0.2405(38)	−1.60(74)
$I(D)$ (exact)	−0.0150	−0.0210	−0.0293	−0.0429	−0.0512	−0.2751	−2.679
$X(D)$	−0.0031	−0.0043	−0.0043	+0.0102	+0.0102	+0.035	+1.1(7)

Three facts stand out. *(i)* The intra part dominates and is exact; numerically $I(D)/D \to -0.0027$, and the measured law $e_4 \approx -0.00228\,D$ of Figure 12(c) splits as $-3.0 = -3.5_{\text{intra}} + 0.5_{\text{cross}}$ in the normalized units of the excess-kurtosis law $-3/D$. *(ii)* X is piecewise constant in D and *moves only at composite squarefree conductors.* This is now a theorem in full: nothing can move at non-squarefree D (a new conductor, or a new term m with $\mu(m) \neq 0$ in a weight $m_q(D/q)$, requires $D = qm$ with q, m squarefree and coprime, so D squarefree), and $X(p) = X(p-1)$ at every prime p (Theorems 12.4 and 12.6 below). *(iii)* X changes sign, exactly: at $D = 15$, where $X(15) = \frac{822653}{81033750} = +0.01015\ldots$; and it stays minute against the trivial $O(D^2)$ scale of the resonance expansion — $X(1000) = +1.1(7)$.

12.4. **The prime-conductor cancellation.** The constancy of X across prime arrivals has a complete and elementary proof, and — unexpectedly — it needs no regularization at all: everything below is an identity between long-run averages of bounded periodic functions over a common period. Three ingredients: the valuation lemma, which is the p-adic skeleton; a one-line *unit-interval lemma*, which decouples the conductor-1 orbit completely; and the disjointness of Fourier supports. Call a zero-sum quadruple of reduced nonzero frequencies *genuine* if no two-term subsum vanishes.

Proposition 12.2 (valuation lemma)**.** *In any zero-sum quadruple of reduced frequencies u_i/q_i, every prime divides either none or at least two of the conductors q_i.*

Proof. If $p \mid q_1$ and $p \nmid q_2 q_3 q_4$, the p-adic valuation of u_1/q_1 is negative and that of the other three terms is ≥ 0; the sum cannot be 0. □

This is the multiplicative matching made precise: cross-orbit resonances require conductors that share their primes in pairs, which is why only composite conductors can couple to older ones in a new way.

Lemma 12.3 (unit-interval lemma)**.** *For squarefree $q > 1$ the function $x \mapsto \sigma_q(x/q)$ is constant on every interval $(j, j+1)$, $j \in \mathbb{Z}$. Consequently, for any function g of integer period that is constant on unit intervals — in particular any polynomial in such σ's, with any conductors > 1 —*

$$\widehat{g}(k) = 0 \quad \textit{for every integer } k \neq 0, \qquad \big\langle B_1(x)^m\, g(x)\big\rangle = \big\langle B_1^m\big\rangle\,\langle g\rangle \quad (m \ge 0).$$

Proof. $\sigma_q(x/q) = \sum_{t|q}(\mu(t)/t)B_1(tx/q)$ jumps only where $tx/q \in \mathbb{Z}$, i.e. at multiples of the integer q/t, and its slope is $\frac{1}{q}\sum_{t|q}\mu(t) = 0$ for $q > 1$. If g has period L and value g_j on $(j, j+1)$, then $\widehat{g}(k) = \frac{1}{L}\sum_{j<L} g_j \int_j^{j+1} e(-kx)\,dx = 0$ for $0 \neq k \in \mathbb{Z}$, and $\langle B_1^m g\rangle = \frac{1}{L}\sum_{j<L} g_j \int_0^1 B_1(t)^m\,dt = \langle g\rangle\langle B_1^m\rangle$. □

Theorem 12.4 (the conductor-1 orbit decouples)**.** *Write* $\widetilde{\Phi}_D = m_1(D)\,B_1 + \Psi_D$ *with* $\Psi_D = \sum_{1<q\le D}\mu(q)\,m_q(D/q)\,\sigma_q(x/q)$*. Then*

$$\langle\widetilde{\Phi}_D^4\rangle = m_1^4\,\langle B_1^4\rangle + 6\,m_1^2\,\langle B_1^2\rangle\langle\Psi_D^2\rangle + \langle\Psi_D^4\rangle, \qquad \langle\widetilde{\Phi}_D^2\rangle = \frac{m_1^2}{12} + \langle\Psi_D^2\rangle,$$

hence

$$e_4(D) = m_1(D)^4\,\varepsilon_4(1) + \Big(\langle\Psi_D^4\rangle - 3\langle\Psi_D^2\rangle^2\Big),$$

and $X(D)$ *does not depend on* $m_1(D)$ *at all:* the resonance classes containing an integer frequency contribute nothing to X in total.

Proof. Ψ_D and its powers are constant on unit intervals (Lemma 12.3), so every mixed term splits: $\langle B_1^a\Psi_D^b\rangle = \langle B_1^a\rangle\langle\Psi_D^b\rangle$, which vanishes for $a = 1, 3$ and equals $\frac{1}{12}\langle\Psi_D^b\rangle$ for $a = 2$. In $e_4 = \langle\widetilde{\Phi}_D^4\rangle - 3\langle\widetilde{\Phi}_D^2\rangle^2$ the terms $6m_1^2\cdot\frac{1}{12}\langle\Psi_D^2\rangle$ cancel exactly, leaving $m_1^4(\frac{1}{80} - \frac{1}{48}) = m_1^4\varepsilon_4(1)$ plus the pure Ψ_D defect. Since the $q = 1$ term of $I(D)$ is exactly $m_1^4\varepsilon_4(1)$, the weight m_1 drops out of $X = e_4 - I$. □

Remark 12.5. The cancellation is also exact *per conductor multiset*: a single multiset involves finitely many fixed functions, so the bounded-convergence expansion applies to it alone, and Lemma 12.3 gives $\langle B_1^a\,\sigma_{q_1}\cdots\sigma_{q_b}\rangle = \langle B_1^a\rangle\langle\sigma_{q_1}\cdots\sigma_{q_b}\rangle$. For instance, the symmetric family sum of $\{1, 2, 3, 6\}$ (witness $\frac{1}{2} + \frac{1}{3} + \frac{1}{6} - 1 = 0$) vanishes, as direct summation confirms.

Theorem 12.6 (prime arrivals leave X unchanged)**.** $X(p) = X(p-1)$ *for every prime* p*.*

Proof. When D moves from $p-1$ to p, the only coefficients of the primitive decomposition that change are $m_1(D)$, which jumps by $\mu(p)/p$, and the new $c_p = \mu(p)m_p(1) = \mu(p)$: a new term in $m_q(D/q)$ requires $qm = D$ with $\mu(m) \neq 0$, which at $D = p$ forces $q \in \{1, p\}$. The change in m_1 is invisible to X by Theorem 12.4. Write $\Psi = \Psi_{p-1}$ for the unchanged part; every conductor in Ψ is $\le p-1$, hence coprime to p. The Fourier support of $\sigma_p(x/p)$ is $\{u/p : p \nmid u\}$; the supports of Ψ, Ψ^2, Ψ^3 are contained in the fractions with denominator coprime to p (sumsets of the Ψ-frequencies); the support of σ_p^3 lies in $\frac{1}{p}\mathbb{Z}$. All functions here are bounded and piecewise polynomial with a common period, so their Fourier coefficients are square-summable and Parseval applies. Then $\langle\sigma_p\Psi\rangle = \langle\sigma_p\Psi^3\rangle = 0$ (a reduced denominator equal to p never equals one coprime to p: Proposition 12.2 in L^2 form) and $\langle\sigma_p^3\Psi\rangle = 0$ ($\frac{1}{p}\mathbb{Z}$ meets the support of Ψ only in the integers, which Ψ does not carry). Therefore adding $c_p\sigma_p$ changes e_4 by

$$4c_p\langle\sigma_p\Psi^3\rangle + 6c_p^2\big(\langle\sigma_p^2\Psi^2\rangle - \langle\sigma_p^2\rangle\langle\Psi^2\rangle\big) + 4c_p^3\langle\sigma_p^3\Psi\rangle + c_p^4\varepsilon_4(p),$$

and I by $c_p^4\varepsilon_4(p)$, so ΔX is the three mixed brackets. The first and third vanish as above. For the middle one, the supports of σ_p^2 ($\subset \frac{1}{p}\mathbb{Z}$) and Ψ^2 (denominators coprime to p) meet only in the integers — and both functions are constant on unit intervals, so by Lemma 12.3 neither carries any nonzero integer frequency: $\langle\sigma_p^2\Psi^2\rangle = \langle\sigma_p^2\rangle\langle\Psi^2\rangle$. □

Remark 12.7 (the frequency-space view). Class by class, the same cancellations read as principal-value identities, which we found first and record for their own sake. A pair of conductor-p frequencies summing to an integer $k \neq 0$ carries the factor

$$T_p(k) = \sum_{u\neq 0,\ p\nmid u} \frac{1}{u(pk-u)} = \frac{1}{pk}\sum_{u\neq 0,\ p\nmid u}\Big(\frac{1}{u} + \frac{1}{pk-u}\Big) = 0$$

(two negation-symmetric harmonic sums), and a triple of them summing to $N = pk$ carries $U_p(N) = \sum_{u_1+u_2+u_3=N,\ p\nmid u_i} 1/(u_1u_2u_3)$, which also vanishes: fixing u_3 and using partial fractions with the classical $\sum_{n\equiv a\,(p)} 1/n = \frac{\pi}{p}\cot\frac{\pi a}{p}$ gives $\frac{2\pi}{p(N-u_3)}\cot\frac{\pi u_3}{p}$, and the sum over u_3 pairs the resulting $\sum_a \cot^2\frac{\pi a}{p}$ against its reflection with the opposite sign, by $\cot\frac{\pi(N-w)}{p} = -\cot\frac{\pi w}{p}$. (The restriction to genuine quadruples is vacuous in both: a vanishing two-term subsum would force $k = 0$ or $p \mid u_i$.) These identities hold in the symmetric limit of Section 8; the real-space proof above is what makes that limit, and the bookkeeping it would require across infinitely many classes, unnecessary.

The exact computations confirm the theorems and sharpen the picture. $X(D) = 0$ for $D \le 5$ (no composite conductor yet); the first nonzero value of the cross term is the exact rational $X(6) = -\frac{1}{324}$; and across every prime arrival in range X indeed does not move:

$$X(7) = X(6) = -\tfrac{1}{324}, \qquad X(11) = X(12) = X(13) = -\tfrac{217}{50625} = -0.0042864198,$$
$$X(15) = X(18) = X(19) = +0.0101519799.$$

The composite arrivals move it: $\Delta X = -\frac{1}{324}$ at 6, $-\frac{3}{2500}$ at 10, $-\frac{3}{4802}$ at 14, $\frac{1830934}{121550625}$ at 15, $-\frac{211124}{24310125}$ at 21 and $-\frac{44110277}{22598383500}$ at 22. The first three are *exactly* $-(p^2-1)/(32p^4)$ for the conductors $2p$, $p = 3, 5, 7$ — and the out-of-sample test at $p = 11$ refuted that law, $\Delta X(22) = -0.00195\ldots$ against a predicted -0.00026. When we first found this we recorded it as a caution against single-family formulas; the sector decomposition of the next subsection turns both halves of the experience into theorems: the law is exactly true, at *every* prime p, for one of the four sectors of X, and the correction at $p = 11$ is the other three sectors switching on, each by an exact, explained amount (Theorem 12.18). Only composite conductors drive X, exactly as Proposition 12.2 prescribes.

12.5. **The four sectors of the cross term.** Grouping the four frequencies of a resonance by the multiplicities of their conductors splits X into four *sectors*, each an exact and finite object. Throughout this subsection σ_q abbreviates $\sigma_q(x/q)$, all conductors are squarefree in $(1, D]$, $c_q = \mu(q)\, m_q(D/q)$, and $\mathrm{Cov}(f, g) = \langle fg\rangle - \langle f\rangle\langle g\rangle$.

Proposition 12.8 (sector decomposition)**.** *For every D,*

$$X(D) = X_{31}(D) + X_{22}(D) + X_{211}(D) + X_{1111}(D),$$

$$X_{31} = 4\sum_{q\neq r} c_q^3 c_r\,\langle\sigma_q^3\sigma_r\rangle, \qquad X_{22} = 6\sum_{q<r} c_q^2 c_r^2\,\mathrm{Cov}(\sigma_q^2, \sigma_r^2),$$
$$X_{211} = 12\sum_{q}\sum_{\substack{r<s\\ r,s\neq q}} c_q^2 c_r c_s\,\langle\sigma_q^2\sigma_r\sigma_s\rangle, \qquad X_{1111} = 24\sum_{q<r<s<t} c_q c_r c_s c_t\,\langle\sigma_q\sigma_r\sigma_s\sigma_t\rangle,$$

and each average is a finite exact computation over the common period of its conductors.

Proof. Expand $\langle\Psi_D^4\rangle$ multinomially over the orbits and subtract $3\langle\Psi_D^2\rangle^2 = 3(\sum_q c_q^2 v_q)^2$, $v_q = \langle\sigma_q^2\rangle$; orthogonality kills nothing here but pairs the Wick term $6\sum_{q<r} c_q^2 c_r^2 v_q v_r$ with the two-pair moments, producing the covariances, while the all-equal terms give $\sum_q c_q^4(\langle\sigma_q^4\rangle - 3v_q^2) = I(D) - m_1^4\varepsilon_4(1)$, which cancels in $X = e_4 - I$ by Theorem 12.4. □

Lemma 12.9 (support and parity)**.** (i) $\langle\sigma_q^3\sigma_r\rangle = 0$ *unless $r \mid q$.* (ii) $\sigma_2 = \mp\frac{1}{4}$ *on unit intervals, so $\sigma_2^2 \equiv \frac{1}{16}$ and $\mathrm{Cov}(\sigma_2^2, \cdot) = 0$: conductor 2 is inert in the $(2,2)$ sector.* (iii) *For even squarefree q, σ_q^2 carries no half-integer frequency: its coefficient at ξ with 2ξ odd sums $1/(uu')$ over $u + u' = q\xi$ with u, u' odd units of q, and $q\xi = \frac{q}{2}(2\xi)$ is odd — an odd number is not a sum of two odd ones.* (iv) *Hence $\mathrm{Cov}(\sigma_q^2, \sigma_r^2) = 0$ unless q and r share an* odd *prime: the common frequencies lie in $\frac{1}{(q,r)}\mathbb{Z}$, the integers are blind by Lemma 12.3, and when $(q, r) = 2$ the half-integers are blind by (iii).*

Lemma 12.10 (the parity sieve)**.** *Let $q_1, \dots, q_k > 1$ be distinct squarefree conductors with multiplicities $e_1, \dots, e_k \geq 1$. If $\sum_{i:\, 2|q_i} e_i$ is odd, then*

$$\left\langle \sigma_{q_1}^{e_1} \cdots \sigma_{q_k}^{e_k} \right\rangle = 0.$$

Proof. If no q_i is even the count is zero and there is nothing to prove; otherwise $L = \operatorname{lcm}(q_1, \dots, q_k)$ is squarefree and even, so $2 \parallel L$. Split the product into two blocks, one factor against the rest. Both blocks are bounded periodic step functions, so Parseval applies with absolute convergence (Cauchy–Schwarz), and the coefficient of each block at a frequency ξ is an absolutely convergent sum over the solutions of $\sum u_{ij}/q_i = \xi$ with every u_{ij} coprime to its q_i (Proposition 10.5: each σ_{q_i} is supported on the reduced fractions of denominator exactly q_i). The two blocks pair at opposite frequencies, so every nonzero contribution to the mean comes from a *resonance*

$$\sum_{i=1}^{k} \sum_{j=1}^{e_i} \frac{u_{ij}}{q_i} = 0, \qquad (u_{ij}, q_i) = 1.$$

Multiply by L and reduce modulo 2. For even q_i the cofactor L/q_i is odd and so is u_{ij}: the term contributes 1. For odd q_i the cofactor L/q_i is even: the term contributes 0. The left side is therefore $\equiv \sum_{i:2|q_i} e_i \pmod 2$, odd by hypothesis, while the right side is 0. No resonance exists, and the mean vanishes. □

The sieve is the 2-adic mechanism of parts (iii)–(iv) of Lemma 12.9, promoted from supports to full means: the prime 2 is the only prime with a single unit, so a resonance must use an *even* number of frequency slots of even conductor. Applied to the sectors it halves the problem:

Corollary 12.11 (sector parities)**.** *In Proposition 12.8: every $(3,1)$ term with q even and r odd vanishes, so by Lemma 12.9(i) the $(3,1)$ sector is supported on $r \mid q$ with $2 \mid r$ whenever $2 \mid q$; every $(2,1,1)$ term with $r \not\equiv s \pmod 2$ vanishes; every $(1,1,1,1)$ term with an odd number of even conductors vanishes. The $(2,2)$ sector is untouched, its multiplicities being even.*

Proof. Count the even slots: $3[2 \mid q] + [2 \mid r]$ is odd exactly when q and r have opposite parity — the case q odd, r even is already void by Lemma 12.9(i), since then $r \nmid q$, leaving $2 \mid q$, $2 \nmid r$; $2[2 \mid q] + [2 \mid r] + [2 \mid s]$ is odd exactly when $r \not\equiv s \pmod 2$; $\sum_i [2 \mid q_i]$ odd is the stated condition. □

The corollary is sharp in every clause: $\langle \sigma_6^3 \sigma_3 \rangle = 0$ while $\langle \sigma_6^3 \sigma_2 \rangle = -\frac{1}{432}$; the means $\langle \sigma_q^2 \sigma_r \sigma_s \rangle$ with $r \not\equiv s$ vanish identically while the all-odd $\langle \sigma_{15}^2 \sigma_{21} \sigma_{35} \rangle = \frac{74}{55125}$ and $\langle \sigma_5^2 \sigma_3 \sigma_{15} \rangle = \frac{2}{1125}$ do not; and the *first* admissible quadruple, $\{2, 6, 10, 15\}$ at $D = 15$, carries three even conductors and dies, $\langle \sigma_2 \sigma_6 \sigma_{10} \sigma_{15} \rangle = 0$, while $\langle \sigma_2 \sigma_6 \sigma_{10} \sigma_{30} \rangle = -\frac{1}{1200}$ lives with four. Across the complete term lists to $D = 800$ (ancillary `xsectors.py`), no multiset violating the parity condition carries a nonzero mean. The sieve also re-proves, in one line, the two vanishings of Proposition 12.16 — $\langle \sigma_p \sigma_{2p}^3 \rangle$ and $\langle \sigma_2 \sigma_p \sigma_{2p}^2 \rangle$ each use three even slots — and several entries of the kill list in Theorem 12.18.

Lemma 12.12 (the coefficient bound)**.** *For squarefree $q > 1$ and $0 \neq \xi \in \frac{1}{q}\mathbb{Z}$, the Fourier coefficient of σ_q^2 at ξ equals $-\widetilde{T}_q(q\xi)/(4\pi^2)$, where, with absolute convergence,*

$$\widetilde{T}_q(N) = \sum_{\substack{u+u'=N \\ (uu',q)=1}} \frac{1}{uu'} = \frac{2\pi}{Nq} \sum_{a \in R_q(N)} \cot \frac{\pi a}{q}, \qquad R_q(N) = \{a \bmod q : \ (a(N-a), q) = 1\},$$

and $|\widetilde{T}_q(N)| \leq 4(1 + \log q)/|N|$.

Proof. The coefficient is the convolution of the coefficients of σ_q, and $\sum_u 1/(|u||N-u|) < \infty$. The involution $u \mapsto N - u$ preserves the index set and swaps the partial fractions $\frac{1}{uu'} = \frac{1}{N}(\frac{1}{u} + \frac{1}{u'})$, so $\widetilde{T}_q(N) = \frac{2}{N} \lim_U \sum_{|u| \leq U} 1/u$ over the set (the involution maps the window $|u| \leq U$ to $|u - N| \leq U$,

at a cost $O(N/U)\to 0$), which is $\frac{2}{N}\sum_{a\in R_q(N)}\frac{\pi}{q}\cot\frac{\pi a}{q}$ by $\sum_{u\equiv a\,(q)}\frac{1}{u}=\frac{\pi}{q}\cot\frac{\pi a}{q}$. The bound follows from $\cot x\le 1/x$ on $(0,\pi/2)$: $\sum_{a=1}^{q-1}|\cot\frac{\pi a}{q}|\le\frac{2q}{\pi}(1+\log q)$. □

Theorem 12.13 (two sectors are tame)**.** *Unconditionally,*

$$|X_{31}(D)|\ll D(\log D)(\log\log D)^6,\qquad |X_{22}(D)|\ll D(\log D)^2(\log\log D)^4.$$

Hence $X(D)=X_{211}(D)+X_{1111}(D)+O(D\log^4 D)$: the sectors with at most two distinct conductors obey the conjectured bound up to logarithmic factors, and Conjecture 12.19 reduces, up to those factors, to the sectors with at least three distinct conductors.

Proof. For X_{31}: by Lemma 12.9(i) the sum has $\ll\sum_{q\le D}d(q)\ll D\log D$ terms, each at most $|c_q|^3|c_r|\cdot\|\sigma_q\|_\infty^2(\langle\sigma_q^2\rangle\langle\sigma_r^2\rangle)^{1/2}\ll(\log\log D)^6$ by $\|\sigma_q\|_\infty\le\frac12\prod_{p\mid q}(1+p^{-1})$ and $|c_q|\le q/\varphi(q)$. For X_{22}: by Lemmas 12.9 and 12.12, for $g=(q,r)>2$,

$$|\operatorname{Cov}(\sigma_q^2,\sigma_r^2)|\le\sum_{0\ne\xi\in\frac1g\mathbb{Z}}\frac{4(1+\log q)}{4\pi^2q|\xi|}\cdot\frac{4(1+\log r)}{4\pi^2r|\xi|}=\frac{2\zeta(2)}{\pi^4}\,\frac{g^2(1+\log q)(1+\log r)}{qr},$$

and summing $g^2/(qr)$ over $q=gq_1,r=gr_1\le D$ gives $\sum_{g\le D}\big(\sum_{n\le D/g}1/n\big)^2\ll D$, so the weighted total is $\ll D(\log D)^2(\log\log D)^4$. □

The same spectral method reaches into the $(2,1,1)$ sector. The key is the analogue of Lemma 12.12 for a *pair* of distinct conductors.

Lemma 12.14 (the pair-coefficient bound)**.** *Let $r\ne s$ be squarefree conductors >1, $g=(r,s)$, $r_1=r/g$, $s_1=s/g$. The Fourier coefficients of $\sigma_r(x/r)\,\sigma_s(x/s)$ are supported on the ξ whose reduced denominator is divisible by r_1s_1 and divides gr_1s_1; at any $\xi\ne0$ with $rs\,\xi\in\mathbb{Z}$ the coefficient equals $-W_{r,s}(\xi)/(4\pi^2)$, where*

$$W_{r,s}(\xi)=\sum_{\substack{u/r+v/s=\xi\\(u,r)=(v,s)=1}}\frac{1}{uv}\qquad\textit{(absolutely convergent)},\qquad |W_{r,s}(\xi)|\le\frac{2(1+\log rs)(r+s)}{rs\,|\xi|}.$$

Proof. The coefficient at ξ is the convolution $\sum_{u/r+v/s=\xi}\frac{1}{2\pi iu}\cdot\frac{1}{2\pi iv}$, and the sum converges absolutely because $v=(rs\xi-us)/r$ grows linearly in u. The support: for a prime $p\mid r_1$ we have $p\nmid us$ and $p\,\|\,r$, so p survives at full power in the reduced denominator of $u/r+v/s$, and symmetrically for s_1; all denominators divide $\operatorname{lcm}(r,s)=gr_1s_1$. For the bound, put $N=rs\,\xi=us+vr\in\mathbb{Z}\setminus\{0\}$ and use the partial fractions $\frac{1}{uv}=\frac1N\big(\frac ru+\frac sv\big)$. The solution set is a union of residue classes modulo rs in u (the congruence $us\equiv N\pmod r$), the condition $(u,r)=1$, and the condition $(v,s)=1$ are all determined by u mod rs), and v runs over a union of classes modulo rs as well. Evaluating each half in symmetric windows — the window $|u|\le U$ transports to a window in v centred at N/r, at a recentring cost $O(N/(Us))\to0$ — the classical $\sum_{u\equiv a\,(M)}\frac1u=\frac{\pi}{M}\cot\frac{\pi a}{M}$ and $\sum_{a=1}^{M-1}|\cot\frac{\pi a}{M}|\le\frac{2M}{\pi}(1+\log M)$ give $|\sum{}'1/u|,\ |\sum{}'1/v|\le2(1+\log rs)$, hence $|W|\le2(1+\log rs)(r+s)/|N|$. □

Proposition 12.15 (the $(2,1,1)$ sector is $O(D^{3/2+o(1)})$)**.** *Unconditionally,*

$$|X_{211}(D)|\ \ll\ D^{3/2}\,(\log D)^{7/2}\,(\log\log D)^{O(1)}.$$

Proof. A term $(q;r<s)$ of X_{211} is nonzero only if $r_1s_1\mid q$, where $g=(r,s)$, $r_1=r/g$, $s_1=s/g$: the support of $\sigma_r\sigma_s$ sits over denominators divisible by r_1s_1 (Lemma 12.14) and the support of σ_q^2 over denominators dividing q. Write $q=r_1s_1h$ and $E=(g,h)$; the parametrization $(g,r_1<s_1,h)\mapsto(q;r,s)$ is injective, and the constraints are $gs_1\le D$, $r_1s_1h\le D$. Each term carries two bounds. *Trivial:* $|\langle\sigma_q^2\sigma_r\sigma_s\rangle|\le\|\sigma_q\|_\infty^2\sqrt{v_rv_s}\le\tau:=C(\log\log D)^2$, by $\|\sigma_q\|_\infty\le\frac12\prod_{p\mid q}(1+p^{-1})$ and

$v_q \le \frac{1}{12}$. *Spectral:* the common support of σ_q^2 and $\sigma_r\sigma_s$ lies in $\frac{1}{r_1s_1E}\mathbb{Z}$, so by Parseval, Lemma 12.12 at $q\xi = hk/E$ and Lemma 12.14 at $rs\,\xi = g^2k/E$ (writing $\xi = k/(r_1s_1E)$, $k \neq 0$),

$$|\langle \sigma_q^2\sigma_r\sigma_s\rangle| \le \sum_{k\neq 0} \frac{(1+\log q)E}{\pi^2 h|k|}\cdot\frac{(1+\log rs)(r_1+s_1)E}{2\pi^2 g|k|} \le C(\log D)^2\,\frac{E^2(r_1+s_1)}{hg},$$

where $(r+s)/(rs)\cdot r_1s_1E = (r_1+s_1)E/g$ was used. Since $|c_q^2c_rc_s| \ll (\log\log D)^4$, it remains to sum $\min(\tau, S)$ over the parameters. Fix (g, r_1, s_1) and sum over $h = E\tilde{h}$, $E \mid g$:

$$\sum_{h\le D/(r_1s_1)} \min\Big(\tau,\, C\log^2 D\,\frac{E(r_1+s_1)}{\tilde{h}g}\Big) \le \sum_{E\mid g} \min\Big(\tau\,\frac{D}{r_1s_1E},\, C\log^3 D\,\frac{E(r_1+s_1)}{g}\Big) \le d(g)\sqrt{\tau C\log^3 D}\,\sqrt{\frac{2D}{r_1g}},$$

using $\min(x,y) \le \sqrt{xy}$, in which E cancels, and $(r_1+s_1)/(r_1s_1) \le 2/r_1$. The number of admissible s_1 is at most $\min(D/g, D/r_1)$, so, with the per-term logarithms absorbed into $\mathcal{L} = (\log D)^{3/4}(\log\log D)^{O(1)}$,

$$|X_{211}(D)| \ll \mathcal{L}^2\sum_{g\le D} d(g)\sqrt{\frac{D}{g}}\sum_{r_1\le D}\frac{1}{\sqrt{r_1}}\min\Big(\frac{D}{g},\frac{D}{r_1}\Big) \ll \mathcal{L}^2\sum_{g\le D} d(g)\sqrt{\frac{D}{g}}\cdot\frac{4D}{\sqrt{g}} \ll D^{3/2}\mathcal{L}^2\sum_{g\le D}\frac{d(g)}{g},$$

and $\sum_{g\le D} d(g)/g \ll \log^2 D$: collecting powers, $|X_{211}| \ll D^{3/2}(\log D)^{3/2+2}(\log\log D)^{O(1)}$. (The inner estimate: $\sum_{r_1\le g} r_1^{-1/2}(D/g) + \sum_{r_1>g} r_1^{-3/2}D \le 4D/\sqrt{g}$.) □

This is far from the conjectured $O(D)$, but it is the first unconditional bound on a sector with three distinct conductors, and the proof shows exactly where the loss occurs: the absolute values surrender the sign of the cotangent sums in $W_{r,s}$. The data below suggest the truth is very much smaller.

The remaining structure is best seen in the data. In exact arithmetic (ancillary `sectors.py`):

D	6	10	14	15	21	22
X_{31}	+0.00617	+0.00712	+0.00636	+0.00810	+0.00840	+0.00727
X_{22}	−0.00309	−0.00429	−0.00491	−0.00045	−0.00059	−0.00084
X_{211}	−0.00617	−0.00712	−0.00636	+0.01077	+0.01521	+0.01465
X_{1111}	0	0	0	−0.00826	−0.02155	−0.02156
X	−0.00309	−0.00429	−0.00491	+0.01015	+0.00147	−0.00048

Two striking facts: $X_{31} + X_{211} = 0$ *exactly* for all $D \le 14$, so that below $D = 15$ the cross term *is* the $(2,2)$ sector; and the refuted $2p$ law lives entirely in X_{22}. Both have proofs.

Proposition 12.16 (the $2p$ ladder)**.** *For every odd prime p,*

$$\mathrm{Cov}(\sigma_p^2, \sigma_{2p}^2) = -\frac{p^2-1}{48\,p^4}, \qquad \langle\sigma_2\sigma_p^2\sigma_{2p}\rangle = \frac{p^2-1}{96\,p^3}, \qquad \langle\sigma_2\sigma_{2p}^3\rangle = -\frac{p^2-1}{128\,p^3},$$

$$\langle\sigma_p\sigma_{2p}^3\rangle = \langle\sigma_2\sigma_p\sigma_{2p}^2\rangle = 0, \qquad \textit{and in particular} \qquad 4\langle\sigma_2\sigma_{2p}^3\rangle + 3\langle\sigma_2\sigma_p^2\sigma_{2p}\rangle = 0.$$

Consequently the joint $(3,1)+(2,1,1)$ contribution of the cluster $\{2,p,2p\}$ equals

$$12\,c_2c_{2p}\,\langle\sigma_2\sigma_p^2\sigma_{2p}\rangle\Big(c_p^2 - \frac{c_{2p}^2}{4}\Big),$$

which vanishes — for any *value of c_2 — whenever $m_p(D/p) = \frac{1}{2}\,m_{2p}(D/2p)$.*

Proof of Proposition 12.16. On $(j, j+1)$, $0 \le j < 2p$, the three functions take the explicit values

$$\sigma_p = v_j = \frac{2(j \bmod p) + 1 - p}{2p}, \qquad \sigma_2 = -\frac{(-1)^j}{4}, \qquad \sigma_{2p} = w_j = \frac{2j+1-2p}{4p} + \frac{(-1)^j}{4p} - \frac{v_j}{2}$$

(the last from $B_1(t\,\frac{2j+1}{4p})$, $t \mid 2p$). Each identity is a sum over j of a polynomial in $j, (-1)^j, v_j$, *defined by these formulas for every odd* p; splitting j by the Chinese remainder pair ($j \bmod 2$, $j \bmod p$) and applying Faulhaber's formulas, each side, cleared of its fixed power of p, becomes a polynomial in p of degree at most 24. The identities hold in exact arithmetic at the twenty-nine odd values $3 \le p \le 59$ (ancillary `ladder.py`, which also checks the value formulas against the true σ_p, σ_{2p}), hence identically in p; specializing to odd primes gives the statement. □

Lemma 12.17 (double-ladder orthogonality)**.** *For distinct odd primes* $r \neq p$,

$$\langle \sigma_r \sigma_{2r} \sigma_p \sigma_{2p} \rangle = 0.$$

Proof. On $(j, j+1)$ write $\varepsilon = (-1)^j$ and $\rho = j \bmod r$. From $j \bmod 2r = \rho + r\,\frac{1-\varepsilon(-1)^\rho}{2}$, the value formula of Proposition 12.16 splits as $\sigma_{2r} = A_r(\rho) + \varepsilon B_r(\rho)$ with $B_r(\rho) = \frac{1}{4r} - \frac{(-1)^\rho}{4}$, and $\sigma_r = v_r(\rho)$. As j runs over $2rp$ consecutive integers the triple (ε, $j \bmod r$, $j \bmod p$) is uniform on the product (Chinese remainders), so averaging the product $v_r(A_r + \varepsilon B_r)\, v_p(A_p + \varepsilon B_p)$ kills the ε-odd part and factorizes the rest over the two independent residues:

$$\langle \sigma_r \sigma_{2r} \sigma_p \sigma_{2p} \rangle = \overline{v_r A_r}\;\overline{v_p A_p} + \overline{v_r B_r}\;\overline{v_p B_p}.$$

Both one-prime means vanish: $\overline{v_r A_r} = \langle \sigma_r \sigma_{2r} \rangle = 0$ by orthogonality, and $\overline{v_r B_r} = \frac{1}{4r}\overline{v_r} - \frac{1}{4r}\sum_{\rho<r} v_r(\rho)(-1)^\rho = 0$ since $\overline{v_r} = 0$ and $\sum_{\rho=0}^{r-1}(-1)^\rho(2\rho + 1 - r) = 2 \cdot \frac{r-1}{2} + (1-r) = 0$ for odd r. The vanishing is special to the double-ladder structure: for instance $\langle \sigma_3 \sigma_6 \sigma_{14} \sigma_{21} \rangle = \frac{1}{2646} \neq 0$. □

Theorem 12.18 (the $2p$ jumps, proved and explained)**.** *At* $D = 2p$, p *an odd prime, the* $(2,2)$ *sector jumps by exactly*

$$6\, c_p^2 c_{2p}^2 \operatorname{Cov}(\sigma_p^2, \sigma_{2p}^2) = -\frac{p^2-1}{32\,p^4} \qquad \text{for every } p$$

($c_p = \mu(p) m_p(2) = -\frac{1}{2}$ and $c_{2p} = 1$ at arrival; no older conductor shares an odd prime with $2p$, and conductor 2 is inert). For $p = 3, 5, 7$ *the* joint *contribution of the other three sectors does not move — the individual sectors do move, as the table above shows, but their changes cancel exactly: every new multiset involving $2p$ is annihilated by Proposition 12.2, Lemma 12.9, Proposition 12.16 or Lemma 12.17, except the cluster pair, whose total vanishes by the cluster identity with the relation $m_{p'}(D/p') = \frac{1}{2}$, valid for every live odd prime throughout $D \le 14$. Hence*

$$\Delta X(2p) = -\frac{p^2-1}{32p^4} \qquad \text{exactly, for } p = 3, 5, 7.$$

From $D = 15$ the weight relation breaks ($m_3(5) = \frac{3}{10}$, $m_5(3) = \frac{1}{6}$), and at $D = 22$ the other sectors move through the new term $-\frac{1}{11}$ in $m_2(11)$, by the exact amounts $-\frac{76512673}{67795150500}$ (sector 31), $-\frac{37687127}{67795150500}$ (211) and $-\frac{16}{1414875}$ (1111):

$$\Delta X(22) = -\frac{11^2-1}{32 \cdot 11^4} - \frac{7198}{4244625} = -\frac{44110277}{22598383500} = -0.0019519218\ldots$$

The law, its three exact successes, and its failure at $p = 11$ are thereby explained term by term.

Proof. The $(2,2)$ statement: the new pairs at $D = 2p$ are $\{2p, q\}$, $q < 2p$, and $\operatorname{Cov}(\sigma_{2p}^2, \sigma_q^2) = 0$ unless q shares an odd prime with $2p$ (Lemma 12.9), i.e. $q = p$; older pairs can move only through c_2 and c_p, but conductor 2 is inert and no older pair involving p has a partner sharing the prime p. The weights at arrival are $m_p(2) = 1 - \frac{1}{2}$ and $m_{2p}(1) = 1$, giving $6 \cdot \frac{1}{4} \cdot (-\frac{p^2-1}{48p^4})$. For the other sectors at $p \le 7$: a new non-$(2,2)$ multiset uses $2p$ and, by Proposition 12.2, must contain p or a second even conductor. The complete list of multisets surviving Proposition 12.2, and its fate: $\{p, 2p, 2p, 2p\}$ and $\{2, p, 2p, 2p\}$, zero by the vanishings of Proposition 12.16; $\{2, 2, p, 2p\}$, zero since σ_2^2 is constant and $\langle \sigma_p \sigma_{2p} \rangle = 0$; the cluster pair $\{2, 2p, 2p, 2p\}$ together with $\{2, p, p, 2p\}$ —

jointly $4c_{2p}^3c_2\langle\sigma_2\sigma_{2p}^3\rangle + 12c_p^2c_2c_{2p}\langle\sigma_2\sigma_p^2\sigma_{2p}\rangle$, which vanishes at $c_p^2 = \frac{1}{4}c_{2p}^2$; the multisets $\{q,q,p,2p\}$ and $\{q,q,r,2r\}$ with q even (older, or $q = 2p$ itself) and r odd, which die because σ_q^2 carries no half-integer and no nonzero integer frequency while $\sigma_p\sigma_{2p}$ (resp. $\sigma_r\sigma_{2r}$) is supported on denominators $2p$ and 2 (resp. $2r$ and 2); and the double ladders $\{r,2r,p,2p\}$, $r<p$ odd prime, zero by Lemma 12.17. Reweighting of older multisets at $D = 2p$ goes only through Δc_2, Δc_p and Δm_1, the last invisible by Theorem 12.4; below $D = 15$ the only nonzero older non-$(2,2)$ multisets are the clusters $\{2,p',2p'\}$ — the older double ladders vanish by Lemma 12.17 — and the clusters are linear in c_2 and zero for any c_2 by the weight relation, while no nonzero older multiset contains p. The values at $D = 22$ are exact computations over the periods of the multisets (ancillary `sectors.py`). □

A prediction worth recording: the $(2,2)$ jump law holds at *every* prime, so at $D = 26$ the $(2,2)$ sector must move by exactly $-(13^2-1)/(32\cdot 13^4) = -\frac{21}{114244}$, whatever the other sectors do.

12.6. **The four sectors at scale.** The sector means $\langle\cdot\rangle$ do not depend on D — only the coefficients $c_q = \mu(q)m_q(D/q)$ do — so each admissible multiset need be integrated once, over its own period, and the four sectors can then be swept in D through the coefficients alone. With the support prefilters ($r \mid q$ for $(3,1)$; an odd common prime for $(2,2)$; $r_1s_1 \mid q$ for $(2,1,1)$; every prime dividing at least two conductors for $(1,1,1,1)$) this reaches $D = 800$ in exact structure and double precision (ancillary `xsectors.py`; the values at $D \le 22$ reproduce the exact rationals of the table above to 10^{-15}). There are 78 336 multisets with nonzero mean in the $(2,1,1)$ sector and 146 211 in the $(1,1,1,1)$ sector up to $D = 800$; the parity sieve of Corollary 12.11 is confirmed across the complete term lists.

D	X_{31}	X_{22}	X_{211}	X_{1111}	$X_{211}+X_{1111}$	$\Sigma\lvert\cdot\rvert_{211}$	$\Sigma\lvert\cdot\rvert_{1111}$
100	+0.0291	−0.0138	+0.0859	−0.0673	+0.0186	0.25	0.68
200	+0.0656	−0.0215	+0.0990	−0.0254	+0.0736	0.64	2.15
300	+0.0918	−0.0288	+0.0419	−0.1149	−0.0730	0.99	4.46
400	+0.1203	−0.0439	+0.1694	−0.4907	−0.3212	1.45	7.33
500	+0.1581	−0.0582	+0.1988	−0.6003	−0.4016	1.86	10.47
600	+0.1864	−0.0720	+0.3129	−0.6391	−0.3261	2.17	13.82
700	+0.2227	−0.0859	+0.3430	−0.5301	−0.1871	2.66	18.08
800	+0.2537	−0.0906	+0.2983	−0.3559	−0.0576	3.07	22.66

The last two columns are the sums of *absolute* term contributions $12\sum|c_q^2c_rc_s\langle\cdot\rangle|$ and $24\sum|c_qc_rc_sc_t\langle\cdot\rangle|$. Three facts stand out (Figure 11).

The conjecture has enormous room. The running maximum of $|X_{211}+X_{1111}|$ on $2 \le D \le 800$ is 0.438, attained at $D = 506$; at $D = 800$ the ratio $(X_{211}+X_{1111})/D$ is $-7\cdot 10^{-5}$. The measured log–log slopes on $100 \le D \le 800$ (window fits, not asymptotics) are ≈ 1.0 for X_{31} and X_{22}, ≈ 1.2 for the amplitude of X_{211}, ≈ 1.5 for X_{1111} — and ≈ 0.26 for the total $X(D)$, which at $D = 800$ is $+0.105$: the four sectors cancel one another to a total that grows, in this window, like $D^{1/4}$.

The $(1,1,1,1)$ *sector cannot be tamed by absolute values.* The absolute sums grow like $D^{1.2}$ ($(2,1,1)$) and $D^{1.65}$ ($(1,1,1,1)$) in the window. Both are still far below D at $D = 800$, but the trend of the second crosses every linear bound: if it persists, no estimate that surrenders the signs of the quadruple means — including the method of Proposition 12.15 — can prove $X_{1111} = O(D)$. The conjecture, if true, is true because of sign cancellation *among* the quadruples, not within them.

The individual means are small in a structured way. Across all 78 336 nonzero $(2,1,1)$ means, the normalized size $|\langle\sigma_q^2\sigma_r\sigma_s\rangle|\cdot gr_1s_1$ has median 10^{-4} and 99th percentile ≈ 1.3, so the typical mean behaves like $(gr_1s_1)^{-1}$ — summing that model over the admissible parameters gives $O(D\log D)$, which is why we read the $D^{3/2}$ of Proposition 12.15 as an artefact of the cotangent absolute values — but a thin tail (maximum 21.7, at $(q;r,s) = (253;374,782)$) keeps the uniform model false. In the $(1,1,1,1)$ sector the largest mean in absolute value is $3.3\cdot 10^{-3}$ and does not grow across the

range, while $|\langle\sigma_q\sigma_r\sigma_s\sigma_t\rangle| \cdot \mathrm{lcm}(q,r,s,t)$ has 99th percentile ≈ 19: the quadruple means live at the scale of one over their own period.

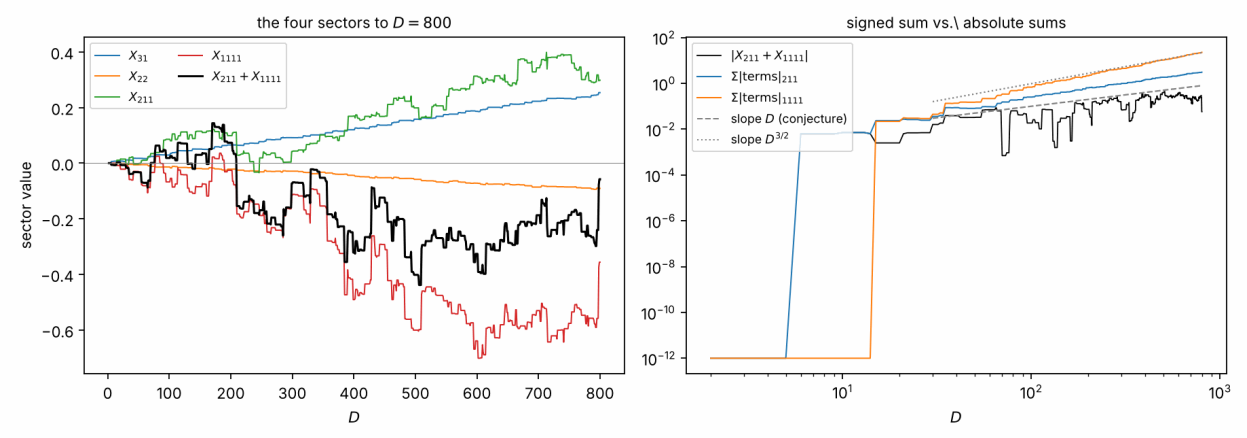

FIGURE 11. The four sectors to $D = 800$. Left: signed values; the black curve is $X_{211} + X_{1111}$, the object of Conjecture 12.19 after Theorem 12.13. Right: $|X_{211} + X_{1111}|$ against the absolute-value sums of the two sectors, log–log, with the conjectured slope D and the proved $D^{3/2}$ for scale.

12.7. **What remains.**

Conjecture 12.19. $X(D) = O(D)$.

The intra part needs no conjecture: $|m_q(Y)| \le q/\varphi(q)$ — from $m_q(Y) = \sum_s s^{-1}\,u(Y/s)$, the sum running over the q-smooth integers s (every prime factor of s divides q), where $u(Z) = \sum_{m\le Z}\mu(m)/m$ satisfies the elementary $|u| \le 1$ and $\sum_s 1/s = \prod_{p\mid q}(1-p^{-1})^{-1} = q/\varphi(q)$ — and $|\varepsilon_4(q)| \ll (\log\log q)^2$, so $I(D) \ll D(\log\log D)^6$ unconditionally, with the data saying $I(D) \approx -0.0027\,D$. Conjecture 12.19 would therefore give $e_4(D) \ll D(\log\log D)^6$, and since $V(D) \ge \sum_{D/2<k\le D}\mu^2(k)J_2(k)/k^2 \gg D$ by positivity, excess kurtosis $O\big(D^{-1}(\log\log D)^6\big)$ — the measured law is the cleaner $-3/D$. By Theorem 12.13 the two sectors with at most two distinct conductors are $O(D\log^4 D)$ — the conjectured bound up to logarithmic factors — so what remains is $X_{211}(D) + X_{1111}(D) = O(D)$, which would give $X(D) = O(D\log^4 D)$. Proposition 12.15 settles a first piece unconditionally, $X_{211} \ll D^{3/2+o(1)}$, so the open problem now has two named parts: improve the $(2,1,1)$ bound from $D^{3/2}$ to D — a problem about the true size of the pair sums $W_{r,s}$, which the data of Section 12.6 put at $(gr_1s_1)^{-1}$ for the typical mean, a model that sums to $O(D\log D)$ — and bound X_{1111}, where the same data say absolute values cannot suffice and the cancellation *among* quadruple resonances must enter. Unlike everything else in this circle, Conjecture 12.19 is *not* the Riemann hypothesis in disguise: it concerns fixed-D averages of periodic functions — finitely many conductor multisets at each D, each a finite exact computation — and the intra-orbit part of the problem is already settled exactly by (31) and Lemma 12.1. Its proof would make the $-3/D$ law unconditional. It would not touch the moving diagonal $D = x$.

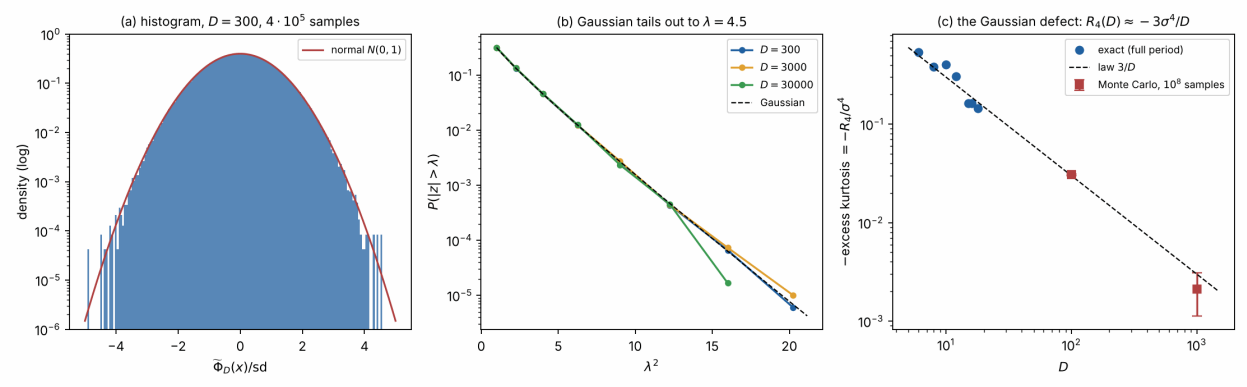

FIGURE 12. Gaussian statistics. (a) Histogram at $D = 300$ against the normal density, log scale. (b) Tails against the Gaussian law, to $\lambda = 4.5$. (c) The Gaussian defect of the fourth moment: exact values for $D \leq 18$, high-precision Monte Carlo at $D = 100, 1000$, and the law $-3/D$, which splits as intra -3.5 plus cross $+0.5$ by (31).

Part III
One constant, three mirrors

13. One constant

Write ρ for the nontrivial zeros of the Riemann zeta function, and

$$\widehat{B} := \sum_{\rho} \frac{1}{\rho} = 1 + \frac{\gamma}{2} - \frac{1}{2}\log 4\pi = 0.023095708966\ldots, \tag{32}$$

the sum taken with zeros paired with their conjugates, as usual; the closed form is classical (see [Dav, §12]). Three statements, proved in three different languages, put this one number at the centre of the Riemann hypothesis:

(1) (*primes*) RH $\iff$ a specific weighted mean of the error $\psi(t) - t$ stays below $-\widehat{B}$ for ever (Section 15);
(2) (*Möbius*) RH $\iff$ the constant function 1 can be approximated in a weighted L^2 by Möbius-weighted combinations of the sawtooths $\{1/nt\}$ — and the squared error is expected to decay exactly like $2\widehat{B}/\log N$ (Section 16);
(3) (*zeros*) the doubling $2\widehat{B} = \sum_\rho 1/(\rho(1-\rho))$ is an identity, equal to $\sum_\rho |\rho|^{-2}$ precisely on RH (Section 14).

The three languages are the three corners of the convolution identity $\Lambda = \mu * \log$: prime side, Möbius side, zero side. The purpose of this part is modest and expository-plus-numerical: to state the three mirrors with matching normalizations — the halving between $\widehat{B}$ and the Báez-Duarte constant $2 + \gamma - \log 4\pi$ seems not to be pointed out explicitly in the literature — to recompute the Nyman–Beurling distances from scratch, and to record how early the constant becomes visible in the data. We claim no progress towards RH, and the last section of this part says precisely why none of the three mirrors, verified however far, could supply any.

Throughout, γ is Euler's constant, $\psi(t) = \sum_{q \le t} \Lambda(q)$, and all numerical work is reproducible from the ancillary code.

14. The mirror of the zeros

Pairing ρ with $1 - \rho$ (the functional equation preserves the zero set),

$$\sum_{\rho} \frac{1}{\rho(1-\rho)} = \sum_{\rho} \Big(\frac{1}{\rho} + \frac{1}{1-\rho}\Big) = 2\sum_{\rho} \frac{1}{\rho} = 2\widehat{B} = 2 + \gamma - \log 4\pi = 0.046191417\ldots, \tag{33}$$

an unconditional, absolutely convergent identity. On RH one has $\bar\rho = 1 - \rho$, so $|\rho|^2 = \rho\bar\rho = \rho(1-\rho)$ and

$$\sum_{\rho} \frac{1}{|\rho|^2} = 2\widehat{B} \qquad \text{(on RH)}. \tag{34}$$

Moreover (34) is itself *equivalent* to RH: each off-line quadruple $\{\rho, 1-\rho, \bar\rho, 1-\bar\rho\}$, $\rho = \beta + i\gamma$ with $\beta \neq \frac{1}{2}$, contributes the strictly positive excess $2(2\beta-1)^2/(|\rho|^2|1-\rho|^2)$ to $\sum_\rho |\rho|^{-2} - \sum_\rho 1/(\rho(1-\rho))$, so $\sum_\rho |\rho|^{-2} \ge 2\widehat{B}$ always, with equality if and only if every zero lies on the critical line. The third mirror is thus an equivalence, exactly like the other two. The quantity $\sum_\rho |\rho|^{-2}$ is exactly the constant of the Nyman–Beurling asymptotics below: the inverse-square energy of the zeros.

15. The mirror of the primes: the Cesàro margin

This mirror is Part I: RH is equivalent to $M(N) \geq 0$ for all large N (a theorem of Suzuki [Su1]), and by the mean identity of Section 7, with $\mu_K(N)$ the K_N-weighted mean of the prime-counting error E^*,

$$M(N) = -2\log N\,\big(\mu_K(N) + \widehat{B}\big), \qquad \text{RH} \iff \mu_K(N) \leq -\widehat{B} \text{ for all large } N, \tag{35}$$

and the verification of Part I gives $\mu_K(N) \leq -\widehat{B} - 0.0171/\log N$ throughout $10^3 \leq N \leq 2.15 \cdot 10^{29}$. The mean of the prime-counting error must hug $-\widehat{B}$ from below, for ever.

16. The mirror of the Möbius sawtooths: Nyman–Beurling

The Nyman–Beurling criterion [Ny, Be], in Báez-Duarte's strengthening to integer dilations [BD], states that RH holds if and only if the constant function 1 lies in the closed span of the dilated sawtooths $t \mapsto \{1/nt\}$, $n \geq 1$, in the natural weighted L^2. In the equivalent Dirichlet-polynomial normalization used by [BCF],

$$d_N^2 = \inf_{A_N} \frac{1}{2\pi} \int_{-\infty}^{\infty} \big|1 - \zeta A_N\big(\tfrac{1}{2} + it\big)\big|^2 \, \frac{dt}{\frac{1}{4} + t^2}, \qquad A_N(s) = \sum_{n \leq N} \frac{a_n}{n^s}, \tag{36}$$

and RH $\iff$ $d_N \to 0$. Burnol [Burn] proved the unconditional lower bound

$$\liminf_{N \to \infty} d_N^2 \log N \;\geq \sum_{\Re\rho = 1/2} \frac{m(\rho)^2}{|\rho|^2}, \tag{37}$$

$m(\rho)$ the multiplicity, sharpening [BBLS]; on RH with simple zeros the right side is $\sum_\rho |\rho|^{-2} = 2\widehat{B}$ by (34), and this value is expected to be the true asymptotic:

$$d_N^2 \;\sim\; \frac{2\widehat{B}}{\log N} = \frac{2 + \gamma - \log 4\pi}{\log N}. \tag{38}$$

Bettin, Conrey and Farmer [BCF] proved that, under RH and a mild growth hypothesis on $\sum_\rho |\zeta'(\rho)|^{-2}$, the asymptotic (38) is already attained by the explicit polynomial

$$V_N(s) = \sum_{n \leq N} \Big(1 - \frac{\log n}{\log N}\Big) \mu(n)\, n^{-s}\,: \tag{39}$$

a *Cesàro-smoothed Möbius polynomial* — the smoothing of Part I and the Möbius weights of Part II in a single object. Thus the Báez-Duarte constant *is* $2\widehat{B}$: the margin constant of Section 15, doubled by the pairing $\rho \leftrightarrow 1 - \rho$. We have not found this halving stated explicitly, trivial as it is; it is the bridge between the one-sided prime criterion and the two-sided L^2 criterion.

17. The computation

We recomputed the Nyman–Beurling quantities from scratch. Expanding (36), with $\langle f, g\rangle = \frac{1}{2\pi}\int f\bar{g}\, dt/(\frac{1}{4} + t^2)$ on the critical line, one needs $\|1\|^2 = 1$, the linear functionals $c_n = \langle 1, \zeta n^{-s}\rangle$, and the Gram matrix $G_{mn} = \langle \zeta m^{-s}, \zeta n^{-s}\rangle$.

The linear term. Writing $\frac{1}{\frac{1}{4}+t^2} = \frac{1}{s(1-s)}$ with $s = \frac{1}{2} + it$,

$$c_n = \frac{1}{2\pi i} \int_{(1/2)} \frac{\zeta(s)\, n^{-s}}{s(1-s)}\, ds = -\operatorname*{Res}_{s=1} \frac{\zeta(s)\, n^{-s}}{s(1-s)} = \frac{\gamma - 1 - \log n}{n},$$

by closing the contour to the right across the double pole at $s = 1$ (expand $\zeta(s) = \frac{1}{s-1} + \gamma + O(s-1)$ and $n^{-s} = n^{-1}(1 - (s-1)\log n + \cdots)$).

The Gram matrix. By Mellin–Plancherel, G_{mn} is the autocorrelation $\int_0^\infty \{1/mt\}\{1/nt\}\,dt$, for which Vasyunin [Vas] (see also [BBLS, p. 141]) gave the exact closed form: for coprime m, n,

$$mn\int_0^\infty \Big\{\frac{1}{nt}\Big\}\Big\{\frac{1}{mt}\Big\}dt = \frac{\log 2\pi-\gamma}{2}(m+n) + \frac{m-n}{2}\log\frac{n}{m} - \frac{\pi}{2}\big(S(m,n)+S(n,m)\big), \tag{40}$$

$$S(m,n) = \sum_{k=1}^{n-1}\Big\{\frac{mk}{n}\Big\}\cot\frac{k\pi}{n},$$

and $G_{mn} = g^{-1}G_{m'n'}$ for $g=(m,n)$, $m=gm'$, $n=gn'$. We validated (40) against direct numerical integration of $|\zeta(\frac12+it)|^2(n/m)^{it}/(\frac14+t^2)$ out to $t=6000$ at six pairs, with agreement to the size of the truncation tail ($\sim 10^{-3}$), and checked that the resulting Gram matrices are positive definite. Then $d_N^2 = 1 - c^{\mathsf T}G^{-1}c$ over all $n \le N$, solved by spectral decomposition with monitored conditioning (relative residuals below 10^{-13} throughout the table), and the energy of the Bettin–Conrey–Farmer polynomial V_N is evaluated directly, with no inversion, over the squarefree indices.

N	20	40	80	160	320	500	640	800
$d_N^2 \log N$	0.0494	0.0469	0.0481	0.0474	0.0459	0.0461	0.0453	0.0449

N	100	400	700	1000	1500	2000
$\|1-\zeta V_N\|^2 \log N$	0.822	0.640	0.587	0.562	0.533	0.512

Two readings (Figure 13). The optimal distance locks onto the constant almost immediately: $d_N^2 \log N$ oscillates between 0.0448 and 0.0495 from $N=20$ on, against $2\widehat{B} = 0.04619$ — and the computation inputs *no zero of zeta*: only fractional parts, cotangent sums and two residues. The inverse-square energy of the zeros emerges from sawtooth geometry alone. One point of candour: the last two entries (0.0453, 0.0449) lie *below* the floor $2\widehat{B}$. Burnol's bound (37) constrains only the liminf, so a finite-N undershoot is logically harmless — the sequence must come back up — but it is a concrete reminder that finite data cannot certify even the floor, which is the moral of Section 18 in miniature. The natural arithmetic approximant V_N, by contrast, is still a factor ~ 11 above the constant at $N=2000$ and descending slowly; its asymptotic optimality [BCF] is invisible at accessible N — a known phenomenon (compare the graphs cited in [BCF]), and a reminder of how slowly the Möbius function surrenders its cancellation.

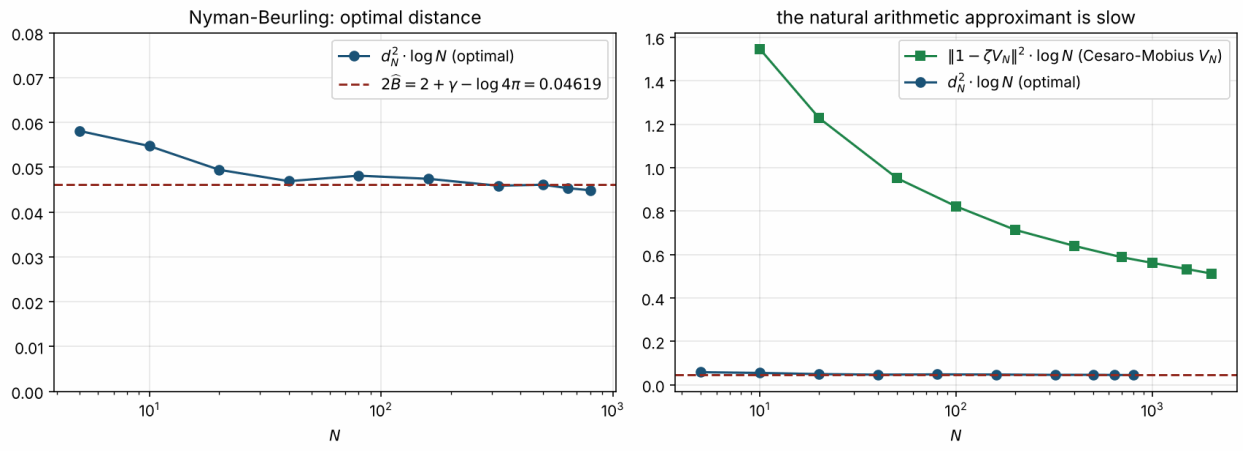

FIGURE 13. Left: $d_N^2 \log N$ against $2\widehat{B} = 0.04619$. Right: the same, with the Cesàro–Möbius polynomial V_N of (39). No zeros of ζ enter the computation.

18. The triangle, and what it cannot do

The convolution identity $\Lambda = \mu * \log$ has three corners, and $\widehat{B}$ is pinned at all three. From the zero corner it is a definition, (32). From the prime corner it is the exact level that the weighted mean of $\psi(t) - t$ must respect one-sidedly, (35). From the Möbius corner it is (doubled) the exact price of building the constant function out of sawtooths, (38). The three statements are not three theorems about three constants; they are one constant crossing the explicit formula in three directions.

None of this, verified however far, approaches a proof of RH, and the reasons are the same two walls described in Parts I and II: the prime mirror is a one-sided inequality to be held *for ever*, and no finite verification or mean-value bound excludes a late, measure-zero failure; the Möbius mirror is a limit statement whose rate carries RH only in the limit, and Burnol's bound (37) shows the constant $2\widehat{B}$ is a *floor*, not a certificate — reaching it numerically is consistent with RH, and proves nothing. What the triangle does provide is calibration: any proposed approach to RH through means, margins or L^2 distances must, at minimum, reproduce $\widehat{B}$ exactly at all three corners. It is a sharp sanity check, and we offer it as such.

Part IV
The cross-energy of Dirichlet partial sums

Let $X_n(s) = \sum_{k\le n} k^{-s}$ be the partial sums of the Dirichlet series of ζ. The *cross-energy* of consecutive partial sums,

$$C_2(n,s) = |X_n(s)|^2 - |X_{n-1}(s)|^2 - n^{-2\sigma} = 2\,\mathrm{Re}\big(X_{n-1}(s)\,\overline{n^{-s}}\big), \tag{41}$$

records how the newest term interacts with all previous ones. A twenty-two–part series of the author studies this object at the zeros of ζ; the parts on the strip, the energy, the slope, the Euler product and the functional reflection are consolidated, with complete proofs, in the companion monograph [CEn], while the window studies quoted in Part V are proved in [CE20, CE21, CE22]. This part states the main results — those the companion's own accounting marks as new, together with the elementary structure (Theorem 18.1) that frames them — and the threads that tie them to Parts I–III. Throughout this part γ_0 is Euler's constant, $\rho = \beta + i\gamma$ runs over the nontrivial zeros, $\Theta = \sup\beta$, and

$$C_B = \sum_\rho \frac{1}{\rho(1-\rho)} = 2 + \gamma_0 - \log 4\pi = 2\widehat{B},$$

the constant of Parts I and III once more. Proofs are in [CEn] at the stated locations; what is classical there is credited there (Cramér's theorem, Turing's method, Lagarias's pairing, Weil's criterion, the density theorems), and we repeat none of those attributions here.

18.1. **The trichotomy and the crossing law.**

Theorem 18.1 (the trichotomy; [CEn, Thm. A, §2])**.** *For $0<\sigma<1$,*

$$X_{n-1}(s)\,\overline{n^{-s}} = \frac{n^{1-2\sigma}}{1-s} + \zeta(s)\,n^{-\sigma}e^{it\log n} - \tfrac12 n^{-2\sigma} + O\big(|s|\,n^{-2\sigma-1}\big).$$

Consequently $\lim_n C_2(n,s)$ equals $1/(\frac14+t^2)$, 0 or $+\infty$ according as $\sigma=\frac12$, $\sigma>\frac12$ or $\sigma<\frac12$, at every point of the strip, and the zeros of ζ appear only in the second term. The first term is the trace of the pole at $s=1$: the critical line is where it survives exactly.

On the line, the planar curve described by $X_{n-1}\overline{n^{-s}}$ is a helix whose natural orientation gives it a *signed* radius, and that signed radius is Hardy's function $Z(t)$: each zero on the line is a crossing of the axis, and an off-line pair is a missing crossing (Turing's method). At a zero, $C_2(\cdot,\rho)$, extended to real n by the Hurwitz tail $X_{n-1}(s) = \zeta(s) - \zeta(s,n)$, changes sign exactly once beyond $|\gamma|^{3/2}$, at

$$n^* = \gamma^2 + \tfrac13 + \tfrac{1}{180\gamma^2} + O(\gamma^{-4}) \qquad (\beta=\tfrac12),$$

displaced by $(2\beta-1)\gamma^2/(2(1-\beta)) + O(1)$ otherwise [CEn, Prop. 2.2]; in the variable $\lambda = s(1-s)$ the law is a universal series with rational coefficients,

$$n^* = \lambda + \frac{1}{12} + \frac{1}{180\lambda} - \frac{53}{5040\lambda^2} + \cdots,$$

and the X–Y energies obey the same rule: at a zero on the line $|X_N^C|^2 = N/\lambda - \frac{1}{6N} + (\frac{1}{90} + \frac{\lambda}{240})N^{-3} + \cdots$ to all orders, two zeros interact only through $\lambda_1 - \lambda_2$, and the normalised energy gap at a zero tends to $-\frac16\,\mathrm{Re}\,\lambda(\rho)/|1-\rho|^2$, which equals $-B_2$ exactly when $\beta=\frac12$. At a zero on the line $C_2(n,\rho) = 1/(\gamma^2+\frac14) - \frac1n - \frac{1}{12n^2} + O(|\gamma|^5 n^{-5})$ increases monotonically to its limit, and *RH is equivalent to the boundedness of $C_2(\cdot,\rho)$ at every zero* [CEn, Cor. 2.1]. Three arithmetical by-products [CEn, §2]: a zero detector on the critical line; the reading of Landau's formula on the circles of the Euler factors — placed on the circle $|p^{-s}| = p^{-1/2}$ of a prime p, the zeros lean toward the angle π, and toward $\pi + \arg\chi(p)$ for a Dirichlet L-function, while on the circle of e they are uniformly distributed; and, for the finite sums $\sum_{k\le N} k^{-s}$, a polygon condition in $\pi(\sqrt{N})$

angles governing the real parts of their zeros, with the rules that adding a prime never moves the right edge left, adding a power of 2 moves it left for $2^k \le 16$ — and 11^2 is an exception to the corresponding rule for prime powers.

18.2. **Reflected pairs and the functional reflection.** For a reflected pair ρ, $1-\bar\rho$ the weights $n^{-2\beta}$ and $n^{-2(1-\beta)}$ compensate.

Proposition 18.2 (reflected pairs; [CEn, Prop. 2.2′, 2.2″])**.** *The product $C_2(n,\rho)\,C_2(n,1-\bar\rho)$ tends to a finite positive limit for every β, and is a square on the line (at every point of the line, zero or not). Off the line, beyond $|\gamma|^{3/2}$, it is negative exactly between the two crossings, an interval that contains the crossing $n^*(\frac12+i\gamma)$ of the line, with minimum asymptotically $-(2\beta-1)^2/(4\beta(1-\beta))$ times the limit; for the zeros above an effective but uncomputed height, the critical line is characterized by the nonnegativity of this product. The product equals $(\operatorname{Re}\mathcal{C}_2)^2-(\operatorname{Re}\mathcal{W})^2$, with $\mathcal{C}_2(N;s,1-s)$ the two-variable cross-energy and $\mathcal{W}$ the Wronskian of consecutive partial sums, purely imaginary on the line. Summed over the zeros with a weight, the product is Weil's quadratic form at test functions built from the integers below n.*

Note the family resemblance to Part III: the off-line excess there was $2(2\beta-1)^2/(|\rho|^2|1-\rho|^2)$; the deformation parameter $(2\beta-1)^2$ is the common currency of every detector in this circle. The zeros of the truncations spiral into their limits: with the endpoint term $\frac12 N^{-s}$ restored, along helices of radius $|\rho|N^{-\beta-1}/(12|\zeta'(\rho)|)$, proved with explicit constants and verified for the first 100 zeros; for the Davenport–Heilbronn function the two zeros of an off-line pair converge at rates that differ by $2(\beta-\frac12)$ [CEn, Prop. 2.3′].

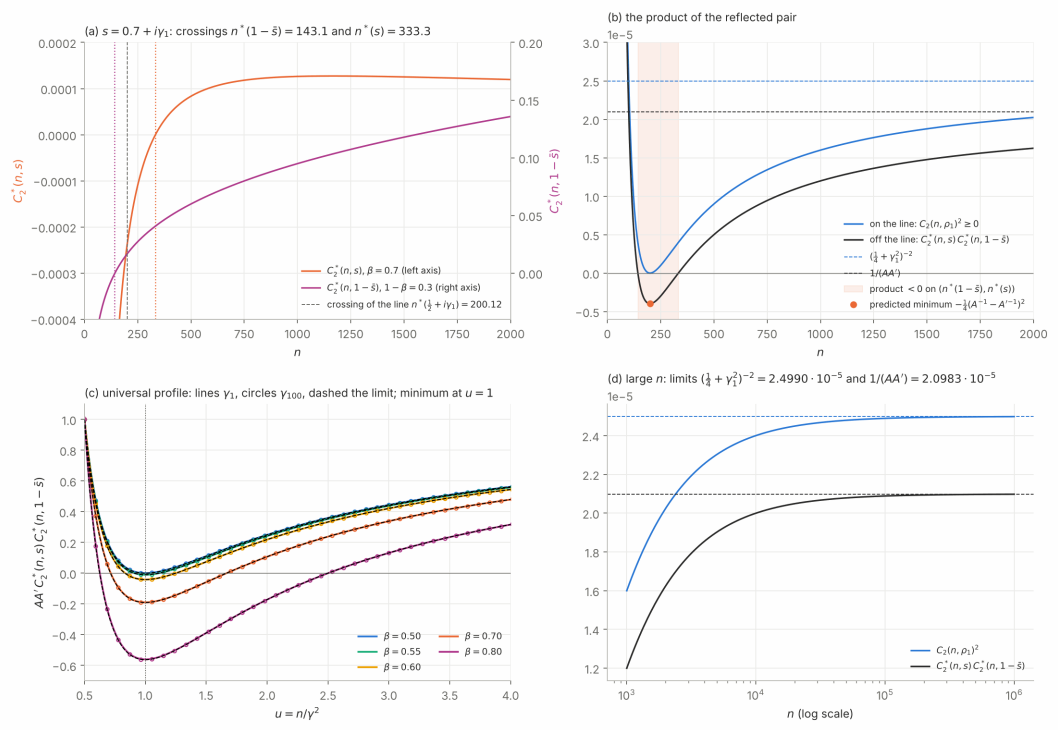

FIGURE 14. The reflected-pair product, wave removed, at a synthetic off-line point $s = 0.7 + i\gamma_1$: negative exactly between the two crossings, with depth governed by $(2\beta-1)^2$. (Figure from [CEn].)

The *functional reflection* of the cross-energy is the deepest tool of the series:

$$C_2^{\mathrm{FR}}(N,s) = G_N\big(s(1-s)\big) + \zeta(s)\,K_N(s), \tag{42}$$

with G_N explicit and K_N the carrier of the functional equation, whose zeros lie on the critical line [CEn, Prop. 2.2‴]. On the line it equals Hardy's function times $2\cos(t\log N - \vartheta(t))$, *computed from the integers below* N, and it sees every zero with its real part. Applied to any number system from its integers alone, it sees the off-line real zero of a Beurling system — which drifts with N, because the system has no functional equation — while the off-line zero of the Davenport–Heilbronn function is fixed [CEn, Prop. $2.2^{(6)}$] (Figure 15). Weil's criterion restricted to the cross-energy's own test functions $f * \mu_n$ is still equivalent to RH [CEn, Prop. $2.2^{(4)}$], its positivity residing in the fluctuation of the primes in the dual range $[n, n^2]$, which neither the density of the primes, nor Chebyshev's identity, nor random or Beurling models reproduce; a variant trades the dual range for the primes up to ne^{2a} and the divisor function [CEn, Prop. $2.2^{(5)}$]. The margin of the restricted form at height T is of order T^{-2} to T^{-4} of the Gamma term, and as $n \to \infty$ the restricted criterion becomes the positivity of the Hankel matrices of the power sums of $1/(\rho(1-\rho))$ — the sequence whose first term is $C_B = 2\widehat{B}$ — computable from ξ at $s = 1$ alone, and verified to size 43 [CEn, Prop. $2.2^{(8)}$].

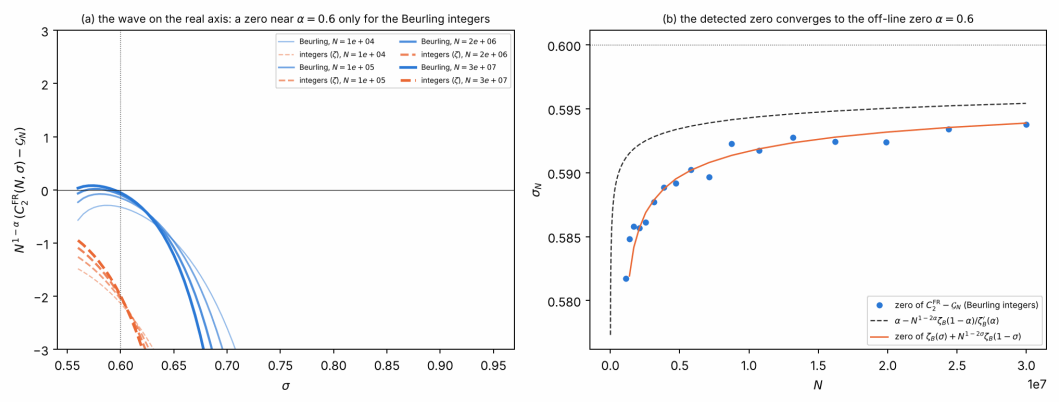

FIGURE 15. The functional reflection applied to a Beurling system from its integers alone: the off-line real zero is seen, and drifts with N because the system has no functional equation. (Figure from [CEn].)

18.3. **The energy law: a fourth mirror of the same constant.**

Theorem 18.3 (the energy law; [CEn, Thm. B, §3])**.** *Let $E(u)$ be the normalized error term of the prime number theorem. When all zeros are simple,*

$$\limsup_{U\to\infty} \frac{1}{U}\int_0^U E(u)^2\,du = \sum_{\rho} \lim_{n\to\infty} C_2(n,\rho) \qquad \text{in } [0,\infty];$$

in general both sides are finite or both are infinite, and RH holds if and only if they are finite. The exponential growth rate of the energy is $2\Theta - 1$. An unconditional local version expresses the energy in a Gaussian window as an absolutely convergent double sum over zeros; under RH it stays in a band of explicit width around C_B.

Corollary 18.4 ([CEn, Cor. 3.7])**.** *The mean energy is at least $C_B = 2\widehat{B}$ unconditionally, with equality if and only if RH holds and all zeros are simple.*

So the constant of Parts I and III has a fourth mirror: it is the minimal possible mean energy of the prime-counting error, attained exactly on RH with simple zeros — the primes up to 10^8 measure 0.0453 against 0.0462. Under RH, simple zeros and a spacing condition, the energy up to X is

$$C_B \log X + K + \Phi(\log X) + o(1), \qquad K = 2.5125\ldots,$$

with K explicit and Φ an almost periodic interference term from pairs of zeros, which the primes up to 10^7 follow with correlation 0.999 [CEn, Prop. 3.8] (Figure 16); the Gram matrix of the zeros has the C_2 limits on its diagonal and as its asymptotic spectrum, and random models of the primes and of μ (Cramér, Denjoy, Wintner) carry 20 to 150 times more energy [CEn, Rem. 3.9] — an echo, in the cross-energy metric, of the 46% gcd cancellation of Part II: arithmetic is far quieter than chance. Each zero contributes an orbit of the prime signal, a closed ellipse if $\beta = \frac{1}{2}$ and an expanding spiral if $\beta > \frac{1}{2}$, whose phase also determines β.

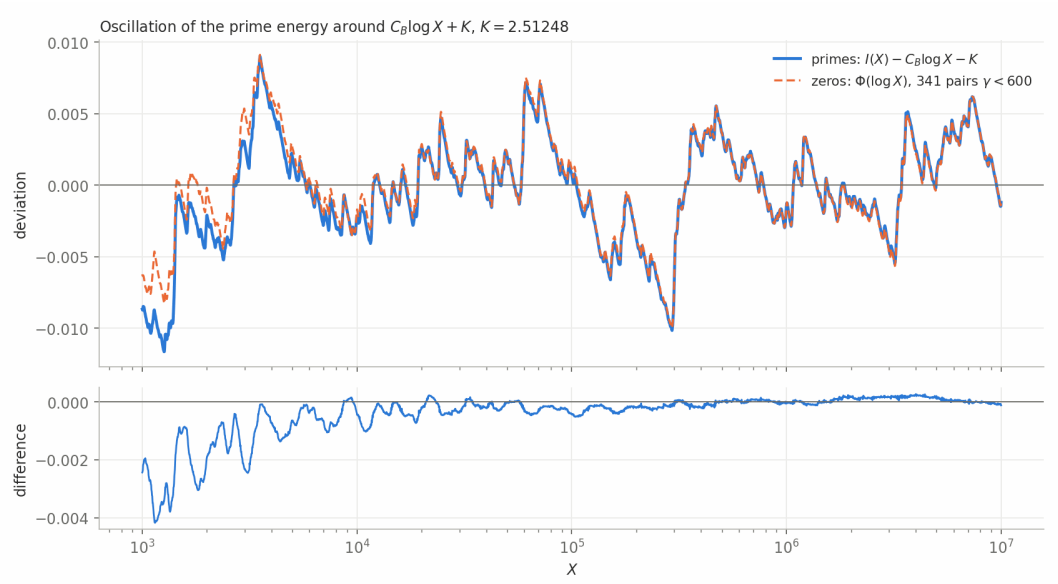

FIGURE 16. The second-order energy law $C_B \log X + K + \Phi(\log X)$, $K = 2.5125\ldots$: the primes up to 10^7 follow the interference term Φ with correlation 0.999. (Figure from [CEn].)

18.4. **The energy slope.** Let $S(\sigma, t) = \partial_\sigma \log|\xi(\sigma + it)|$, so that RH is equivalent to $S > 0$ for $\sigma > \frac{1}{2}$.

Theorem 18.5 ([CEn, Thms. C, D, E])**.** (C) *An off-line pair $\rho, 1 - \bar{\rho}$ contributes negatively to $S(\sigma, t)$ exactly inside the disc $(\sigma - \frac{1}{2})^2 + (t - \gamma)^2 < (\beta - \frac{1}{2})^2$.* (D) *For $\frac{1}{2} < \sigma < 1$, the set of $t \in [0, T]$ where $S(\sigma, t) \le 0$ has measure at most $2\sqrt{\sigma(1-\sigma)}\, N(\sigma, T+1)$: zero-density estimates make it sparse.* (E) *$S(\sigma, t)$ is Weil's functional evaluated at a Poisson kernel of width $\delta = \sigma - \frac{1}{2}$; truncating the kernel to a window of length L gives an exact finite expression in the prime powers $n \le e^L$, which under RH lies between $(1 - e^{-\delta L})^2$ and $(1 + e^{-\delta L})^2$ times $2S(\sigma, t)$. Consequently, under RH, primes up to P resolve S at distance δ from the line once $\delta \log P \gtrsim \log(1/\varepsilon)$; for the Davenport–Heilbronn function, an off-line zero at horizontal distance $\beta - \sigma$ from the probe becomes visible at $L \approx 4.5/(\beta - \sigma)$.*

Two companions of the slope: the argument of de la Vallée Poussin gives $S > 0$ in a region of width $c/\log t$ to the left of $\sigma = 1$ [CEn, Prop. 4.4]; and in the square $v = (\sigma - \frac{1}{2})^2$ of the distance to the line, *RH is equivalent to the concavity of* $\log|\xi|$ *in v at every height* [CEn, Prop. 4.3] — the structural conjecture of Appendix A, concavity of $\log W(\sqrt{u})$ for the Laguerre kernel, is a cognate statement one level up.

18.5. **The Euler product inside the strip.** The companion's Section 6 identifies each partial Euler product with the C_2 energy of the partial sums restricted to smooth integers, and proves that a corrected partial product converges to ζ in $\frac{1}{2} < \sigma < 1$ *exactly when RH holds*, decaying at a zero on the line like $|\zeta'(\rho)|e^{-\gamma_0}/\log x$. On the critical line the corrected partial products, read along $u = \log x$, are a superposition of helices, one per zero, whose mean energy at $\frac{1}{2} + it$ is the curvature $-(\log|\xi|)''(t)$ — the quantity of Appendix A again; computed with care, the helix built from the primes up to 10^8 fixes the real parts of known zeros to within about $4 \cdot 10^{-5}$, with symmetric off-line pairs entering only at second order (Figure 17). Sieving the wave $C_2 + iS_2$ by the primes and removing the pole correction turns its free amplitude into the Euler helix, so the partial sums and the partial products describe the same object; with a smooth weight in $\log x$ the identity becomes local — ζ is reproduced to 10^{-12} by the primes and the zeros within distance 20 of t, and $\zeta'(\rho)$ at every simple zero is given exactly, *without* the Deep Riemann Hypothesis [CEn, Prop. 6.1′]. On the line each prime acts on the wave by the real operator $1 - p^{-1}T_p$, the same at every height; the sieved functional equation shows why this rigidity does not combine with the reflection — it maps the half of the strip where the Euler product converges onto the half where it diverges. Three honesty results frame the section: the amplification that proves RH for curves becomes over $\mathbb{Q}$ a one-sided bound for $\psi(x)$, where, on short intervals near 10^8, Selberg's sieve loses a factor close to 2 exactly along Liouville's function — the parity barrier; the sums over zeros that see β are non-holomorphic and are reached from the primes only through energies, the damped energy having its first singularity at $2\Theta - 1$ and detecting an off-line zero at height T only with the primes up to about $T^{1/\delta}$; and *no combination of cosines without a constant term can be nonnegative at every* $\log p$ [CEn, Prop. 6.12]: positivity prime by prime requires the pole, and without it positivity can only be quadratic. Finally, the Jensen-moment route: mollifying by the Euler–Cáceres factor instead of a Dirichlet polynomial removes the pole barrier 0.0386; numerically, the resulting bound for the off-line excess decays like $1.828\,W(s_0)/\log X$ for ζ and grows with X for the Davenport–Heilbronn function; an attained off-line supremum of the real parts forces it to grow like a power of X, under RH it tends to 0, and RH is equivalent to its convergence to 0 — and, when the supremum of the real parts is attained, to the boundedness of a single mean square of ζ/E^* on the critical line [CEn, Prop. 6.1‴–Cor. $6.1^{(6)}$].

18.6. **The controls.** Every criterion of the series is tested against the Davenport–Heilbronn function f (same type of functional equation, zeros off the line) and its Euler-product twin $L(s, \chi_5)$ (same conductor and gamma factor): f violates each criterion in the way its off-line zeros predict — the energy grows at the rate its zeros dictate, the failure set of the pointwise inequality consists of one interval per off-line zero and fills 86–90% of Theorem 18.5(D)'s bound, the truncated Weil functional goes negative but outside the discs of (C), and a scan of Hardy's function to height 10^5 shows no extremum that fails to cross for ζ against eight such extrema for f, one per off-line pair — while $L(s, \chi_5)$ behaves like ζ in every test. The zeros of the minimiser of Weil's form [Co] approximate those of f as well as those of ζ until shortly after the form loses positivity at $L^* \approx 3.44$ (which high precision locates below the double-precision value 3.66), after which all approximations degrade at once and one zero of the minimiser locks onto the ordinate of the off-line pair. Weil's form with two windows detects an isolated zero $\beta = \frac{1}{2} + \delta$ at height γ with the prime powers up to $(C \log \gamma)^{1/\delta}$; the corresponding positivity holds for ζ up to $3 \cdot 10^6$ but fails for a random multiplicative model. And a Beurling system of Broucke–Debruyne–Révész type, built to $3 \cdot 10^7$, shares every zero of ζ, has a real zero at 0.6, and violates Chebyshev's identity at ρ_1 with remainder of order $N^{-0.4}$: *the regularity of the integers cannot replace the functional equation* [CEn, Prop. 6.3′]. These controls are the experimental counterpart of the no-go theorem of Part V: criteria blind to f, or satisfied by zero-sharing Beurling systems, cannot prove RH.

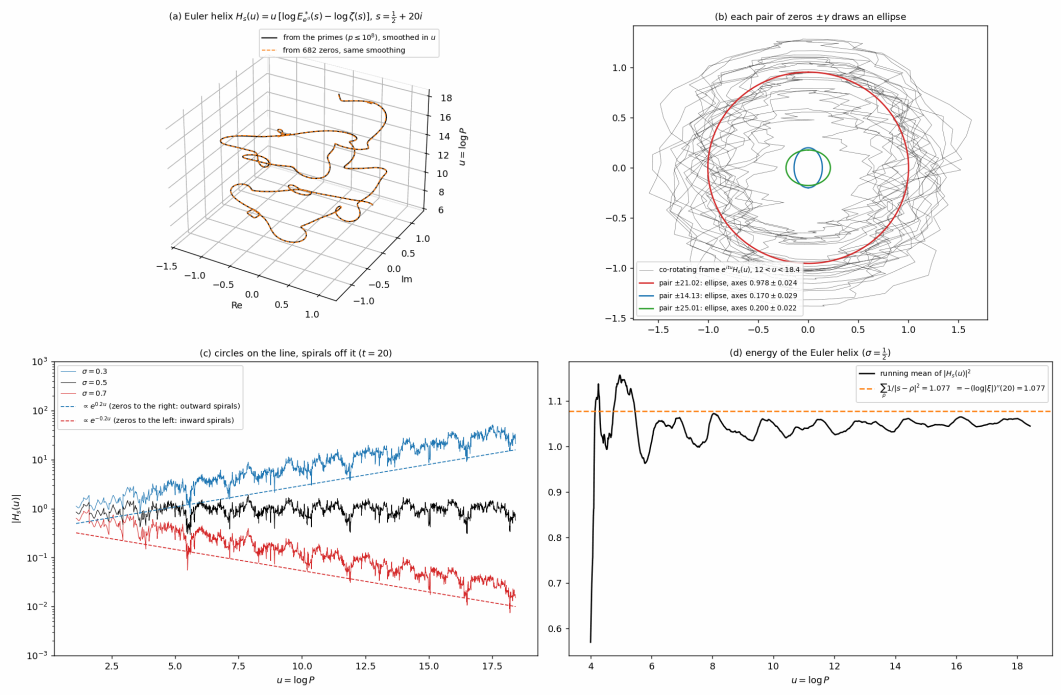

FIGURE 17. The Euler helix on the critical line at $s = \frac{1}{2} + 20i$: corrected partial products over the primes, one helix per zero, with the running mean of the helix energy against the curvature $-(\log|\xi|)''(20)$. (Figure from [CEn].)

Part V
The walls

The four parts above share a discipline: everything proved is an identity, a verified inequality, or an average. This part collects what that discipline cannot buy; the Davenport–Heilbronn and Beurling controls of Part IV are its experimental face. Section 19 proves it on the prime side; Section 20 reports the computational study of Weil positivity; Section 21 states the two walls on the Möbius side; Appendix A closes a tempting door on the ξ side.

19. What the arguments cannot see, and what they must use

19.1. **What the energy programme cannot see: a planted zero.** Remark 4.8 says that every route of Section 4.1 ends at the holomorphy of $\frac{\zeta'}{\zeta}$. This subsection makes that statement precise in the opposite direction. We build a second system of weights on the prime powers. It satisfies every structural ingredient the programme uses: positivity, the energy identity, the prime number theorem with the classical error term, the local recurrence, and Chebyshev's identity up to $O(x^{1/2+\delta})$. Yet its slope carries an oscillation of an off-line zero. So any proof of the energy bound must use the exact values of Λ at a precision that these ingredients do not give.

The construction. Fix $\delta \in (0, \frac{1}{2})$, $\gamma_0 > 0$, a switch-on point x_0 and a smooth function χ with $\chi = 0$ on $(-\infty, 0]$, $\chi = 1$ on $[1, \infty)$. Put

$$f(x) = \frac{2}{\gamma_0}\, x^{\delta} \sin(\gamma_0 \log x)\, \chi\Big(\log \frac{x}{x_0}\Big), \qquad \phi = f + \lambda h,$$

where h is a fixed smooth bump with compact support in $(1, x_0)$ and λ is a real constant chosen below. Define new weights on the prime powers,

$$w(q) = \frac{\Lambda(q)}{\sqrt{q}}\,\big(1 + \varepsilon(q)\big), \qquad \varepsilon(q) = \sqrt{q}\,\phi'(q), \tag{43}$$

and let A_w, $\alpha_w = A_w - 2\sqrt{x} + \frac{c_0}{2}$, C_w, $M_w = R - C_w$ and m_w be built from w exactly as A, α, C, M and m are built from $\Lambda(n)n^{-1/2}$.

Theorem 19.1 (a planted zero)**.** *With the notation above, for x_0 large enough:*

(i) positivity, and the weights are asymptotically the true ones: $|\varepsilon(q)| \le C_\phi\, q^{\delta - \frac{1}{2}} < 1$, *so* $w(q) > 0$ *and* $w(q)\sqrt{q}/\Lambda(q) \to 1$;
(ii) the energy identity: *Theorem 4.4 holds verbatim for m_w, with $\Lambda(q)q^{-1/2}$ replaced by $w(q)$;*
(iii) the planted oscillation: *unconditionally,* $\alpha_w(x) = \alpha(x) + f(x) + o(x^{\delta})$;
(iv) the prime number theorem is unchanged at the classical scale: $\sum_{q \le x} \Lambda(q)(1 + \varepsilon(q)) = x + O\big(x\, e^{-c\sqrt{\log x}}\big)$;
(v) Chebyshev's identity up to $x^{1/2+\delta}$: *for a suitable* λ, $\sum_{q \le x} \Lambda(q)(1 + \varepsilon(q))\lfloor x/q \rfloor = \log \lfloor x \rfloor ! + O(x^{1/2+\delta})$;
(vi) one of the two systems violates the energy bound: *for all large X,*

$$\max\Big(\int_1^X \alpha^2 \frac{dx}{x},\ \int_1^X \alpha_w^2 \frac{dx}{x} \Big) \ \ge\ \frac{X^{2\delta}}{7\,\delta\,\gamma_0^2}.$$

Proof. (i) $|f'(x)| \le \frac{2}{\gamma_0}(\delta + \gamma_0 + \|\chi'\|_\infty) x^{\delta - 1}$ and h' is bounded with compact support, so $|\varepsilon(q)| \le C_\phi q^{\delta - 1/2}$. On the support of h the bound is $|\lambda|\sqrt{q}\|h'\|_\infty$, which is < 1 for the small λ needed in (v). For $q \ge x_0$, $C_\phi x_0^{\delta - 1/2} < 1$ once x_0 is large.

(ii) The proof of Theorem 4.4 uses only that $m_w = r - 2A_w L + 2B_w$ is smooth with $m_w'' = r''$ between prime powers and that m_w' jumps by $-2w(q)$ at q. Nothing else about the weights enters.

(iii) $A_w(x)-A(x)=\sum_{q<x}\Lambda(q)\phi'(q)=\int_{2^-}^{x}\phi'(t)\,d\psi(t)$. Write $d\psi=dt+d(\psi(t)-t)$ and integrate the second part by parts:

$$A_w(x)-A(x)=\phi(x)+\phi'(x)\big(\psi(x)-x\big)-\int_2^x\big(\psi(t)-t\big)\phi''(t)\,dt.$$

Here $\phi(2)=0$, and for $x\ge x_0$, $\phi=f$. With $\psi(t)-t\ll te^{-c\sqrt{\log t}}$ and $\phi''(t)\ll\gamma_0 t^{\delta-2}$, the last two terms are $\ll x^{\delta}e^{-c'\sqrt{\log x}}=o(x^{\delta})$.

(iv) The same computation without the factor $q^{-1/2}$ gives $\sum_{q\le x}\Lambda(q)\varepsilon(q)=\int_2^x\sqrt{t}\,\phi'(t)\,d\psi(t)\ll x^{1/2+\delta}$, which is smaller than the classical error term.

(v) Write $\lfloor x/q\rfloor=x/q-\{x/q\}$. The defect is $D(x)=x\sum_{q\le x}\Lambda(q)\varepsilon(q)/q-\sum_{q\le x}\Lambda(q)\varepsilon(q)\{x/q\}$. The second sum is $\ll\sum_{q\le x}\Lambda(q)q^{\delta-1/2}\ll x^{1/2+\delta}$. The series $S=\sum_q\Lambda(q)\varepsilon(q)/q$ converges absolutely, and its tail beyond x is $\ll x^{\delta-1/2}$. The bump changes S by $\lambda\sum_q\Lambda(q)h'(q)q^{-1/2}\neq 0$ for a suitable h, so one value of λ gives $S=0$, and then $D(x)\ll x^{1/2+\delta}$.

(vi) By (iii), $f=\alpha_w-\alpha+o(x^{\delta})$, so $\int_1^X f^2\frac{dx}{x}\le 3\int\alpha_w^2\frac{dx}{x}+3\int\alpha^2\frac{dx}{x}+o(X^{2\delta})$. Since $\sin^2$ has mean $\frac12$, $\int_1^X f^2\frac{dx}{x}\sim\frac{4}{\gamma_0^2}\int^{\log X}e^{2\delta u}\sin^2(\gamma_0 u)\,du\sim\frac{X^{2\delta}}{\delta\gamma_0^2}$. □

What the theorem says. Every structural fact the programme uses holds for both systems: positivity of the weights, the energy identity, the recurrence of the minima (Proposition 3.6 is also pure bookkeeping), the prime number theorem with its classical error, and the arithmetic identity $\Lambda*1=\log$ up to $O(x^{1/2+\delta})$. Yet at most one of the two systems satisfies the energy bound of Proposition 4.2. So a proof of RH along these lines must use the exact values of Λ at relative precision $x^{\delta-1/2}$, for every $\delta>0$. That is the strength of the statement $\psi(x)=x+O(x^{1/2+\delta})$ itself, the zero-free strip of width $\frac12-\delta$.

Beurling systems give a stronger form of the same obstruction. Their weights come from a genuine Euler product, the integers satisfy $N(x)=cx+O(x^{\theta})$ with $\theta<1$, and yet the error term of the prime number theorem can be as large as $x\exp(-c\sqrt{\log x})$ [DMV]. The weights (43) do not come from an Euler product. Their point is that they stay within $q^{\delta-1/2}$ of the true ones and violate $\Lambda*1=\log$ only at the scale $x^{1/2+\delta}$. The information that rules them out is the exact identity $\sum_{n\le x}\Lambda(n)\lfloor x/n\rfloor=\log\lfloor x\rfloor!$, equivalently $\zeta(s)=\sum n^{-s}$ with its analytic continuation. No step of Sections 3 and 4 uses it beyond the prime number theorem.

The numbers. Figure 4(b) plants two zeros into the true prime powers below 2^{26}. The bump $h=\sin^2$ in $\log x$ on $[30,300]$ fixes $S=0$ on the computed range. In (A), $\delta=0.15$, $\gamma_0=10$, $x_0=10^3$: $\max|\varepsilon(q)|=0.112$, the Chebyshev defect satisfies $|D(x)|\le 0.31\,x^{0.65}$ for $10^6\le x\le 6\cdot10^7$, and the block energy grows like $2^{2\delta k}$. In (B), $\delta=0.25$, $\gamma_0=1000$, $x_0=10^5$: $\max|\varepsilon(q)|=0.077$, and the block energies agree with those of ζ within 0.002 up to $k=16$. Case (B) is the realistic one; a true off-line zero has $\gamma_0>T_0$ and is invisible below e^{54}.

Remark 19.2 (status). Theorem 19.1 is proved and unconditional. It does not say that RH is false, or that the energy programme is useless. It says exactly which input is missing: an argument that uses $\Lambda=\mu*\log$ exactly, not only through the prime number theorem. The same test applies to any other argument in this paper: an argument that remains valid for the weights (43) cannot prove RH.

19.2. **Using $\Lambda=\mu*\log$ exactly.** Theorem 19.1 says that a proof of the energy bound must use the identity $\Lambda=\mu*\log$ exactly, not only through the prime number theorem. This subsection carries that identity through the slope α and asks what it gives.

The identity. For every $n\ge1$,

$$\Lambda(n)=\sum_{d\mid n}\mu(d)\log\frac{n}{d},\qquad\text{equivalently}\qquad\sum_{d\mid n}\Lambda(d)=\log n,\qquad -\frac{\zeta'}{\zeta}(s)=\frac{1}{\zeta(s)}\cdot\big(-\zeta'(s)\big).\tag{44}$$

The divisor sum is a massive cancellation (Figure 18(a)). For n with two distinct prime factors, the positive and negative terms are each as large as $\log n$, and they cancel exactly to 0. Summed up to x, (44) is Chebyshev's identity $\sum_{m\le x}\psi(x/m)=\log\lfloor x\rfloor!$. Its Möbius inverse $\psi(x)=\sum_{d\le x}\mu(d)\log\lfloor x/d\rfloor!$ adds terms of size $x\log x$ that cancel to x (Figure 18(b): at $x=10^4$ the first term is 82 109 and the sum is 10 013.4). The cleanest exact form for what follows is

$$\psi(x)+\log x=\sum_{d\le x}\mu(d)\left\lfloor\frac{x}{d}\right\rfloor\log\frac{x}{d}, \tag{45}$$

which follows from (44) by summing $\sum_{d|n}\mu(d)\log(x/d)=[n=1]\log x+\Lambda(n)$ over $n\le x$. It holds to 10^{-11} in floating point at $x=10^4$, 12 345.6 and 99 999.5.

The slope through Möbius. Write, for $x>1$ not an integer,

$$m_1(x)=\sum_{d<x}\frac{\mu(d)}{d},\quad m_2(x)=\sum_{d<x}\frac{\mu(d)\log d}{d},\quad m_{1/2}(x)=\sum_{d<x}\frac{\mu(d)}{\sqrt{d}},\quad S(y)=\sum_{m<y}\frac{\log m}{\sqrt{m}},$$

and $S_0(y)=2\sqrt{y}\log y-4\sqrt{y}-\zeta'(\frac12)$. Euler–Maclaurin gives $S(y)-S_0(y)\to0$; the constant is $-\zeta'(\frac12)$ because $\sum_{m\le y}m^{-s}\log m=\frac{y^{1-s}\log y}{1-s}-\frac{y^{1-s}}{(1-s)^2}-\zeta'(s)+o(1)$ for $0<s<1$.

Proposition 19.3 (the slope through Möbius)**.** *For every $x>1$ that is not an integer, exactly,*

$$\begin{aligned}\alpha(x)&=\mathfrak{T}_1(x)+\mathfrak{T}_2(x)+\mathfrak{T}_3(x),\\ \mathfrak{T}_1&=2\sqrt{x}\,\big(m_1\log x-m_2-2m_1-1\big),\qquad \mathfrak{T}_2=-\zeta'(\tfrac12)\Big(m_{1/2}-\frac{1}{\zeta(\frac12)}\Big),\\ \mathfrak{T}_3&=\sum_{d<x}\frac{\mu(d)}{\sqrt{d}}\,(S-S_0)\Big(\frac{x}{d}\Big).\end{aligned}$$

Moreover $(S-S_0)(y)=-B_1(\{y\})\frac{\log y}{\sqrt{y}}+O\big(\frac{\log ey}{y^{3/2}}\big)$ with $B_1(t)=t-\frac12$, so that

$$\mathfrak{T}_3(x)=-\frac{1}{\sqrt{x}}\sum_{d<x}\mu(d)\,B_1\Big(\Big\{\frac{x}{d}\Big\}\Big)\log\frac{x}{d}+\sum_{d<x}\frac{\mu(d)}{\sqrt{d}}\,O\Big(\frac{\log(ex/d)}{(x/d)^{3/2}}\Big). \tag{46}$$

Proof. By (44), $A(x)=\sum_{dm<x}\mu(d)\log m\,(dm)^{-1/2}=\sum_{d<x}\mu(d)d^{-1/2}S(x/d)$. Write $S=S_0+(S-S_0)$. The S_0 part is

$$\sum_{d<x}\frac{\mu(d)}{\sqrt{d}}\Big[2\sqrt{\tfrac{x}{d}}\big(\log x-\log d\big)-4\sqrt{\tfrac{x}{d}}-\zeta'(\tfrac12)\Big]=2\sqrt{x}\,(m_1\log x-m_2-2m_1)-\zeta'(\tfrac12)\,m_{1/2}.$$

Subtracting $2\sqrt{x}$ and adding $\frac{c_0}{2}=\frac{\zeta'}{\zeta}(\frac12)$ gives $\alpha=\mathfrak{T}_1+\mathfrak{T}_2+\mathfrak{T}_3$. For (46), apply Euler–Maclaurin to $f(t)=t^{-1/2}\log t$ with the first Bernoulli function: $S(y)=\int_1^y f+C-B_1(\{y\})f(y)-\int_y^\infty B_1(\{t\})f'(t)dt$. The last integral is $O(|f'(y)|)$ after one more integration by parts, and C is fixed by $S-S_0\to0$. □

Status: proved; the identity is checked to $4.4\cdot10^{-9}$ at 400 points in $[10^3,10^7]$ (Figure 18(c)). The three pieces behave very differently. $\mathfrak{T}_1$ and $\mathfrak{T}_2$ are each large: their root mean squares on $[10^3,10^7]$ are 0.694 and 0.680, against 0.218 for α. But they are almost exact negatives of each other: their correlation is -0.9999999, and $\mathfrak{T}_1+\mathfrak{T}_2$ has root mean square 0.0140. So α is essentially $\mathfrak{T}_3$, and $\mathfrak{T}_3$ is essentially the *Möbius-weighted sawtooth* in (46): the difference has root mean square 0.012. This difference in turn nearly cancels $\mathfrak{T}_1+\mathfrak{T}_2$: α minus the sawtooth has root mean square 0.0020, against 0.249 for α at the same 50 points. To within 1%,

$$\alpha(x)\approx-\frac{1}{\sqrt{x}}\sum_{d<x}\mu(d)\Big(\Big\{\frac{x}{d}\Big\}-\frac12\Big)\log\frac{x}{d}.$$

The slope of the margin is the fractional parts of x/d, weighted by $\mu(d)\log(x/d)$. This is the one place where the additive structure of $\mathbb{N}$, through $\lfloor\cdot\rfloor$, enters the margin. The 1% agreement is not a coincidence. By Remark 19.6, the sawtooth sum is $-(\psi(x)-x)$ up to Mertens-type sums. Under RH, $(\psi(x)-x)/\sqrt{x}$ has the coefficients $-1/\rho = \frac{i}{\gamma} + O(\gamma^{-2})$, the same as α up to $O(\gamma^{-2})$.

The spectral content of the pieces. Under RH the coefficient of $e^{i\gamma u}$ in $\alpha(e^u)$ is $\frac{im_\gamma}{\gamma}$ (Proposition 4.1). It is rigid: it depends only on γ and the multiplicity. The pieces of Proposition 19.3 are not rigid.

Proposition 19.4 (the pieces carry $1/\zeta'(\rho)$)**.** *Assume RH, that the zeros are simple, and that the Perron expansions of m_1, m_2 and $m_{1/2}$ over the zeros converge in B^2 (as they do for $M(x)$ under Ng's hypothesis below). Then the coefficients of $e^{i\gamma u}$, γ the ordinate of $\rho = \frac12 + i\gamma$, are*

$$c_1(\gamma) = \frac{4(1-i\gamma)}{(i\gamma-\frac12)^2\,\zeta'(\rho)}, \qquad c_2(\gamma) = -\frac{\zeta'(\frac12)}{i\gamma\,\zeta'(\rho)}, \qquad c_3(\gamma) = \frac{i}{\gamma} - c_1(\gamma) - c_2(\gamma),$$

for $\mathfrak{T}_1$, $\mathfrak{T}_2$ and $\mathfrak{T}_3$. In particular, as $\gamma\to\infty$,

$$c_1(\gamma) + c_2(\gamma) = \frac{\big(4+\zeta'(\frac12)\big)\,i}{\gamma\,\zeta'(\rho)}\big(1+O(\gamma^{-1})\big), \qquad 4+\zeta'(\tfrac12) = 0.0774.$$

Proof. Perron's formula for $m_{1/2}$ is $\frac{1}{2\pi i}\int \frac{x^w}{w\,\zeta(\frac12+w)}dw$. The zero ρ gives the residue $\frac{x^{i\gamma}}{i\gamma\zeta'(\rho)}$, so $\mathfrak{T}_2$ has c_2. For m_1 the residue is $\frac{x^{w_0}}{w_0\zeta'(\rho)}$, $w_0=\rho-1$. For m_2, the Dirichlet series is $-(1/\zeta)'(1+w)$, which has a double pole with principal part $\frac{1}{\zeta'(\rho)(w-w_0)^2}$; the residue of $\frac{x^w}{w}\cdot\frac{1}{\zeta'(\rho)(w-w_0)^2}$ is $\frac{x^{w_0}}{\zeta'(\rho)}\big(\frac{\log x}{w_0}-\frac{1}{w_0^2}\big)$. Hence $m_1\log x - m_2 - 2m_1$ contributes $\frac{x^{w_0}}{\zeta'(\rho)}\big(\frac{1}{w_0^2}-\frac{2}{w_0}\big)$. Multiplying by $2\sqrt{x}$ gives c_1, since $1-2w_0 = 2-2i\gamma$. The sum of the three coefficients is the coefficient of α, which gives c_3. The asymptotic form follows from $c_1 = \frac{4i}{\gamma\zeta'(\rho)}(1+O(\gamma^{-1}))$. □

Status: the residues are a computation; the convergence in B^2 is an assumption of the same type as Ng's. A least-squares fit of the first fourteen frequencies to the computed pieces on $\log x \in [\log 10^3, \log 10^7]$ gives amplitudes $2|c|$ that agree with the predictions to the accuracy of the fit. At γ_1 the predicted and fitted values are 0.6998 and 0.7058 for $\mathfrak{T}_2$, 0.7145 and 0.7206 for $\mathfrak{T}_1$, and 0.0147 and 0.0148 for the residual $\mathfrak{T}_1+\mathfrak{T}_2$, which is 50 times smaller than either piece.

What the factorization does to the energy. The identity (44) factors $-\frac{\zeta'}{\zeta}$ into $\frac{1}{\zeta}$ times $-\zeta'$. The second factor is harmless: it is holomorphic in $\Re s > \frac12$ except at $s=1$, and of polynomial growth there. The first factor carries all the zeros, with residues $\frac{1}{\zeta'(\rho)}$. The rigidity of α is the identity $\zeta'(\rho)\cdot\frac{1}{\zeta'(\rho)} = 1$, and it reappears only when the three pieces are added. The energy shows the difference.

Proposition 19.5 (the Möbius energy)**.** *Let $M(x) = \sum_{n\le x}\mu(n)$.*

(i) *Unconditionally: if $\int_1^X M(x)^2x^{-2}dx = X^{o(1)}$, then RH holds.*

(ii) *(Ng [Ng]) If RH holds and $J_{-1}(T) = \sum_{0<\gamma\le T}|\zeta'(\rho)|^{-2} \ll T$, then $e^{-u/2}M(e^u)$ has a limiting distribution with mean 0 and mean square*

$$\big\langle (M(x)/\sqrt{x})^2\big\rangle = 2\sum_{\gamma>0}\frac{1}{|\rho\,\zeta'(\rho)|^2}.$$

(iii) *Numerically, the sum over the first 96 zeros is 0.0282. The Gonek–Hejhal conjecture $J_{-1}(T)\sim\frac{6T}{\pi^3}$ [Go] estimates the tail beyond $\gamma = 229.3$ at 0.0017, for a total of 0.0299. The block energies of $M(x)/\sqrt{x}$ for $k = 10,\dots,25$ have mean 0.0296 (Figure 18(d)).*

FIGURE 18. $\Lambda = \mu * \log$. (a) The divisor sum (44) for $n \le 60$: positive and negative terms, and their sum $\Lambda(n)$. (b) Partial sums of $\psi(x) = \sum_{d \le x} \mu(d) \log\lfloor x/d \rfloor!$ at $x = 10^4$. (c) The exact decomposition of Proposition 19.3: $\mathfrak{T}_1$ and $\mathfrak{T}_2$ are mirror images, and α follows the Möbius-weighted sawtooth (46) (dots). (d) Energy per dyadic block of α (to 10^{10}) and of the Mertens function $M(x)/\sqrt{x}$ (to 2^{26}), with their predicted means $(\log \xi)''(\frac{1}{2})$ and $2\sum_{\gamma>0} |\rho\zeta'(\rho)|^{-2}$.

Proof. (i) For $\Re s > 1$, $\int_1^\infty M(x) x^{-s-1} dx = \frac{1}{s\zeta(s)}$. In the variable $u = \log x$, $\int_0^\infty M(e^u) e^{-u/2} e^{-wu} du = \frac{1}{(w+\frac{1}{2})\zeta(w+\frac{1}{2})}$. If the energy is $e^{o(U)}$, the Cauchy–Schwarz argument of Proposition 4.2 makes this Laplace integral converge absolutely in $\Re w > 0$. So $\frac{1}{\zeta}$ is holomorphic in $\Re s > \frac{1}{2}$. (ii) is proved in [Ng], and (iii) is a computation. □

The comparison is the point of this subsection. Under RH, the energy of α is $2\sum m_\gamma^2/\gamma^2$. It is finite, equal to $(\log \xi)''(\frac{1}{2})$ for simple zeros, and needs nothing more. The Möbius energy needs control of $\zeta'(\rho)$ from below, which is not known even under RH. Through (44), the energy problem for α becomes an energy problem for M together with a harmless smooth factor. That problem is no easier: it is strictly less rigid.

Three levels of exactness. Which arithmetic facts exclude the systems of Theorem 19.1?

(0) *The planted weights* (43) have an Euler product only formally. Defining $\zeta_w(s) = \exp\big(\sum_n \Lambda(n)(1+\varepsilon(n)) n^{-s}/\log n\big)$, the local factor at p is $\exp\big(\sum_k (1+\varepsilon(p^k)) p^{-ks}/k\big)$. So $\mu_w(p) = -(1+\varepsilon(p))$ and $\mu_w(p^2) = \frac{1}{2}\big((1+\varepsilon(p))^2 - (1+\varepsilon(p^2))\big)$. For the system (A) of Figure 4(b), $\max|\mu_w(p)| = 1.112$ on $10^3 < p < 2^{13}$, where 55% of the primes have

$|\mu_w(p)| > 1$, and $\max|\mu_w(p^2)| = 0.113$. The local facts $\mu(p) = -1$ and $\mu(p^2) = 0$ already exclude them.

(1) *Beurling systems* have $\mu_B(p) = -1$ and $\mu_B(p^k) = 0$ for $k \geq 2$, because they come from a genuine Euler product over generalized primes. Diamond, Montgomery and Vorhauer [DMV] construct such systems with integer counting function $N_B(x) = \kappa x + O(x^\theta)$, $\frac{1}{2} < \theta < 1$, whose primes satisfy $\pi_B(x) = \mathrm{li}(x) + \Omega\big(x\exp(-c\sqrt{\log x})\big)$. So the multiplicative structure, even with integers known to precision x^θ for any $\theta > \frac{1}{2}$, does not give the energy bound. In Proposition 19.3 for such a system, the sawtooth $B_1(\{x/d\})$ is replaced by the error of N_B at x/d, which can be as large as $(x/d)^\theta$ instead of $O(1)$.

(2) *The integers* $\mathbb{N}$: $N(x) = \lfloor x \rfloor$, with error $O(1)$. This is the only information left, and Proposition 19.3 shows where it enters: the sawtooth in (46) and the two constants $\zeta(\frac{1}{2})$, $\zeta'(\frac{1}{2})$.

Remark 19.6 (status: what an exact argument must do). By Proposition 19.3 and Theorem 19.1, a proof along the lines of this paper must show that the Möbius-weighted sawtooth

$$\mathcal{S}(x) = \sum_{d \leq x} \mu(d)\, B_1\Big(\Big\{\frac{x}{d}\Big\}\Big) \log\frac{x}{d}$$

is $x^{1/2+o(1)}$ in mean square on $[X, 2X]$, and it must use the cancellation between $\mu(d)$ and the fractional parts of x/d. The identity (45) shows that nothing is gained for free. Inserting $\lfloor y \rfloor = y - \frac{1}{2} - B_1(\{y\})$ gives, exactly,

$$\mathcal{S}(x) = x\sum_{d \leq x} \frac{\mu(d)}{d}\log\frac{x}{d} - \frac{1}{2}\sum_{d \leq x}\mu(d)\log\frac{x}{d} - \psi(x) - \log x,$$

so $\mathcal{S}$ is $-\big(\psi(x) - x\big)$ up to Mertens-type sums that are $O(x^{1/2+\varepsilon})$ precisely when RH holds. The identity $\Lambda = \mu * \log$ is reversible, and rewriting ψ through it changes the form of the problem but not its content.

The form is still informative. $\mathcal{S}$ is a type I sum: Möbius against the explicit, bounded, additive function $B_1(\{x/d\})$. Expanding B_1 in its Fourier series turns $\mathcal{S}$ into exponential sums $\sum_d \mu(d)e(hx/d)$. The best unconditional bounds for $\psi(x) - x$, and so for $\mathcal{S}$, save only a factor $\exp\big(-c(\log x)^{3/5}(\log\log x)^{-1/5}\big)$ over the trivial bound. That is the Vinogradov–Korobov zero-free region, seen from the margin. A saving of $x^{-1/2+\varepsilon}$ would be RH. We do not claim any progress on closing this gap.

20. Windows and Weil positivity: a summary

The margin is Weil's form Q for one test function. The companion papers [CE20, CE21, CE22] study Q on whole spaces of test functions. We summarize the results, with their status, and refer to those papers for proofs and data.

Dirichlet polynomials [CE20]. Because $\Lambda \geq 0$, the prime part of Q on Gaussian-smoothed Dirichlet polynomials is unbounded below. So positivity is a cancellation between the archimedean place and the coherent sum of all primes, not a diagonal dominance (proved, elementary). With ball arithmetic, Q is positive definite on these polynomials for every length $N \leq 179$ (computer-assisted, for smoothing width 0.005). No translation-invariant square $\|Xk\|^2$ with X a Fourier multiplier fits below Q on the whole line; in particular, no Euler-product or mollifier multiplier does (proved).

The margin as $W - E$ [CE21]. Following Weil's proof for curves, write $Q = W - E$ with the pole $W = 2\widehat{g}(\frac{i}{2})\widehat{g}(-\frac{i}{2})$. Then W is a hyperbolic plane, like the two fibres of $C \times C$, and $Q \geq 0$ on a space implies that E has at most one positive direction there (proved; not conversely). The archimedean place alone has $(t_0/\pi + o(1))\log N$ positive directions, $t_0 = 6.2898$ (proved, of Szegő type). Numerically, the primes remove them only all together and only with their powers.

Windows [CE22]. In a window of length $L = \log N$ the form is one symbol. A positive completion of that symbol gives an exact sum of squares, and a coercive form has the completion $\pi/|E_L|^2$ with E_L the de Branges function of the window (proved, classical in substance). For even k of bounded variation supported in $[-\frac{L}{2}, \frac{L}{2}]$, with $K = \widehat{k}$, the explicit formula applied to $k * \tilde{k}$ gives

$$Q(k) = \sum_{\rho} K(\gamma_\rho)^2,$$

so under RH the least eigenvalue of a window is the minimum of $2\sum_{\gamma>0} K(\gamma)^2/\|k\|^2$ (proved). Connes observed that the minimizer reproduces the zeros [Co]. In high precision, it does so to 10^{-2} up to 1.3 to 1.9 times $T^* = 2\pi N$, the height where the density $\frac{1}{2\pi}\log\frac{t}{2\pi}$ of the zeros equals the Nyquist density $L/2\pi$ of the window (numerical; Figure 19). For the Davenport–Heilbronn function, the agreement stops at the corresponding scale $2\pi N/5$. The zeros below T^* carry a fraction between 10^{-25} and 10^{-8} of the least eigenvalue.

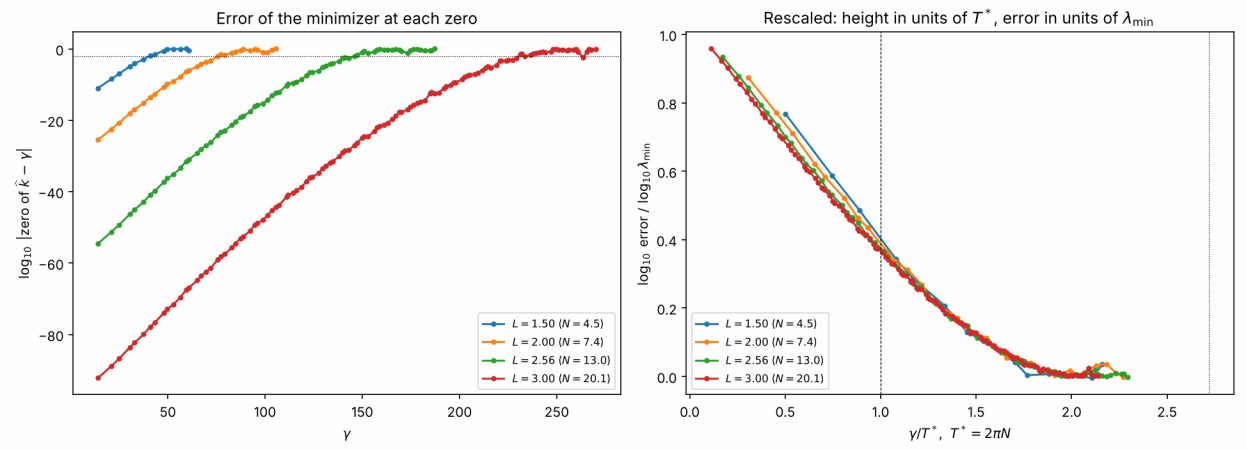

FIGURE 19. The minimizer of Weil's form in the windows $N = 4.5, 7.4, 13, 20.1$ [CE22]. Left: $\log_{10}$ of the distance from each zero of ζ to the nearest zero of the minimizer's transform. Right: the same data with the height in units of $T^* = 2\pi N$ and the error in units of $\lambda_{\min}$; dashed T^*, dotted eT^*.

21. The wall, stated precisely

Everything proved in Part II — the bounded smooth part, the variance identity and its constants, the exact intra-orbit fourth moment, the decoupling of the conductor-1 orbit, the prime-arrival invariance of the cross term, and (modulo Conjecture 12.19) the full Gaussian fourth moment — concerns averages over x with the truncation D fixed. The Riemann hypothesis (27) is a different kind of statement: a *pointwise* bound along the moving diagonal $D = x$, one value of $\widetilde{\Phi}$ per window, with no average left to take.

It is worth spelling out what a full Gaussian theorem would and would not give. If the tails of Section 11 were a theorem uniform in D — $P\big(|\widetilde{\Phi}_D| > \lambda \cdot 0.166\sqrt{D}\big) \le Ce^{-\lambda^2/2}$ — then $\lambda = \sqrt{2(1+\varepsilon)\log x}$ and a union over windows would give

$$|\Phi(x)| \le 0.24\sqrt{x\log x} \qquad \text{outside a sparse exceptional set,}$$

the Gaussian-heuristic version of RH, with a logarithmic loss and almost everywhere. Two gaps separate this from RH, and they are the whole distance. *Uniformity*: our moments average at fixed D; the diagonal couples the argument to the truncation. *The exceptional set*: no mean-value

control excludes isolated large values, and a zero off the critical line is precisely a measure-zero phenomenon in this metric — it would show in Φ as a coherent oscillation of size x^{β} visible only at the right scales, costing nothing in any fixed moment. Part I proves a no-go theorem (Section 19) on the other side of the same coin: weights indistinguishable from the primes' up to $q^{\delta-1/2}$ satisfy every such average statement and still carry an off-line zero, so no argument built from size estimates and fixed-order moments alone can close (27); the exact identities — (25) here, $\Lambda = \mu * \log$ there — must enter beyond their first moments. What Part II adds to that program is the exact second-moment structure (29), the identification of all the observed constants, and one new finite problem, Conjecture 12.19, that can be attacked without meeting the wall.

22. Discussion

Every statement in Part I that is equivalent to RH is equivalent through a Landau-type argument, and the same is true of the criteria of Parts II–IV, and each reduces to the holomorphy of $\frac{\zeta'}{\zeta}$, or of $\frac{1}{\zeta}$, to the right of the critical line. The unconditional results are of two kinds. Some are exact identities between the primes and the margin: the Fejér formula, the description of the minima, the energy identity and the Möbius representation. The others are statements of what cannot work: no uniform positive lower bound for M under RH (Proposition 3.3), and no argument that survives planted weights (Theorem 19.1). The second kind gives a test that any proposed argument can be run against: if it also works for the weights (43), for a Beurling system [DMV], or for the Davenport–Heilbronn function of the controls of Part IV, it cannot prove RH. What remains is an estimate that uses $N(x) = \lfloor x \rfloor$ exactly, through the sawtooth of Proposition 19.3. The best unconditional estimates of that kind save $\exp(-c(\log x)^{3/5}(\log\log x)^{-1/5})$, the Vinogradov–Korobov region. A saving of $x^{-1/2+\varepsilon}$ would be RH. We do not prove the Riemann hypothesis.

Remark 22.1 (Weil positivity for small support, and where this family sits). The margin evaluates the Weil functional on the one-parameter Fejér family: the autocorrelation underlying $M(N)$ has multiplicative support $[N^{-1}, N]$, so the primes enter only through the powers $q < N$. The regime in which Weil positivity is *proved* for all test functions is the archimedean one: Connes and Consani [CC] establish, through the prolate and Sonin operators, a positivity inequality for every g supported in $[2^{-1/2}, 2^{1/2}]$ (autocorrelation in $[\frac{1}{2}, 2]$), precisely the range in which no prime occurs; the semilocal programme [CCM] is to add the places $2, 3, \dots, p$ and the support (p^{-1}, p) one prime at a time. Specialized to our family, the proved range is $N \le 2$, where $M = R$ and positivity is the elementary estimate of Proposition 2.8(1); their theorem is of course much stronger there, being an operator bound valid for all g. The map between the two languages is Section 6: adding a place is admitting the next batch of prime powers into the walk, and the question both programmes face is the same uniformity – positivity must survive as the support, and with it every prime, enters. Theorem 19.1 constrains the operator route too: any semilocal positivity argument robust under the planted weights (43) cannot close RH, so the arithmetic identity $\Lambda = \mu * \log$ must enter the operator estimates at some stage. We do not see how to do this, and record the translation in the hope that the one-variable family, where every object is computable, is a useful test bench for that programme.

Appendix A. A cognate positivity problem: the Laguerre kernel of ξ and the Gaussian-mixture route

Write $F(t) = \xi(\frac{1}{2} + it) = \int_{\mathbb{R}} \Phi(x) e^{itx}\, dx$ with the classical theta kernel $\Phi > 0$, and let $L_1(t) = F'(t)^2 - F(t)F''(t)$ be the first Laguerre expression; RH implies $L_1 \ge 0$ on $\mathbb{R}$, and proving this unconditionally is an open problem of the Csordas circle [Cs]. One has the exact representation

$L_1 = \widehat{W}$ with the nonnegative, even kernel

$$W(v) = \tfrac{1}{2}\int_{\mathbb{R}} (2x - v)^2\, \Phi(x)\Phi(v - x)\, dx = (\Phi_2 * \Phi_0 - \Phi_1 * \Phi_1)(v), \qquad \Phi_k(x) = x^k \Phi(x),$$

so $L_1 \geq 0$ is equivalent, by Bochner, to W being positive definite. A natural sufficient condition would be a representation $W(v) = \int_0^\infty e^{-av^2}\, d\mu(a)$ with $\mu \geq 0$: positive mixtures of Gaussians are positive definite. *This route is closed.* If such a μ existed, $g(u) = W(\sqrt{u}) = \int e^{-au} d\mu(a)$ would be logarithmically convex in u, by Cauchy–Schwarz on μ. Computing W in high-precision arithmetic (Simpson quadrature of the doubly exponentially decaying integrand, step-halving agreement to 49 digits), the second differences of $\log W(\sqrt{u})$ on $u \in [0.4, 8.8]$ (step 0.2) are *strictly negative throughout*, between -0.034 and -0.093 — some 38 orders of magnitude above the numerical error. No positive Gaussian mixture represents W. The data suggest more (Figure 20): $\log W(\sqrt{u})$ appears to be *concave* on all of $u > 0$, with monotone second differences. We record this as a structural conjecture — $(\log W(\sqrt{u}))'' < 0$ for $u > 0$ — which rhymes with the doubly exponential decay $\log W(v) \asymp -2\pi e^{|v|}$ inherited from Φ, and leave the positive-definiteness of W where we found it: open, and not approachable by termwise positivity.

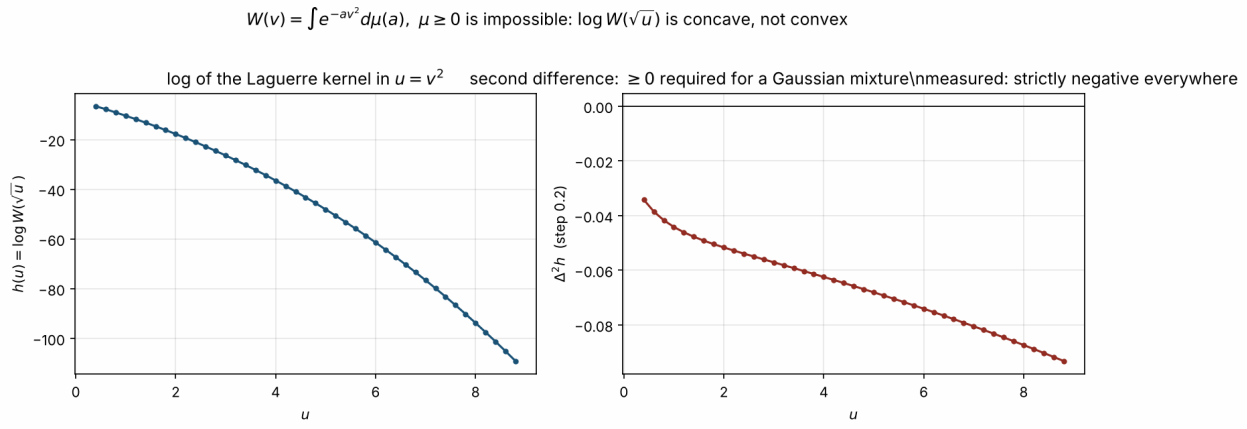

FIGURE 20. The Laguerre kernel of ξ in the variable $u = v^2$: $\log W(\sqrt{u})$ is concave, not convex, so W is not a positive mixture of Gaussians.

Data and code

Each part carries its own ancillary code, included with this paper: for Part I, the verification, certification, walk, mean and projection scripts of the Cesàro margin (including `verify_margin.py`, run by the author to 10^{12}, and `certify_zeros.py`, Arb ball arithmetic); for Part II, the exact-arithmetic suite (`xexact.py`, `decouple.py`, `sectors.py`, `ladder.py`, `xsectors.py`, `parity.py`, and the sampling and figure scripts); for Part III, `nbmain.py` and the Vasyunin–Gram data; for Part IV, the code and data travel with the companion monograph [CEn]; for the appendix, the high-precision kernel computation (`wkernel.py`, in the ancillary). All exact rational claims were reproduced independently in separate arithmetic during review.

Use of AI assistance

The research programme and the choice of objects are the author's. Claude (Anthropic) was used throughout as a research and writing aid: to derive and check lemmas, to write and run the computations, to referee drafts in independent sessions, and to draft text and figures. The author checked all results and is responsible for the content.

References

[CC] A. Connes, C. Consani, Weil positivity and trace formula, the archimedean place, *Selecta Math.* 27 (2021), 77; arXiv:2006.13771.

[CCM] A. Connes, C. Consani, H. Moscovici, Zeta zeros and prolate wave operators, arXiv:2310.18423, 2023.

[Bu] J. Büthe, An analytic method for bounding $\psi(x)$, *Math. Comp.* 87 (2018), 1991–2009.

[Arb] F. Johansson, Arb: efficient arbitrary-precision midpoint-radius interval arithmetic, *IEEE Trans. Comput.* 66 (2017), 1281–1292.

[ANS] A. Akbary, N. Ng, M. Shahabi, Limiting distributions of the classical error terms of prime number theory, *Q. J. Math.* 65 (2014), 743–780.

[Bes] A. S. Besicovitch, *Almost Periodic Functions*, Cambridge University Press, 1932.

[Bo] E. Bombieri, Remarks on Weil's quadratic functional in the theory of prime numbers I, *Rend. Mat. Acc. Lincei* (9) 11 (2000), 183–233.

[Co] A. Connes, The Riemann hypothesis: past, present and a letter through time, arXiv:2602.04022, 2026.

[DH] H. Davenport, H. Heilbronn, On the zeros of certain Dirichlet series I, II, *J. London Math. Soc.* 11 (1936), 181–185, 307–312.

[Go] S. M. Gonek, On negative moments of the Riemann zeta-function, *Mathematika* 36 (1989), 71–88.

[DMV] H. G. Diamond, H. L. Montgomery, U. M. A. Vorhauer, Beurling primes with large oscillation, *Math. Ann.* 334 (2006), 1–36.

[IK] H. Iwaniec, E. Kowalski, *Analytic Number Theory*, AMS Colloquium Publications 53, 2004.

[In] A. E. Ingham, *The Distribution of Prime Numbers*, Cambridge Tracts 30, Cambridge University Press, 1932.

[Jn] D. R. Johnston, Improving bounds on prime counting functions by partial verification of the Riemann hypothesis, *Ramanujan J.* 59 (2022), 1307–1321.

[RuSa] M. Rubinstein, P. Sarnak, Chebyshev's bias, *Experiment. Math.* 3 (1994), 173–197.

[MV] H. L. Montgomery, R. C. Vaughan, *Multiplicative Number Theory I. Classical Theory*, Cambridge University Press, 2007.

[KLN] H. Kadiri, A. Lumley, N. Ng, Explicit zero density for the Riemann zeta function, *J. Math. Anal. Appl.* 465 (2018), 22–46; arXiv:2101.12263.

[Ng] N. Ng, The distribution of the summatory function of the Möbius function, *Proc. London Math. Soc.* (3) 89 (2004), 361–389.

[PT] D. Platt, T. Trudgian, The Riemann hypothesis is true up to $3 \cdot 10^{12}$, *Bull. Lond. Math. Soc.* 53 (2021), 792–797.

[Su1] M. Suzuki, Aspects of the screw function corresponding to the Riemann zeta-function, arXiv:2206.03682, 2022 (v4, 2023).

[Su2] M. Suzuki, On variants of Chebyshev's conjecture, *Ramanujan J.* (2025), doi:10.1007/s11139-025-01238-9, arXiv:2411.07436.

[T] E. C. Titchmarsh, *The Theory of the Riemann Zeta-Function*, 2nd ed., Oxford University Press, 1986, Section 10.25.

[Tr] T. S. Trudgian, An improved upper bound for the argument of the Riemann zeta-function on the critical line II, *J. Number Theory* 134 (2014), 280–292.

[Vo] A. Voros, *Zeta Functions over Zeros of Zeta Functions*, Lecture Notes of the Unione Matematica Italiana 8, Springer, 2010.

[We52] A. Weil, Sur les "formules explicites" de la théorie des nombres premiers, *Comm. Sém. Math. Univ. Lund* (1952), tome supplémentaire, 252–265.

[Wi] D. V. Widder, *The Laplace Transform*, Princeton University Press, 1941.

[CPF] P. Caceres, The Caceres Prime Function: a square-root damped prime weight, its exact Prime Zeta structure, and a dual spectrum of the Riemann zeros, preprint, 2026.

[CE20] P. Caceres, Cross-energy of Dirichlet partial sums. Part 20: The Euler cross-energy, Weil positivity and a one-variable form of the Riemann hypothesis, preprint, 2026.

[CE21] P. Caceres, Cross-energy of Dirichlet partial sums. Part 21: The margin as $W - E$, a Hodge index for the integers, preprint, 2026.

[CE22] P. Caceres, Cross-energy of Dirichlet partial sums. Part 22: The minimizer of the Weil form in a window, and the scale $2\pi N$, preprint, 2026.

[Ca26] P. Caceres, The $6k \pm 1$ race: Chebyshev's bias, the margin of $L(s, \chi_{-3})$ and an exact weighted inequality, preprint, 2026.

[FL] J. Franel, E. Landau, Les suites de Farey et le problème des nombres premiers, *Gött. Nachr.* (1924), 198–206.

[BD] L. Báez-Duarte, A strengthening of the Nyman–Beurling criterion for the Riemann hypothesis, *Atti Accad. Naz. Lincei Rend. Lincei Mat. Appl.* 14 (2003), 5–11.

[BBLS] L. Báez-Duarte, M. Balazard, B. Landreau, E. Saias, Notes sur la fonction ζ de Riemann, 3, *Adv. Math.* 149 (2000), 130–144.

[BCF] S. Bettin, J. B. Conrey, D. W. Farmer, An optimal choice of Dirichlet polynomials for the Nyman–Beurling criterion, *Proc. Steklov Inst. Math.* 280 (2013), suppl. 2, S30–S36; arXiv:1211.5191.

[Be] A. Beurling, A closure problem related to the Riemann zeta-function, *Proc. Nat. Acad. Sci. USA* 41 (1955), 312–314.

[Burn] J.-F. Burnol, A lower bound in an approximation problem involving the zeros of the Riemann zeta function, *Adv. Math.* 170 (2002), 56–70.

[CEn] P. Caceres, Cross-energy of Dirichlet partial sums: the critical line as the trace of the pole, the energy of the prime-counting error, the energy slope as a Weil functional, and the functional reflection, consolidated monograph (v3), preprint, 2026.

[Cs] G. Csordas, Fourier transforms of positive definite kernels and the Riemann ξ-function, arXiv:1309.0055, 2013.

[Dav] H. Davenport, *Multiplicative Number Theory*, 3rd ed., Springer, 2000.

[Ny] B. Nyman, *On the One-Dimensional Translation Group and Semi-Group in Certain Function Spaces*, Thesis, Uppsala, 1950.

[Vas] V. I. Vasyunin, On a biorthogonal system associated with the Riemann hypothesis, *Algebra i Analiz* 7 (1995), no. 3, 118–135; English transl., *St. Petersburg Math. J.* 7 (1996), 405–419.

Universidad Europea de Valencia, Valencia, Spain
Email address: pedrojesus.caceres@universidadeuropea.es