

On the proportions of simple and of distinct zeros of the Riemann zeta function

Oliver D’Souza*

Release version, October 2026 (repository tag `v1.0`)

Abstract

We prove, without hypotheses on the zeros and with computer-assisted steps, that, as $T \rightarrow \infty$, at least 0.676102666 of the zeros ρ of $\zeta(s)$ with $T < \text{Im } \rho \leq 2T$, counted with multiplicity, are simple, and at least 0.838051333 of them are distinct. The largest previous unconditional bounds that we located in the refereed and arXiv literature are, truncated, 0.67250077 for simple zeros, which is the bound for simple zeros on the critical line, and 0.83625038 for distinct zeros. Unrefereed computer-assisted claims for simple zeros on the line at bandwidth one reach 0.67349239; the three largest of them use the fixed-test-function pair-correlation formula, as Lamzouri does, rather than the inputs of Alpöge and Furman used here. Two unrefereed works claim considerably more by other means, and we have not verified them. The argument extends the rank–trace method of Alpöge and Furman. With a smaller computer-certified local inequality ($K = 5$ consecutive zeros instead of 7) it gives the slightly smaller pair 0.675158622 and 0.837579311. For this pair the whole proof, the local certificate included, is checked by the kernel of Lean 4 with no hypotheses. For the $K = 7$ pair and for the two statistics below, everything except the local certificate (the analytic and algebraic steps, the majorant certificate and the closing arithmetic) is checked by the kernel, and the local certificate enters as a single displayed hypothesis, verified outside Lean by interval arithmetic. Two further statistics are recorded. At least 0.6733736895 of the zeros are simple and on the critical line, by a method that others published first; this is not a record, since for this statistic six unrefereed works announce larger values, from 0.673399 to 0.67349239, one of them (0.6734) as a hypothesis-free Lean theorem, whereas our formal statement carries a certificate hypothesis. At least 0.8879195 of the zeros are simple or on the line, which is Lamzouri’s formula evaluated at 0.6733736895. The results are lower bounds only and say nothing about the Riemann hypothesis itself. None of them has been independently refereed, and the Lean development has not been rebuilt outside this project.

Contents

1	Introduction	3
1.1	Results	3
1.2	What is new	4
1.3	What the results are not	5
1.4	Relation to other work	5
1.5	Formal verification	7
1.6	Plan	9

*© 2026 Oliver D’Souza. This paper is licensed under CC BY 4.0.

2	Setting and the inputs from Alpöge–Furman	9
2.1	The window, the test family and the matrix	9
2.2	The kernel and the window constant	10
2.3	The inputs	10
2.4	Normalised ordinates and the Gram identity	10
2.5	The limit $\lambda \rightarrow 1^-$	11
3	Linear algebra	12
3.1	The stability rank–trace inequality	12
3.2	Energy and the function Φ_m	13
3.3	Pinching and monotonicity	13
3.4	Pair removal	14
3.5	Localisation of the large eigenvalues	14
4	The local inequality and simple zeros on the line	15
4.1	The local inequality	15
4.2	From the zeros to the local inequality	16
4.3	Proof of Theorem 1.2	17
5	Simple and distinct zeros	17
5.1	The all-marks local inequality	18
5.2	The slack of an on-line set	18
5.3	Removing off-line pairs and heavy zeros	19
5.4	The on-line spectral condition	19
5.5	Proof of Theorem 1.1	21
6	Zeros that are simple or on the line	23
7	The certificates	24
7.1	What is certified	24
7.2	Interval branch-and-bound	25
7.3	Kernel replay inside Lean	26
7.4	How to reproduce	28
8	Remarks on the limits of the method	28
A	The Lean statements	30
A.1	The counting functions	30
A.2	The definitions of this paper	31
A.3	The fourteen statements of record	34
A.4	The library’s headline theorems	36
A.5	The chain for the $K = 5$ certificate	37
A.6	The replay’s theorems: the $K = 5$ pair without hypothesis	37
A.7	The majorant	38
A.8	Axiom census and sorry	39
A.9	Repository layout	39
B	How the work was done	41

1 Introduction

1.1 Results

Write a nontrivial zero of $\zeta(s)$ as $\rho = \beta + i\gamma$ and let $m_\rho \geq 1$ be its multiplicity. As in Alpöge and Furman [2], for $0 \leq T_1 < T_2$ let

$$\begin{aligned} N(T_1, T_2) &:= \sum_{T_1 < \gamma \leq T_2} m_\rho && \text{(zeros, counted with multiplicity),} \\ N^d(T_1, T_2) &:= \#\{\rho : T_1 < \gamma \leq T_2\} && \text{(distinct zeros),} \\ N^s(T_1, T_2) &:= \#\{\rho : T_1 < \gamma \leq T_2, m_\rho = 1\} && \text{(simple zeros),} \\ N_0^s(T_1, T_2) &:= \#\{\rho : T_1 < \gamma \leq T_2, \beta = \tfrac{1}{2}, m_\rho = 1\} && \text{(simple zeros on the critical line),} \end{aligned}$$

and let $N^{\text{sc}}(T_1, T_2)$ be the number of zeros with $T_1 < \gamma \leq T_2$ that lie on the critical line, counted with multiplicity, plus the number of those off the line that are simple. Thus $N_0^s \leq N^s \leq N^d \leq N$ and $N_0^s \leq N^{\text{sc}} \leq N$. All ratios below have $N(T, 2T)$, which counts multiplicity, in the denominator.

Theorem 1.1 (Simple and distinct zeros). *As $T \rightarrow \infty$,*

$$N^s(T, 2T) \geq (0.676102666 - o(1)) N(T, 2T), \quad N^d(T, 2T) \geq (0.838051333 - o(1)) N(T, 2T).$$

With the smaller certificate of Section 7 ($K = 5$) the same argument gives 0.675158622 and 0.837579311.

The constants are $\Sigma = (H + \frac{a_2}{2} - \nu)/(1 - a_1 + \frac{a_2}{2})$ and $D = \frac{1}{2}(1 + \Sigma)$, where $H = H(\cos 1.6s) = 0.671981551000\dots$ is the Alpöge–Furman constant of the window $\cos(1.6s)$ and (a_1, a_2, ν) are the rational claims of a certified local inequality (Theorem 5.8). For $K = 7$, $\Sigma = 0.67610266696413\dots$ and $D = 0.83805133348206\dots$; for $K = 5$, $\Sigma = 0.67515862219923\dots$ and $D = 0.83757931109961\dots$; the ninth digits above are truncated. In Lean 4 [8], the $K = 5$ pair is proved with no hypotheses and the $K = 7$ pair with its certificate as the one hypothesis (§1.5).

Theorem 1.2 (Simple zeros on the critical line). *As $T \rightarrow \infty$, $N_0^s(T, 2T) \geq (0.6733736895 - o(1)) N(T, 2T)$.*

Corollary 1.3 (Simple or on the line). *As $T \rightarrow \infty$, $N^{\text{sc}}(T, 2T) \geq (0.8879195 - o(1)) N(T, 2T)$.*

For distinct zeros Theorem 1.1 improves Wang’s bound [27], 0.83625038 truncated (§1.4), by $1.80 \cdot 10^{-3}$ ($1.32 \cdot 10^{-3}$ with $K = 5$). For simple zeros there is, to our knowledge, no earlier bound specific to the statistic: since a simple zero on the line is simple, every lower bound for N_0^s is one for N^s , and Theorem 1.1 exceeds the largest refereed or arXiv value that we located, 0.67250077 [27], by $3.60 \cdot 10^{-3}$, and the largest unrefered bandwidth-one claim for N_0^s that we know of (0.67349239, §1.4) by $2.61 \cdot 10^{-3}$; larger unrefered claims obtained by other methods are discussed there. Both values of Theorem 1.1 lie below what is known under the Riemann hypothesis: 0.6792 [7] and $19/27$ [5] for simple zeros, 0.8477 [7] for distinct zeros.

Theorem 1.1 is the result that this paper claims as new, as far as we have found in the public record: apart from the unrefered claims of Yang and Yang by other methods (§1.4), we found no earlier bound for N^s or N^d beyond the standard consequences of a bound for N_0^s (§1.2). We read 24 related works at source and made three web searches; this says only that we did not find such a bound, not that none exists. Theorem 1.2 is not a record: for simple zeros on the line, six unrefered works announce larger values, from 0.673399 to 0.67349239, and one of them, Zhaoyi-Tian’s 0.6734, is announced as a hypothesis-free Lean theorem, while the formal statement of Theorem 1.2 carries a certificate hypothesis (§1.4, §1.5).

Table 1: Lower bounds for $\liminf_{T \rightarrow \infty} X(T, 2T)/N(T, 2T)$. “Previous” lists refereed or arXiv values, and for N_0^s also the range of unrefered values; the unrefered values are discussed in §1.4. The last column is explained in §1.5. Lamzouri’s and Wang’s values are stated for $0 < \gamma \leq T$; a lower bound for every dyadic interval $(T, 2T]$ implies the same lower bound for $0 < \gamma \leq T$. Wang’s values are our truncations of the bounds $C_0 + \delta_0$ and $C_1 + \delta_0/2$ that it states (§1.4); Lamzouri’s 0.8876200 is its formula at the Montgomery–Taylor constant, evaluated by us (it states 88.76%).

X	previous	this paper	status of the proof in Lean
N^d (distinct)	$\frac{5}{6}$ and 0.8362503 [2]; 0.8362 [15]; 0.83625038 [27]	0.838051333 ($K = 7$) 0.837579311 ($K = 5$)	$K = 7$: certificate as hypothesis $K = 5$: no hypothesis (the certificate is replayed by the kernel)
N^s (simple)	0.6725007 [2, 15]; 0.67250077 [27]; all implied by $N_0^s \leq N^s$	0.676102666 ($K = 7$) 0.675158622 ($K = 5$)	as for N^d
N^{sc} (simple or on the line)	0.8876200 [15]	0.8879195 (Cor. 1.3)	certificate as hypothesis
N_0^s (simple, on the line)	0.6725007 [2, 15]; 0.67250077 [27]; unrefered: six values from 0.673399 to 0.67349239, one announced as a hypothesis-free Lean theorem (§1.4)	0.6733736895 (Thm 1.2; not a record)	certificate as hypothesis

1.2 What is new

Alpöge and Furman [2] restrict Weil’s explicit formula to a finite family of modulated copies of a window and obtain a real symmetric matrix G whose trace and Hilbert–Schmidt norm are known asymptotically from the prime side, and whose zero side splits into a positive semidefinite part P from the zeros on the line and a part Q from the off-line pairs with at most one positive eigenvalue per pair. A rank–trace inequality then gives $N_0^s \geq (2 - R(\psi) - o(1))N$ for every admissible window ψ .

The main results of this paper, Theorem 1.1, keep more of the matrix. For N^s and N^d a simple off-line pair counts twice in the numerator, and this is exactly what its negative eigenvalue costs in the rank–trace accounting (the “simple pair” entry in Step 3 of the proof of Theorem 5.8 is $2 - 2a_1 - c_* = 0$). Off-line pairs and zeros of multiplicity at least three can therefore be removed at no loss (Lemma 5.4), and what remains is an inequality about the zeros on the line of multiplicity one or two, the on-line spectral condition (Theorem 5.7). We prove it with four ingredients:

- (i) an exact identity for the slack of the rank–trace inequality on a set of on-line zeros (Lemma 5.2), which isolates the Gram energy among *all* of them, whatever their multiplicities;
- (ii) a local inequality over all patterns of multiplicities $\mathbf{m} \in \{1, 2\}^K$ of K consecutive zeros, certified by computer (Definition 5.1, Section 7), and its aggregation along the ordered zeros (Proposition 5.3);
- (iii) a localisation lemma for the eigenvalues above a threshold (Theorem 3.7), which reduces the spectral condition to a set of zeros with a close neighbour, where it is paid for by the zeros themselves (Lemma 5.6);

- (iv) a majorant of Beurling–Selberg type for the window, with Fourier transform supported in $[-1, 1]$, which bounds the largest eigenvalue on the remaining separated set (Lemma 5.5).

A split of the zeros into close pairs and a separated remainder, whose Gram norm is bounded by a sinc-square Fourier majorant, was published earlier for simple zeros on the line, by gfreund123 in pull request #888 to GettysburgResearch/riemann (14 September 2026) [13] [V~]; ingredients (iii) and (iv) use a device of the same kind for a different spectral condition, and the priority is theirs. Related devices appear, for N_0^s , in trmdy’s second campaign [24] (off-line pairs priced as virtual on-line doubles, 11 August 2026, left open there) and in RakitinD’s exact decomposition of the slack with off-line and multiple-zero residuals [18] [V~]; neither is used there for N^s or N^d . In one line, for X the number of simple zeros: with $\text{Slack}_X := X - 2 \operatorname{tr} G + \|G\|_{\text{HS}}^2$, the prime side gives $X \geq (H(\psi) - o(1))N + \text{Slack}_X$, exactly as in [2]; the four ingredients show

$$\text{Slack}_X \geq a_1 X + \frac{a_2}{2}(N - X) - \nu N - o(N),$$

where a_1, a_2, ν are the certified constants; and solving the two inequalities for X gives $X \geq \Sigma N - o(N)$ with $\Sigma = (H + \frac{a_2}{2} - \nu)/(1 - a_1 + \frac{a_2}{2})$. The bound of [2] corresponds to $a_1 = a_2 = \nu = 0$.

Apart from the unrefereed claims of Yang and Yang discussed in §1.4, we know of no public work that treats the proportions of simple zeros or of distinct zeros beyond the standard consequences of a bound S for N_0^s : the bound S itself for N^s (since $N_0^s \leq N^s$), and for N^d the value $\frac{1}{2}(1 + S)$ that the chain of [2, Theorem A(ii)] gives for $S = H(\psi)$ and that of Theorem 1.2 gives for its constant (Remark 6.2; for Theorem 1.2 this is 0.8366868).

Theorem 1.2 and Corollary 1.3 are secondary, and we state their status plainly in §1.4. Corollary 1.3 is Lamzouri’s formula [15] for zeros that are simple or on the line, evaluated at the constant of Theorem 1.2; it is not an independent advance. For Theorem 1.2 the method appeared publicly before our work on it, and several unrefereed values exceed ours.

1.3 What the results are not

The theorems are lower bounds only. They do not show that the remaining zeros are multiple or off the line; they are merely not reached by the certificates. Theorem 1.1 counts simple zeros *whether or not* they lie on the critical line: a configuration in which a positive proportion of the zeros are simple off-line pairs is consistent with it, and it gives no information on how many zeros are on the line. It says nothing about the Riemann hypothesis. The proportion of zeros that are simple *and* on the line is not improved by the route of Theorem 1.1 (Remark 5.9); for that statistic this paper has only Theorem 1.2. The analytic inputs are those of [2]; in particular no zero-density estimate, zero-free region or mollifier is used. As observed in [2, §1.4], these inputs are insensitive to $o(N)$ zeros off the line, and the local inequalities added here concern only the kernel k_ψ . The local inequalities behind the constants are verified by computer; how, and at what level of trust, is stated in §1.5 and Section 7.

1.4 Relation to other work

The method of this paper is that of Alpöge and Furman [2], whose argument makes Montgomery’s pair-correlation deduction unconditional. Lamzouri [15] gave a second proof of their bound by a Hilbert-space inequality, and obtained the only previous value, 0.8876200, for the proportion of zeros that are simple or on the line. Wang [27] raised the constants, “by refining the method of Lamzouri”, to $C_0 + \delta_0$ for simple zeros on the line and $C_1 + \delta_0/2$ for distinct zeros, with $C_0 = \frac{3}{2} - \frac{1}{\sqrt{2}} \cot \frac{1}{\sqrt{2}} = 0.67250\dots$, $C_1 = \frac{1}{2}(C_0 + 1) = 0.83625\dots$ and $\delta_0 = 6.66624\dots \cdot 10^{-8}$, for $0 < \gamma \leq T$; truncated, these are 0.67250077 and 0.83625038 (our evaluation). We have read its abstract and the statement of its main theorem [V~]. Santibáñez-Leal [19] (24 September 2026, unrefereed) sharpens one step of Wang’s argument to 0.6725007995 and 0.8362503997 [V~]. Before 2026 the unconditional bounds for distinct zeros

included Farmer’s 0.63952 [11] and Wu’s 0.66036 [28], the latter cited as the record in [2]; Lamzouri [15] attributes a value above 70% to Ki and Lee [14], which we have not been able to check.

Simple zeros on the line: others’ work and what is ours. In August and September 2026 many unrefereed, computer-assisted claims above [2] for N_0^s/N appeared on GitHub and Zenodo. In this paragraph and throughout the paper, statements about other work carry the tag [V~] when we read them only through the authors’ descriptions or repository files, and [M] when they are our inference; we have built or run none of these works. As far as their descriptions show [V~], most of them and Theorem 1.2 share one mechanism, published first, as far as we found, by ainta [1] on 11 August 2026: the stability form of the rank–trace inequality, a local K -point functional with a length penalty, an offset-averaged block aggregation, and an interval branch-and-bound. The step from the local inequality to the global bound differs between works (a piecewise-linear subaction in Devine [10], a cap-one block assembly in Zhaoyi-Tian [32], a Bellman coboundary in tawanerguo-cn [21], a fusion polyhedron in RakitinD [18], a separation-or-pair transfer in pull request #888 [13]); we have not compared them [M]. Position weights and windows other than a single cosine appeared the same day in trmdy [24], and also in sxuff [20] (11 August), AMTOPA [3] (12 August) and Gebendorfer [12] (7 September); the aggregation function Φ_m of Lemma 3.3, with a proof, appeared the same day in tawanerguo-cn [21] and sxuff [20], and on 12 August in trmdy [V~]. Box pruning by convexity certificates and tangent planes appears in the verifiers of sxuff, AMTOPA and npip99 [17], and interval-Hessian basins in Zhaoyi-Tian’s certificate; we have not checked whether our box bound differs from these in substance [M]. Six unrefereed values that are, as far as we can tell, bandwidth-one certificates [M] exceed Theorem 1.2, by $2.5 \cdot 10^{-5}$ to $1.18 \cdot 10^{-4}$: gfreund123, in pull request #888 to the repository GettysburgResearch/riemann [13], 0.67349239, a draft pull request, which states that independent analytic review is required and that “no independently reviewed world-record claim is made”; it states its bound for $0 < \gamma \leq T$ and builds on the window and one local inequality of Gebendorfer’s draft [V~]; RakitinD [18], 0.6734775921119, described as a frozen computer-assisted result that builds on a comparator from Gebendorfer’s work, with, in its author’s words, a “partial Lean formalization with explicit remaining boundaries”; stated for $0 < \gamma \leq T$ [V~, M]; Jonas Jakob Gebendorfer [12], 0.6734760237, called by its author an internally checked research draft; stated for $0 < \gamma \leq T$ [V~]; AMTOPA [3], 0.673416490971, described by its authors as an interval-certified research-draft record [V~]; Zhaoyi-Tian [32], 0.6734, announced as a Lean theorem with no hypotheses, checked by the kernel with the three standard axioms, on the same Lean and Mathlib [23] versions as ours, with a kernel-replayed certificate of 3,203 leaves, crediting ainta (the retained spectral defect and seven consecutive zeros), sxuff and AMTOPA’s earlier repository (position-dependent and exact shifted pressure) and AMTOPA (the corrected window, pair loads and pressure data) for ingredients it uses [V~]; we have not built it. By its description it is a formally verified bound above Theorem 1.2. Zhaoyi-Tian announced 0.673 as a hypothesis-free Lean theorem on 31 August and 0.6734 on 3 September 2026; MichaelMobius [16] announces a hypothesis-free Lean theorem for 0.6731175 (latest commits 21–22 September 2026), and we have not established which of these was complete first; an end-to-end Lean reduction of the bound 0.673195 to a single local-certificate premise was published by npip99 [17] on 11–12 August 2026 [V~]; and Michael Devine [10], 0.673399 (stated also for simple and for distinct zeros, as a consequence), whose aggregation by a piecewise-linear subaction appears to differ from the block aggregation used here and by the others [V~, M]. Devine’s later record [9] announces a computation-free improvement of the Montgomery–Taylor constant, with a Lean formalisation in which the published pair-correlation and Riemann–von Mangoldt inputs are represented by explicit interfaces [V~]. Further claims lie between 0.67300 and 0.67332 (among them ainta’s 0.673008527927

and trmdy’s 0.6733127422)¹; teal-sea/zeta-lab [22] reports Lean proofs of small- K bounds (0.6728470197), which, according to its README, the Palomar registry rebuilt and replayed [V~]. What is ours for Theorem 1.2, as far as we have found, is the even polynomial window with exact rational constants as a window for ζ (its H is an exact rational; Alpöge and Furman used a quartic polynomial window for the zeros of ξ' [2, Remark 7.1]), and the box bound LB3 in its exact form (§7.2), which we have not compared in substance with the pruning of other verifiers. The aggregation function Φ_m of Lemma 3.3, which Theorem 1.2 needs on its concave branch since its block energy $A = 1.1144$ exceeds $m/(m-1)$, is not ours: it was published with proofs by tawanerguo-cn [21] and sxuff [20] on 11 August 2026 and by trmdy [24] on 12 August, each using it on its concave branch; yuhangshi888 [31] uses a supporting plane of it, and Zhaoyi-Tian prove in Lean a cap-one variant [V~]. We have verified none of these proofs.

Claims by other methods. The ceiling ≈ 0.682 of [2, §7.2] concerns bandwidth-one certificates that use only the first two trace moments (Section 8). For N^s the largest value that the bandwidth-one claims above would give is 0.67349239; for N^d , if their chains give $\frac{1}{2}(1+S)$ as those of [2, Theorem A(ii)] and of Theorem 1.2 do (Remark 6.2), it is $\frac{1}{2}(1+0.67349239) = 0.8367461$. Both lie below Theorem 1.1. Hongyi Yang and Shihua Yang claim much more, by methods that use trace moments beyond the first two [V~]. In Zenodo record 21975237 [29] they claim 0.7962 for simple zeros on the line and 0.8981 for distinct zeros, with ineffective constants, a result they declare “at the certified-candidate level, not as an established theorem”; according to the record, its identity and certificate layers are compiled and kernel-checked in Lean 4, while its analytic layer is not yet formalised and awaits external peer review [V~]. In Zenodo record 22065921 [30], Shihua Yang and Hongyi Yang claim that almost all zeros of ζ are simple and lie on the critical line, in the density sense, unconditionally and with ineffective constants, with the explicit bounds 0.910460 for simple zeros on the line and 0.955230 for distinct zeros; the record describes the analytic chain as “a certified candidate pending external review” and states that its certificates are kernel-checked in Lean 4 [V~]. The record page of [30] shows its files under embargo, with the stated reason that the version contains errors and needs revision [V~, read 28 September 2026]. We have not verified either claim; if they are correct, they supersede Theorems 1.1–1.2 and Corollary 1.3.

1.5 Formal verification

The four results are formalised in Lean 4 [8], in a library **ZetaS** built on the Lean artifact **Zeta23** of [2] (toolchain v4.33.0-rc2, Mathlib [23] revision 51e6992), against the zero-counting functions that [2] define directly from Mathlib’s **riemannZeta** (Appendix A). We say that a declaration is *kernel-checked* if `#print axioms` reports only `propext`, `Classical.choice` and `Quot.sound`: no `sorry`, no added axiom, no `native_decide`. The formalisation is in the repository <https://github.com/Oliverds321/simple-zeros>, at tag v1.0 (Appendix A.9). The trust level of each result is as follows.

- (a) *Kernel-checked with no hypotheses.* The two headline theorems for the $K = 5$ pair (0.675158622, 0.837579311; Appendix A.6), which are in the library **ZetaSReplay** of the repository, built separately from **ZetaS**. Further, the inputs of Proposition 2.1 for the two windows used, at every bandwidth $\lambda < 1$ (from the artifact of [2] and a window-general version of its pipeline); every lemma, proposition and theorem of Sections 2–6 in its abstract form (not the remarks, among them Remark 6.2, and not the examples

¹The others we know of [V~]: yuhangshi888 [31] (0.6733169771), AMTOPA’s earlier 0.673262375585, sxuff [20] (0.673205978423), npip99 [17] (0.673195198901), tawanerguo-cn [21] (0.6731929114731422), MichaelMobius [16] (0.6731175265, announced with a hypothesis-free Lean proof), vstaln [26] (0.67306445176), Zhongshan-Big-Jun [33] (0.6730664726 of 15 August; its README also lists 0.6730536459 of 14 August) and uwe-schwarz [25] (0.6730266625).

of sharpness), with the qualifications stated after Lemmas 2.2 and 3.3, where the formal version is weaker than the printed statement but covers every use: the linear algebra of Section 3; the local-to-global steps of Sections 4–5 (aggregation, packing, transfer, tight chains, the room lemma; the on-line spectral condition and the deduction of Theorem 1.1 from its certificates only for the two certificates used, Theorems 5.7 and 5.8 for arbitrary certificate data being a mechanical generalisation not yet formalised; the deductions of Theorem 1.2 and Corollary 1.3 from their certificate); the assembly from the zeros of `riemannZeta` to the matrices of Section 2 at every bandwidth $\lambda < 1$ (Poisson and Gram identities, positive definiteness of the kernel, the truncation defect, the inertia of off-line pairs), together with the limit $\lambda \rightarrow 1^-$; the closing arithmetic of all constants; and the majorant certificate of Lemma 5.5, replayed by the kernel in 614 cells (about 88 seconds).

- (b) *Kernel-checked with one displayed hypothesis, a numerical certificate.* In the library the six headline theorems each have one hypothesis: `CertAM5` for the $K = 5$ pair, `CertAM7` for the $K = 7$ pair (0.676102666, 0.838051333), and `CertS8` for Theorem 1.2 and Corollary 1.3. Each of these is one `Prop` asserting that a local inequality holds for all gap vectors (Appendix A); the data that realise it are identified by sha256 in Table 3, and they are verified outside Lean by interval arithmetic (§7.2). Nothing else is assumed: in the library at commit `bb7223b`, `#print axioms` for each of the six headline theorems gives only the three standard axioms (Appendix A.8). The fourteen statements of record (the six headline statements, their six cumulative forms on $0 < \text{Im } \rho \leq T$ and two non-vacuity companions) are in the comparator package of the library (Appendix A.3), and all fourteen are proved with the three standard axioms. For the $K = 5$ pair the hypothesis is discharged by the kernel replay, (a) and (c). The $K = 7$ certificate and the certificate behind Theorem 1.2 are too large for kernel replay on the hardware used (§7.3) and stay as displayed hypotheses.
- (c) *The kernel replay of the $K = 5$ certificate.* The kernel checker of the replay is version 3, which admits a tangent line only when the tangent point lies in the span interval (without this guard a leaf check could accept a false claim, §7.3). The checker, its soundness nodes and the chain `ZetaS.CertV2.ChainV3.certAM5_of_checks`, which derives `CertAM5` from the acceptance of the 20 class certificates by the checker and the identity of their data with the data of `CertAM5`, are in the library with the three standard axioms (Appendix A.5). The full replay, against the library at commit `bb7223b`, whose Lean files the release changes only in comments (Appendix A.9), compiled all 1,478 files with no failure, and the kernel checked the 46,047 leaves and 89,362 point records in exact rational arithmetic (§7.3). Its top files prove `ZetaS.CertV2.certAM5_replayed : ZetaS.CertAM5` and, from it, the two theorems of the $K = 5$ pair without hypothesis (Appendix A.6). Each of the three depends only on `propext`, `Classical.choice` and `Quot.sound` and rests on no open declaration (Appendix A.8). The 1,478 source files match their manifest file by file, and a scan of all of them finds no `sorry`, `native_decide`, `axiom`, `unsafe`, `implemented_by` or `extern`. The replay files, byte-identical to those checked, and the two theorems are in the repository as the library `ZetaSReplay`, outside the default build targets; `lake build ZetaSReplay` repeats the replay (§7.4). The certificates themselves were found by floating-point and interval computation; the kernel checked them.

Earlier formal work. Formal proofs of bounds of this kind were announced before ours [V~]: Zhaoyi-Tian’s hypothesis-free Lean theorems for N_0^s , at 0.673 (31 August) and 0.6734 (3 September 2026) [32]; a hypothesis-free Lean theorem for N_0^s at 0.6731175 announced by MichaelMobius [16] (latest commits 21–22 September 2026; we could not establish when its Lean theorem was first complete); npip99’s end-to-end Lean reduction of 0.673195 to one

local-certificate premise [17] (11–12 August); Lean-proved small- K bounds for N_0^s in teal-sea/zeta-lab [22], which, according to its README, the Palomar registry rebuilt; the finite deduction of yuhangshi888 [31] in Lean; Devine’s formalised computation-free improvement of the Montgomery–Taylor constant [9], with the pair-correlation and Riemann–von Mangoldt inputs as interfaces; partial formalisations by RakitinD [18] and of the certificate layers of Yang and Yang [29, 30]; and a formalisation of Lamzouri’s proof, with the Riemann–von Mangoldt formula and a pair-correlation asymptotic as stated hypotheses, by AxiomMath [4]. Our formal theorem for simple zeros on the line, Theorem 1.2, has a smaller constant than the one Zhaoyi-Tian announces, and its certificate is a hypothesis. Its formal status is that of npip99’s earlier theorem. The Lean artifact of [2], on which our library is built, already proves hypothesis-free bounds for N^d ($\frac{3}{4}$ and $\frac{5}{6}$, and $0.75329\dots$ and $0.83625\dots$ with the Montgomery–Taylor window) as well as for N_0^s and for the distinct zeros on the line; apart from these and, under its hypotheses, AxiomMath’s formalisation of Lamzouri’s bounds, we know of no earlier formal result for N^s , N^d or N^{sc} (Yang and Yang state that the identity and certificate layers of their claims, which include N^d , are kernel-checked in Lean 4).

1.6 Plan

Section 2 fixes the setting of [2] and the inputs taken from it, at bandwidth $\lambda < 1$. Section 3 collects the linear algebra. Section 4 proves Theorem 1.2, Section 5 Theorem 1.1, and Section 6 Corollary 1.3. Section 7 describes the certificates and Section 8 the limits of the method. Appendix A reproduces the Lean statements; Appendix B records how the work was done.

2 Setting and the inputs from Alpöge–Furman

We follow the notation of [2, §§1–2]. Throughout, $\ell := \log(T/2\pi)$ and $N := N(T, 2T) = \frac{T}{2\pi}(\ell + 2\log 2 - 1) + O(\ell)$. For a zero ρ put $\gamma_\rho := (\rho - \frac{1}{2})/i \in \mathbb{C}$, which is real if and only if ρ lies on the critical line. Off-line zeros come in pairs $\pi = \{\rho, 1 - \bar{\rho}\}$ with $\gamma_{1-\bar{\rho}} = \bar{\gamma}_\rho$ and equal multiplicities.

2.1 The window, the test family and the matrix

Fix a *bandwidth* $\lambda \in (0, 1]$ and put $L := \lambda\ell$. A *window* is an even function $\psi \in C^2([-\frac{1}{2}, \frac{1}{2}])$ with $0 < \psi \leq 1$. (The hypotheses of [2, §2.2] are that ψ is even, C^2 and positive; the bound $\psi \leq 1$ is used there tacitly, through $0 \leq \varphi \leq 1$.) Fix a smooth nondecreasing function χ with $\chi = 0$ on $(-\infty, 0]$ and $\chi = 1$ on $[1, \infty)$, and let $\chi_w(t) := \chi(t/w)$, the ramp of width w , for $1 \leq w = o(L)$; [2] take a C^∞ ramp of width 1. Implied constants may depend on χ and ψ . Put, as in [2, (2.7)],

$$\varphi(u) := \chi_w\left(\frac{L}{2} + u\right) \chi_w\left(\frac{L}{2} - u\right) \psi(u/L)^{1/2}, \quad (1)$$

so that φ is real and even, $\text{supp } \varphi = [-\frac{L}{2}, \frac{L}{2}]$ and $\varphi(\pm \frac{L}{2}) = 0$. Let $a := L^{-1} \int \varphi^2 = \int \psi + O(w/L)$ and $\widehat{\varphi}(z) := \int \varphi(u) e^{-iuz} du$, an even entire function, real on $\mathbb{R}$, with $\widehat{\varphi}(\bar{z}) = \overline{\widehat{\varphi}(z)}$. The grid is $\alpha_k := T + 2\pi k/L$ for $0 \leq k < d := \lfloor LT/2\pi \rfloor$, and for a zero ρ we put

$$v_\rho := (\widehat{\varphi}(\gamma_\rho - \alpha_k))_{0 \leq k < d} \in \mathbb{C}^d, \quad u_\rho := (aL^2)^{-1/2} v_\rho.$$

Let $I := [T, 2T)$ and $I' := [T - \sqrt{T}, 2T + \sqrt{T})$. The matrix of [2, (2.10)] is

$$G := \sum_{\text{Re } \gamma_\rho \in I'} m_\rho u_\rho u_\rho^\top, \quad (2)$$

the sum running over distinct zeros, with the bilinear transpose. It is real symmetric: $v_{1-\bar{\rho}} = \overline{v_\rho}$ because the grid is real and $\widehat{\varphi}(\bar{z}) = \overline{\widehat{\varphi}(z)}$, the multiplicities of ρ and $1 - \bar{\rho}$ agree, and both have

the same $\operatorname{Re} \gamma$, so that they enter (2) together. An on-line zero contributes $m_\rho u_\rho u_\rho^\top \succeq 0$. An off-line pair with $u_\rho = x + iy$ ($x, y \in \mathbb{R}^d$) contributes $m_\pi(u_\rho u_\rho^\top + \bar{u}_\rho \bar{u}_\rho^\top) = 2m_\pi(xx^\top - yy^\top)$, which has at most one positive and at most one negative eigenvalue.

2.2 The kernel and the window constant

For a window ψ put $v_\psi := \psi / \int \psi$ and

$$k_\psi(x) := \int_{-1/2}^{1/2} v_\psi(s) \cos(2\pi xs) ds,$$

an even, positive definite function with $k_\psi(0) = 1$ and $|k_\psi| \leq 1$. Let

$$J(\psi) := \iint |u - v| \psi(u) \psi(v) du dv, \quad R_\lambda(\psi) := \frac{\int \psi^2 + \lambda^2 J(\psi)}{\lambda (\int \psi)^2}, \quad H_\lambda(\psi) := 2 - R_\lambda(\psi),$$

all integrals over $[-\frac{1}{2}, \frac{1}{2}]$, and $R := R_1$, $H := H_1$. Substituting $v = u + r$ gives the one-integral form $J(\psi) = 2 \int_0^1 r (\psi \star \psi)(r) dr$ with $(\psi \star \psi)(r) := \int \psi(u) \psi(u + r) du$; the prime side of [2] produces this form, and the two agree for every continuous ψ (evenness is not needed). R is the constant of [2, Lemma 5.6]; for $\psi = \cos(\sqrt{2}s)$, $H = \frac{3}{2} - \cot(1/\sqrt{2})/\sqrt{2} = 0.6725007036\dots$, which is optimal among constants of the form $2 - R(\psi)$ [6]. Every quantity used below is homogeneous of degree 0 in ψ , so a window may be rescaled by a positive constant without changing k_ψ , R_λ or H_λ . For $\psi(s) = \cos(\alpha s)$,

$$k_\psi(x) = \frac{\operatorname{sinc}(\pi x - \frac{\alpha}{2}) + \operatorname{sinc}(\pi x + \frac{\alpha}{2})}{2 \operatorname{sinc}(\frac{\alpha}{2})}, \quad \operatorname{sinc} y = \frac{\sin y}{y} \ (y \neq 0), \quad \operatorname{sinc} 0 = 1. \quad (3)$$

2.3 The inputs

Proposition 2.1 (Inputs from [2]). *Let ψ be a window and $0 < \lambda < 1$. As $T \rightarrow \infty$:*

- (AF1) ([2, Lemma 2.1]) $\|u_\rho\|^2 \leq 1$ for every on-line zero.
- (AF2) ([2, Lemma 3.1 and Proposition 4.1]) *Each off-line pair contributes to G a real symmetric matrix with at most one positive and at most one negative eigenvalue.*
- (AF3) ([2, Proposition 4.2]) $\operatorname{tr} G = N(I') + O(T^{1/2}L^2)$, and $N(I') = N + O(\sqrt{T} \log T)$.
- (AF4) ([2, Theorem 5.7] at $\lambda = 1$; for $\lambda < 1$ its proof, as indicated in [2, Remark 6.1], which states the case $\psi = \psi_0$; the general formula for $R_\lambda(\psi)$ is formally verified for the two windows used here) $\|G\|_{\text{HS}}^2 = (R_\lambda(\psi) + o(1))N$.

The prime-side input behind (AF3)–(AF4) is the Montgomery–Vaughan mean value theorem for a Dirichlet polynomial of length $(T/2\pi)^\lambda$, as in [2]. [2] prove (AF4) also at $\lambda = 1$, with error $O(L^{-1})$; all arguments below go through verbatim there. The formal proof works at $\lambda < 1$, where its prime-side error term is $o(N)$, and then lets $\lambda \rightarrow 1^-$ (§2.5); we follow it.

2.4 Normalised ordinates and the Gram identity

For an on-line zero put $x_\rho := \gamma_\rho L / 2\pi$ (the *normalised ordinate*) and, for $k \in \mathbb{Z}$, $U_\rho(k) := (aL^2)^{-1/2} \widehat{\varphi}(\gamma_\rho - \alpha_k)$, so that $u_\rho = (U_\rho(k))_{0 \leq k < d}$. Let

$$\delta_\rho := \sum_{k \notin [0, d)} U_\rho(k)^2, \quad \widehat{E}_{\rho\rho'} := \sum_{k \notin [0, d)} U_\rho(k) U_{\rho'}(k), \quad k_\varphi(\xi) := \frac{\int \varphi(Ls)^2 e^{2\pi i \xi s} ds}{\int \varphi(Ls)^2 ds}.$$

Let $v_\varphi := \varphi(L\cdot)^2 / \int \varphi(L\cdot)^2$, $\delta_0 := \int v_\psi (1 - \varphi(L\cdot)^2 / \psi)$, and

$$\varepsilon_T := \|v_\varphi - v_\psi\|_{L^1}, \quad \varepsilon'_T := \frac{\delta_0}{1 - \delta_0}.$$

Both tend to 0 because $w = o(L)$. Finally let $\Lambda := |I'|L/2\pi$ be the normalised length of I' .

Lemma 2.2 (Gram identity and residues). *Let T be so large that $a \geq \frac{1}{2} \int \psi$, so that $\int \varphi^2 > 0$. For on-line zeros ρ, ρ' with $\gamma_\rho, \gamma_{\rho'} \in I'$:*

- (i) $\sum_{k \in \mathbb{Z}} U_\rho(k) U_{\rho'}(k) = k_\varphi(x_\rho - x_{\rho'})$; hence $\|u_\rho\|^2 = 1 - \delta_\rho$ with $0 \leq \delta_\rho \leq 1$, and $\langle u_\rho, u_{\rho'} \rangle = k_\varphi(x_\rho - x_{\rho'}) - \widehat{E}_{\rho\rho'}$, where $\widehat{E} \succeq 0$ and $|\widehat{E}_{\rho\rho'}| \leq (\delta_\rho \delta_{\rho'})^{1/2}$.
- (ii) k_φ is even and positive definite with $k_\varphi(0) = 1$, so $|k_\varphi| \leq 1$; moreover $\sup_{\mathbb{R}} |k_\varphi - k_\psi| \leq \varepsilon_T$ and $v_\varphi \leq (1 + \varepsilon'_T) v_\psi$, with $\varepsilon'_T \geq 0$.
- (iii) $\sum_\rho m_\rho \delta_\rho = o(N)$ (indeed $\ll \sqrt{T} \log T$), the sum over on-line zeros with $\gamma_\rho \in I'$.
- (iv) The consecutive gaps between the normalised ordinates of any set of zeros with $\operatorname{Re} \gamma \in I'$ add up to at most $\Lambda = \lambda N(1 + o(1))$; every multiplicity is $\ll \log T$.

Proof. (i) With $f(u) := \varphi(u) e^{-iu(\gamma - T)}$, supported in $[-\frac{L}{2}, \frac{L}{2}]$ and vanishing at the ends, $\widehat{\varphi}(\gamma - \alpha_k)$ is L times a Fourier coefficient of f on the period L ; Parseval gives the identity [2, Lemma 2.1], and the case $\rho = \rho'$ gives $\sum_{k \in \mathbb{Z}} U_\rho(k)^2 = 1$. Subtracting the off-grid part gives the rest; $\widehat{E}$ is a Gram matrix, and Cauchy–Schwarz gives the bound on its entries. (The normalisation needs $\int \varphi^2 > 0$: for the zero window the identity fails, $k_\varphi(0)$ being 0.) (ii) Every kernel matrix $(k_\varphi(x_i - x_j))$ is the Gram matrix of full-lattice vectors by (i), so k_φ is positive definite, and it is even because φ is. Next $v_\varphi = (1 - \delta_0)^{-1} (\varphi(L\cdot)^2 / \psi) v_\psi \leq (1 + \varepsilon'_T) v_\psi$, since the ramp only removes mass; and $|k_\varphi - k_\psi| \leq \|v_\varphi - v_\psi\|_1$. (iii) For real z , two integrations by parts give $|\widehat{\varphi}(z)| \leq |z|^{-2} \int |\varphi''|$, and $|\widehat{\varphi}(z)| \leq \int |\varphi| \leq L\sqrt{a}$; with $\xi = (\gamma_\rho - \alpha_k)L/2\pi$ this is $|U_\rho(k)| \leq \min\{1, \lambda_\varphi \xi^{-2}\}$ with $\lambda_\varphi := L \int |\varphi''| / (4\pi^2 \sqrt{a}) \ll L/w + 1$ (here $a \geq \frac{1}{2} \int \psi$ and $\int |\chi''_w| = w^{-1} \int |\chi''|$ are used). The grid covers $[T, 2T)$, and $\delta_\rho \leq 1$ always. The zeros with $\gamma_\rho \in I' \setminus [T, 2T)$ number $\ll \sqrt{T} \log T$, and those in $[T, 2T)$ within distance 1 (in γ) of T or $2T$ number $\ll \log T$. The others lie in shells of unit length in γ at distance $j \geq 1$ from the ends of $[T, 2T)$, each with $\ll \log T \ll L$ zeros; for such a zero every off-grid k has $|\xi| \geq jL/2\pi$, so $\delta_\rho \ll \lambda_\varphi^2 (jL)^{-3}$, and these zeros contribute $\ll \sum_j L \lambda_\varphi^2 (jL)^{-3} \ll 1$. (iv) is the Riemann–von Mangoldt formula and $N(t+1) - N(t) \ll \log t$. $\square$

The formal proof has (ii) in a weaker form, the positive definiteness of $(1 + \varepsilon'_T)k_\psi - k_\varphi$ in place of the pointwise bound $v_\varphi \leq (1 + \varepsilon'_T)v_\psi$, and (iii) with $o(N)$ only; these are all that is used below.

At bandwidth λ the normalised ordinates have mean spacing λ , not 1. Every local inequality below is stated in these units, where the kernel between two on-line zeros is $k_\varphi(x_\rho - x_{\rho'})$ at every λ ; the certificates therefore apply at every λ without change, and $\Lambda \leq N(1 + o(1))$ is the only place where the density enters.

2.5 The limit $\lambda \rightarrow 1^-$

Each of our bounds is proved in the form: for every $\lambda \in (0, 1)$ and every $\varepsilon > 0$, $X(T, 2T) \geq (f(\lambda) - \varepsilon)N$ for $T \geq T_0$, where $f(\lambda)$ is obtained from the bandwidth-one constant by replacing $H(\psi)$ by $H_\lambda(\psi)$. Since $R_\lambda(\psi) \rightarrow R(\psi)$ as $\lambda \rightarrow 1^-$ and each constant is continuous (indeed affine) in H , letting $\lambda \rightarrow 1^-$ gives the bandwidth-one constants. For instance the constant Σ of Theorem 1.1 at $K = 7$ equals 0.5950, 0.66902 and 0.675404 at $\lambda = 0.9, 0.99$ and 0.999. (Equivalently, in units of the mean spacing the certificates transfer by the exact identity $F_{\gamma, \lambda\mu}^{(\lambda)}(h) = F_{\gamma, \mu}(\lambda h)$, where $F^{(\lambda)}$ is built from $k_\psi(\lambda \cdot)$; so no slack is consumed. The formal proof does not need this, because it works in normalised units throughout.)

3 Linear algebra

All matrices are finite. For a Hermitian matrix X , $\lambda_1(X) \geq \lambda_2(X) \geq \dots$ are its eigenvalues, $n_{\pm}(X)$ the numbers of positive and negative eigenvalues, $\|X\|_{\text{HS}}^2 = \text{tr } X^2$, and $\text{tr } f(X) = \sum_i f(\lambda_i(X))$ for a function f on the spectrum.

3.1 The stability rank–trace inequality

Let

$$\Psi(t) := \begin{cases} (t-1)^2, & 0 \leq t \leq 2, \\ 2t-3, & t \geq 2, \end{cases} \quad (4)$$

so that $\Psi \geq 0$, $\Psi(0) = 1$, Ψ is 2-Lipschitz on $[0, \infty)$, and $\Psi(t) = (t-1)^2 - ((t-2)_+)^2$. It is convex, being the supremum of the tangent lines to $(t-1)^2$ at the points $s \leq 2$.

Lemma 3.1 (Stability rank–trace inequality). *Let P, Q be Hermitian $d \times d$ matrices with $P = VV^*$, $V \in \mathbb{C}^{d \times r}$, and $n_+(Q) \leq b$. Put $M := V^*V \in \mathbb{C}^{r \times r}$. Then*

$$r \geq 2 \text{tr } P + 4 \text{tr } Q - 4b - \|P + Q\|_{\text{HS}}^2 + \text{tr } \Psi(M). \quad (5)$$

Dropping $\text{tr } \Psi(M) \geq 0$ gives [2, Lemma 3.2]. Equality holds for $P = \Pi_1$, $Q = 2\Pi_2$ with orthogonal projections $\Pi_1 \perp \Pi_2$. This form of the inequality, with the term $\text{tr } \Psi(M)$, was published by ainta [1] on 11 August 2026; we know of no earlier source [V~].

Proof. The nonzero eigenvalues of VV^* and V^*V agree with multiplicity, so $\sum_{i \leq r} g(\lambda_i(M)) = \sum_{j \leq d} g(\lambda_j(P))$ for every g with $g(0) = 0$. Applied to $g = \Psi - 1$ this gives $\text{tr } \Psi(M) - r = \sum_{j \leq d} (\Psi(p_j) - 1)$, where $p_1 \geq \dots \geq p_d \geq 0$ are the eigenvalues of P ; no case distinction between $r \leq d$ and $r > d$ is needed. Write $Q = Q_+ - Q_-$ with $Q_{\pm} \succeq 0$, $Q_+Q_- = 0$; let $q_1, \dots, q_{b'}$ ($b' \leq b$) be the positive eigenvalues of Q_+ and $n_1 \geq n_2 \geq \dots \geq 0$ those of Q_- . Since $\text{tr}(PQ_+) \geq 0$ and $\text{tr}(Q_+Q_-) = 0$,

$$\|P + Q\|_{\text{HS}}^2 \geq \|P - Q_-\|_{\text{HS}}^2 + \sum_j q_j^2 \geq \sum_i (p_i - n_i)^2 + \sum_j q_j^2,$$

the second step by von Neumann's trace inequality $\text{tr}(PQ_-) \leq \sum_i p_i n_i$. Using $\text{tr } P = \sum p_i$, $\text{tr } Q = \sum q_j - \sum n_i$ and the identity $1 - 2p + 4n + (p - n)^2 = (p - n - 1)^2 + 2n$,

$$d - [2 \text{tr } P + 4 \text{tr } Q - 4b - \|P + Q\|_{\text{HS}}^2] \geq \sum_{i \leq d} [(p_i - n_i - 1)^2 + 2n_i] + \sum_{j \leq b'} (q_j - 2)^2 + 4(b - b').$$

For $p \geq 0$ the minimum over $n \geq 0$ of $(p - n - 1)^2 + 2n$ is attained at $n = (p - 2)_+$ and equals $\Psi(p)$. Hence the right side is at least $\sum_{i \leq d} \Psi(p_i)$, and subtracting $d - r = \sum_{i \leq d} 1 - r$ gives (5) by the first sentence. $\square$

Lemma 3.2 (Two-parameter form). *In the setting of Lemma 3.1, let $0 \leq s \leq t$ and*

$$\Psi_{s,t}(p) := \begin{cases} (p-s)^2, & 0 \leq p \leq t, \\ (t-s)^2 + 2(t-s)(p-t), & p \geq t. \end{cases}$$

Then $s^2r + t^2b \geq 2s \text{tr } P + 2t \text{tr } Q - \|P + Q\|_{\text{HS}}^2 + \text{tr } \Psi_{s,t}(M)$. For $(s, t) = (1, 2)$ this is Lemma 3.1.

Proof. As above, by the identity $s^2 - 2sp + (p - n)^2 + 2tn = (p - n - s)^2 + 2(t - s)n$,

$$s^2d + t^2b - [2s \text{tr } P + 2t \text{tr } Q - \|P + Q\|_{\text{HS}}^2] \geq \sum_{i \leq d} [(p_i - n_i - s)^2 + 2(t - s)n_i] + \sum_{j \leq b'} (q_j - t)^2 + t^2(b - b').$$

For $p \geq 0$ the minimum over $n \geq 0$ of $(p-n-s)^2 + 2(t-s)n$ is attained at $n = (p-t)_+$ and equals $\Psi_{s,t}(p)$, so the right side is at least $\sum_{i \leq d} \Psi_{s,t}(p_i)$. Since $\Psi_{s,t}(0) = s^2$ (as $t \geq 0$), the first sentence of the proof of Lemma 3.1, applied to $g = \Psi_{s,t} - s^2$, gives $\text{tr } \Psi_{s,t}(M) - s^2 r = \sum_{i \leq d} \Psi_{s,t}(p_i) - s^2 d$; subtracting $s^2(d-r)$ from both sides of the display gives the claim. $\square$

$\Psi_{s,t}$ is convex (the supremum of the tangents to $(p-s)^2$ at the points $\leq t$), C^1 and nonnegative, and $\Psi_{1,t} \geq \Psi$ for $t \geq 2$ and every real p . The one-parameter case $(s,t) = (c/2, c)$, which contains Lemma 3.1 ($c = 2$), is proved in Lean in Zhaoyi-Tian's development [32], in the form $c \text{tr } P + \sum_i \text{gc}_c(\lambda_i) + 2c \text{tr } Q - c^2 b \leq \|P + Q\|_{\text{HS}}^2$ [V~; that this is the case $(c/2, c)$ of Lemma 3.2 is our reading, M]. Corollary 1.3 needs the point $(1, 2 + \sqrt{2})$, which lies off that ray.

3.2 Energy and the function Φ_m

Lemma 3.3 (Aggregation function). *Let $G \succeq 0$ be $m \times m$ with unit diagonal, $m \geq 2$, and $E := \|G - I\|_{\text{HS}}^2$. Then $\text{tr } \Psi(G) \geq \Phi_m(E)$, where*

$$\Phi_m(E) := \begin{cases} E, & E \leq \frac{m}{m-1}, \\ E - \left(\sqrt{\frac{(m-1)E}{m}} - 1 \right)^2, & E \geq \frac{m}{m-1}. \end{cases}$$

Φ_m is increasing and concave on $[0, \infty)$, $\Phi_m(0) = 0$ and $\Phi_m(E) \leq E$; equality holds for $G = (1-t)I + tJ$, $0 \leq t \leq 1$.

Proof. Let $1 + u_i$ be the eigenvalues of G ; then $\sum u_i = 0$, $\sum u_i^2 = E$ and $u_i \geq -1$, and $\text{tr } \Psi(G) = E - \sum_{i \in S} (u_i - 1)^2$ with $S := \{i : u_i > 1\}$. If $S = \emptyset$ we are done. Otherwise let $j := |S|$ (so $1 \leq j < m$), $\sigma := \sum_S u_i > j$, $F := \sum_S (u_i - 1)^2$ and $V := 1 + \sqrt{F}$. The numbers $u_i - 1$, $i \in S$, are positive with sum $\sigma - j$, so $\sqrt{F} \leq \sigma - j$; also $\sum_S u_i^2 = F + 2(\sigma - j) + j$, and the $m - j$ remaining u_i have sum $-\sigma$, hence sum of squares $\geq \sigma^2/(m-j)$. So

$$E \geq (V-1)^2 + 2(V-1) + j + \frac{(V+j-1)^2}{m-j} = V^2 + (j-1) + \frac{(V+j-1)^2}{m-j} \geq V^2 \frac{m}{m-1},$$

with equality in the last step for $j = 1$, and for $j \geq 2$ because $V+j-1 \geq V$ and $m-j < m-1$. Thus $V \leq U := \sqrt{E(m-1)/m}$. If $E \leq m/(m-1)$ then $U \leq 1 \leq V$ forces $F = 0$; otherwise $F = (V-1)^2 \leq (U-1)^2$. The two branches have slope 1 at $E = m/(m-1)$ and the second is concave. $\square$

The inequality $\text{tr } \Psi(G) \geq \Phi_m(E)$ is kernel-checked (the formal version does not need $m \geq 2$); the properties of Φ_m listed after it in the lemma and the equality case are not formalised, and the formal proofs do not use them. The function Φ_m and this inequality were published earlier by tawanerguo-cn [21] and sxuff [20] (11 August 2026) and by trmdy [24] (12 August) [V~].

We use Lemma 3.3 on its concave branch in Theorem 1.2. A consequence used there: for $0 \leq E$ and $0 < A$, $\Phi_m(E) \geq \frac{\Phi_m(A)}{A} \min(E, A)$, by concavity and $\Phi_m(0) = 0$.

3.3 Pinching and monotonicity

Lemma 3.4 (Pinching). *Let $G \succeq 0$ be real symmetric and let f be convex on $[0, \infty)$. For every partition of the index set into blocks b , $\text{tr } f(G) \geq \sum_b \text{tr } f(G_{bb})$.*

Proof. For each block let $w_{b,1}, w_{b,2}, \dots$ be an orthonormal eigenbasis of G_{bb} , extended by zero outside b . These vectors form an orthonormal basis, in particular a Parseval frame of unit vectors, and $\langle w_{b,j}, G w_{b,j} \rangle = \lambda_j(G_{bb})$. For a unit vector w , $\langle w, G w \rangle = \sum_k c_k^2 \lambda_k(G)$ with $\sum_k c_k^2 = 1$, where c_k are the coordinates of w in an eigenbasis of G ; by Jensen, $f(\langle w, G w \rangle) \leq \sum_k c_k^2 f(\lambda_k(G))$. Summing over the frame, and using $\sum_{b,j} c_k(w_{b,j})^2 = 1$ for each k , $\sum_b \text{tr } f(G_{bb}) \leq \text{tr } f(G)$. $\square$

No convexity of $X \mapsto \operatorname{tr} f(X)$ is needed. The lemma applies to Ψ , to $\Psi_{s,t}$ and to $(2-t)_+^2$.

Lemma 3.5 (Monotonicity). *Let A, B be positive semidefinite of the same size with $A - B \succeq 0$, and let f be c -Lipschitz on $[0, \infty)$. Then $|\operatorname{tr} f(A) - \operatorname{tr} f(B)| \leq c \operatorname{tr}(A - B)$.*

Proof. By Weyl's inequality $\lambda_i(A) \geq \lambda_i(B)$ for every i , so $|\operatorname{tr} f(A) - \operatorname{tr} f(B)| \leq c \sum_i (\lambda_i(A) - \lambda_i(B)) = c \operatorname{tr}(A - B)$. $\square$

This is the form that is formalised and used (in Lemma 4.4); the same proof gives it for Hermitian A, B and f Lipschitz on an interval containing both spectra.

3.4 Pair removal

Lemma 3.6 (Pair removal). *Let $G \succeq 0$ and Q be real symmetric of the same size, with $n_+(Q) \leq b$ and $n_-(Q) \leq p$. Then*

$$\|G + Q\|_{\text{HS}}^2 - \|G\|_{\text{HS}}^2 - 2 \operatorname{tr} Q \geq 2(\operatorname{tr} Q - 2b) - \sum_{i \leq p} (\lambda_i(G) - 2)_+^2.$$

Proof. Write $Q = Q_+ - Q_-$ with $Q_{\pm} \succeq 0$, $Q_+ Q_- = 0$, $\tau_+ := \operatorname{tr} Q_+$, $\tau_- := \operatorname{tr} Q_-$. Since $\operatorname{tr}(G Q_+) \geq 0$ and $\operatorname{rank} Q_+ \leq b$, $\|G + Q\|_{\text{HS}}^2 \geq \|G\|_{\text{HS}}^2 - 2 \operatorname{tr}(G Q_-) + \|Q_-\|_{\text{HS}}^2 + \tau_+^2/b$. For every real τ_+ , $\tau_+^2/b - 2\tau_+ = 2(\tau_+ - 2b) + (\tau_+ - 2b)^2/b \geq 2(\tau_+ - 2b)$ (if $b = 0$ then $Q_+ = 0$ and this step is void), and $\tau_+ = \operatorname{tr} Q + \tau_-$. So the left side is at least $2(\operatorname{tr} Q - 2b) + 4\tau_- + \|Q_-\|_{\text{HS}}^2 - 2 \operatorname{tr}(G Q_-)$. By von Neumann, $\operatorname{tr}(G Q_-) \leq \sum_{i \leq p} \lambda_i(G) n_i$, where $n_1 \geq \dots \geq n_p \geq 0$ are the eigenvalues of Q_- , and $n^2 + 4n - 2\lambda n \geq -(\lambda - 2)_+^2$ for all $n \geq 0$. $\square$

The inequality is sharp: $G = \operatorname{diag}(3.5, 0, 0)$, $Q = \operatorname{diag}(-1.5, 2, 0)$.

3.5 Localisation of the large eigenvalues

Fix $c_* \geq 0$ and put $\lambda_c := 2 + \sqrt{c_*}$,

$$\kappa(\lambda) := [(\lambda - 2)_+^2 - c_*]_+, \quad \operatorname{ch}(X) := \sum_i \kappa(\lambda_i(X))$$

for real symmetric X . The *charge* $\operatorname{ch}(X)$ sees only the eigenvalues above λ_c ; κ is nondecreasing.

Theorem 3.7 (Localisation). *Let $M = \begin{pmatrix} A & Y \\ Y^\top & R \end{pmatrix}$ be real symmetric, $A \in \mathbb{R}^{n \times n}$, $R \in \mathbb{R}^{r \times r}$, with $\lambda_{\max}(R) \leq \lambda_c$ (no condition if $r = 0$). Then*

$$\operatorname{ch}(M) \leq \|A - 2I\|_{\text{HS}}^2 + 2\|Y\|_{\text{HS}}^2 - \operatorname{tr}(2I - A)_+^2.$$

Proof. Let $\widetilde{M} := \begin{pmatrix} A & Y \\ Y^\top & \lambda_c I_r \end{pmatrix}$. Then $\widetilde{M} - M = 0 \oplus (\lambda_c I - R) \succeq 0$; this is the only use of the hypothesis. By Weyl, $\lambda_l(M) \leq \lambda_l(\widetilde{M}) =: \tilde{\mu}_l$ for every l , and since κ is nondecreasing, $\operatorname{ch}(M) \leq \operatorname{ch}(\widetilde{M})$. By Cauchy interlacing in counting form: the subspace $0 \oplus \mathbb{R}^r$, on which the form of $\widetilde{M}$ is $\lambda_c \|x\|^2$, shows that $\widetilde{M}$ has at least r eigenvalues $\geq \lambda_c$, so $\tilde{\mu}_l \geq \lambda_c$ for $l \leq r$; and the number of eigenvalues of $\widetilde{M}$ above any θ is at most that of A plus r , so $\tilde{\mu}_{r+j} \leq \alpha_j := \lambda_j(A)$ for $1 \leq j \leq n$. Let $h(\mu) := (\mu - 2)^2 - \kappa(\mu) \geq 0$. Then $h(\tilde{\mu}_l) = c_*$ for $l \leq r$, and $h(\tilde{\mu}_{r+j}) \geq (2 - \alpha_j)_+^2$ (if $\tilde{\mu}_{r+j} \leq 2$ then $h = (2 - \tilde{\mu}_{r+j})^2$ and $\tilde{\mu}_{r+j} \leq \alpha_j$; otherwise $\alpha_j > 2$). Since $\sum_l (\tilde{\mu}_l - 2)^2 = \|\widetilde{M} - 2I\|_{\text{HS}}^2 = \|A - 2I\|_{\text{HS}}^2 + 2\|Y\|_{\text{HS}}^2 + r c_*$,

$$\operatorname{ch}(\widetilde{M}) = \sum_l (\tilde{\mu}_l - 2)^2 - \sum_l h(\tilde{\mu}_l) \leq \|A - 2I\|_{\text{HS}}^2 + 2\|Y\|_{\text{HS}}^2 - \sum_{j \leq n} (2 - \alpha_j)_+^2. \quad \square$$

No positivity is used, and no Schur complement or integral over thresholds: the domination by $\widetilde{M}$ above suffices. The hypothesis $\lambda_{\max}(R) \leq \lambda_c$ cannot be dropped.

4 The local inequality and simple zeros on the line

This section proves Theorem 1.2. The mechanism is not ours (§1.4); we give it in full because Section 5 builds on it and Corollary 1.3 uses its proof.

4.1 The local inequality

Definition 4.1 (Local inequality). Let $K \geq 2$. *Position weights* are numbers $\gamma_{s,i} \geq 0$, $1 \leq s \leq K-1$, $1 \leq i \leq K-s$, with $\sum_i \gamma_{s,i} = 2$ for every s ; *penalties* are numbers $\mu_i \geq 0$, $1 \leq i \leq K-1$, with $\nu := \sum_i \mu_i$. For an even function k with $k(0) = 1$ and $c \in \mathbb{R}$ we say that (LI) holds for k with data (γ, μ, c) if

$$F(g) := \sum_{i=1}^{K-1} \mu_i g_i + \sum_{s=1}^{K-1} \sum_{i=1}^{K-s} \gamma_{s,i} k(g_i + \dots + g_{i+s-1})^2 \geq c \quad \text{for all } g \in [0, \infty)^{K-1}. \quad (6)$$

The variables g_i are the gaps between K consecutive simple zeros on the line, in normalised units, and $\mu \cdot g$ charges the length that such a run occupies. With $\gamma_{s,i} = 2/(K-s)$ and $\mu_i = \mu$ this is the uniform functional of [1]. Since $k^2 \geq 0$, $F(g) \geq \mu_i g_i$, so (6) holds automatically outside a bounded box when every $\mu_i > 0$; this is what makes it a finite computation (§7.2). Coincident points ($g_i = 0$) belong to the domain.

Lemma 4.2 (Local to block). Suppose (LI) holds for k with data (γ, μ, c) . Let $y_1 < \dots < y_m$ be real, $m \geq K$, with gaps $g_i := y_{i+1} - y_i$. Then

$$E := \sum_{i \neq j} k(y_i - y_j)^2 \geq c(m - K + 1) - \sum_{t=1}^{m-K+1} \sum_{i=1}^{K-1} \mu_i g_{t+i-1}.$$

Proof. Apply (6) to each run $y_t, \dots, y_{t+K-1}$, $1 \leq t \leq m - K + 1$, and sum. A pair at index distance $s < K$ receives total weight at most $\sum_i \gamma_{s,i} = 2$, which is its contribution to E (two ordered pairs); pairs at index distance $\geq K$ are dropped, since $k^2 \geq 0$. $\square$

Proposition 4.3 (Packing). Let k be even and positive definite with $k(0) = 1$, and suppose (LI) holds for k with data (γ, μ, c) . Let $y_1 < \dots < y_S$ lie in an interval of length $\Lambda \geq 0$ and $G := (k(y_i - y_j))_{i,j}$. For $m \geq K$ put

$$A := c(m - K + 1), \quad B := \Phi_m(A), \quad \tau := \frac{\nu(m - K + 1)}{m}. \quad (7)$$

If $A > 0$, then $\text{tr } \Psi(G) \geq \frac{B}{m} S - 2B - \frac{B}{A} \tau \Lambda$.

The hypothesis $\Lambda \geq 0$ matters only when $S = 0$; without it the statement is false ($K = 2$, $k \equiv 1$, $c = 2$, $m = 2$, $S = 0$, $\Lambda = -10$).

Proof. $G \succeq 0$ has unit diagonal. Fix an offset $o \in \{0, \dots, m-1\}$ and cut the ordered points into consecutive blocks of m , starting after the first o points; there are at least $S/m - 2$ full blocks. By Lemma 3.4 (discarding the partial blocks, for which $\text{tr } \Psi \geq 0$) and Lemma 3.3, $\text{tr } \Psi(G) \geq \sum_{\text{full } b} \Phi_m(E_b)$. By Lemma 4.2, $E_b \geq A - P_b$ with $P_b \geq 0$ the penalty sum of the runs inside b , and $\Phi_m(E_b) \geq \frac{B}{A} \min(E_b, A) \geq B - \frac{B}{A} P_b$. Hence $\text{tr } \Psi(G) \geq B(S/m - 2) - \frac{B}{A} \sum_b P_b$ for every o . Now average over o . A given gap occupies each relative position $j \in \{0, \dots, m-1\}$ in its block for exactly one offset ($j = 0$ meaning that it straddles two blocks, where it is not charged); at position $j \geq 1$ it is charged $\sum_t \mu_{j-t+1}$ over the runs t of the block that contain it, and summing over j gives $\sum_{t=1}^{m-K+1} \sum_{i=1}^{K-1} \mu_i = (m - K + 1)\nu$. So the average of $\sum_b P_b$ is at most $\tau \sum_i g_i \leq \tau \Lambda$. $\square$

4.2 From the zeros to the local inequality

Let on_1 be the set of simple on-line zeros with $\gamma \in I'$, $s_1 := \#\text{on}_1$, and let $M^\circ := (\langle u_\rho, u_{\rho'} \rangle)_{\rho, \rho' \in \text{on}_1}$ be their Gram matrix.

Lemma 4.4 (Transfer). *Let ψ be a window and suppose (LI) holds for k_ψ with data (γ, μ, c) , $c > 0$. Let $0 < c' < c$ and let A', B' be as in (7) with c' in place of c . Then for $T \geq T_0(c')$,*

$$\text{tr } \Psi(M^\circ) \geq \frac{B'}{m} s_1 - \frac{B'}{A'} \tau \Lambda - o(N).$$

Proof. By Lemma 2.2(i), $M^\circ = K^\circ - \widehat{E}^\circ$ with $K^\circ := (k_\varphi(x_\rho - x_{\rho'}))$ and $\widehat{E}^\circ \succeq 0$. Both M° and K° are positive semidefinite, Ψ is 2-Lipschitz on $[0, \infty)$, and $\text{tr } \widehat{E}^\circ = \sum \delta_\rho = o(N)$, so Lemma 3.5 gives $|\text{tr } \Psi(M^\circ) - \text{tr } \Psi(K^\circ)| = o(N)$. Next $|k_\varphi^2 - k_\psi^2| \leq 2\varepsilon_T$ by Lemma 2.2(ii), so the functional (6) for k_φ differs from that for k_ψ by at most $2\varepsilon_T \sum_{s,i} \gamma_{s,i} = 4(K-1)\varepsilon_T$, and (LI) holds for k_φ with c' once T is large. The normalised ordinates of on_1 are distinct and lie in an interval of length Λ ; apply Proposition 4.3 to K° . $\square$

Theorem 4.5 (Stability plus packing). *Let ψ be a window and suppose (LI) holds for k_ψ with data (γ, μ, c) , $c > 0$. Let $m \geq K$ with $1 - B/m > 0$ (automatic when $A < m$, since $B \leq A$). Then*

$$\liminf_{T \rightarrow \infty} \frac{N_0^s(T, 2T)}{N(T, 2T)} \geq \frac{H(\psi) - (B/A) \tau}{1 - B/m}.$$

Proof. Fix $\lambda \in (0, 1)$. Apply Lemma 3.1 with $V := [u_\rho]_{\rho \in \text{on}_1}$ (real, since γ_ρ is real), $P := VV^\top$, $r = s_1$, $M = V^\top V = M^\circ$, $Q := G - P$ and $b := s_2 + p$, where s_2 is the number of distinct multiple on-line zeros and p the number of off-line pairs with $\text{Re } \gamma \in I'$: each multiple on-line zero contributes a rank-one positive semidefinite term and each pair at most one positive eigenvalue (AF2), so $n_+(Q) \leq b$ by subadditivity of n_+ . Since $\text{tr } P + 2b \leq s_1 + 2s_2 + 2p \leq N(I')$ by (AF1),

$$2 \text{tr } P + 4 \text{tr } Q - 4b = 4 \text{tr } G - 2(\text{tr } P + 2b) \geq 4 \text{tr } G - 2N(I'),$$

and (AF3)–(AF4) give $s_1 \geq (H_\lambda(\psi) - o(1))N + \text{tr } \Psi(M^\circ)$. Insert Lemma 4.4 with $\Lambda \leq N(1 + o(1))$: $s_1(1 - B'/m) \geq (H_\lambda(\psi) - (B'/A')\tau)N - o(N)$. Since $N_0^s(T, 2T) \geq s_1 - O(\sqrt{T} \log T)$, the bound holds with H_λ and c' ; let $c' \uparrow c$ and $\lambda \rightarrow 1^-$ (§2.5); the right side is continuous in both. $\square$

The quantity $\text{tr } \Psi(M^\circ)$ is the Gram energy among the simple on-line zeros that the argument of [2] discards: if every eigenvalue of M° is at most 2, then $\text{tr } \Psi(M^\circ) = \|M^\circ - I\|_{\text{HS}}^2$ up to the diagonal defects δ_ρ . For $\psi = \cos(\alpha s)$ with $0 < \alpha < \pi$ the positive zeros of k_ψ form a sum-free set (by (3) each of them satisfies $x \tan(\pi x) = c$ with $c = \frac{\alpha}{2\pi} \tan \frac{\alpha}{2} > 0$, and if x, y and $x+y$ were all zeros, the addition formula for the tangent would give $x^2 + xy + y^2 + c^2 = 0$), so three simple zeros are never pairwise orthogonal and the energy cannot vanish locally (for the Montgomery–Taylor window this observation, with its proof, was published by ainta [1] on 11 August 2026, and it is the basis of Devine’s computation-free improvement of the Montgomery–Taylor constant [9]); it can vanish for the flat window, whose kernel vanishes on $\mathbb{Z} \setminus \{0\}$. No lower bound for this energy can come from the prime side alone: the periodic configuration with $\frac{4}{6}$ of the sites simple, $\frac{1}{6}$ double and $\frac{1}{6}$ empty [2, §7.2] has the flat-window trace moments and zero energy. The bound is therefore configuration-wise, which is why Theorem 4.5 solves an implicit linear inequality for s_1 .

4.3 Proof of Theorem 1.2

Let

$$\tilde{\psi}(s) := \frac{50000}{50037} \left(1 + \frac{127}{2500} u - \frac{8997}{10000} u^2 + \frac{6187}{10000} u^3 + \frac{3}{1250} u^4 \right), \quad u = (2s)^2,$$

on $[-\frac{1}{2}, \frac{1}{2}]$. It is even and C^∞ , with $0.7679 \leq \tilde{\psi} \leq 1$; it is not monotone on $[0, \frac{1}{2}]$ (the bracket has maximum 1.000731 at $u = 0.0291$, which is why it is normalised by $\frac{50000}{50037}$, and minimum 0.768517 at $u = 0.9357$). Its window constant is the exact rational

$$H(\tilde{\psi}) = \frac{26957030199857}{40103091391077} = 0.6721933213820763\dots,$$

with $\int \tilde{\psi}$, $\int \tilde{\psi}^2$ and $J(\tilde{\psi})$ exact rationals (J through the kernel $\int |s - s'| \tilde{\psi}(s') ds'$, a polynomial of degree 10 in s on $|s| \leq \frac{1}{2}$); this is $2.5 \cdot 10^{-4}$ below $H(\cos 1.48s)$, and the gain comes from the packing term. Take $K = 8$, the exact rational position weights and penalties of the certificate file named in Table 3, with $\nu = \frac{7}{1700}$, and $c = \frac{199}{25000} = 0.00796$. Then (LI) holds for $k_{\tilde{\psi}}$ (Section 7). Take $m = 147$: then $A = 0.00796 \cdot 140 = \frac{1393}{1250} = 1.1144 > \frac{147}{146}$, so $B = \Phi_{147}(A) = A - (\sqrt{14527/13125} - 1)^2$ lies on the concave branch of Lemma 3.3, and $\tau = \frac{7}{1700} \cdot \frac{140}{147} = \frac{1}{255}$. Theorem 4.5 gives

$$\frac{H(\tilde{\psi}) - (B/A)\tau}{1 - B/m} = 0.673373689541190\dots > 0.6733736895.$$

In the formal proof the square root is enclosed in an interval of width 10^{-18} , and the final inequality, linear in the enclosed quantity, is checked in exact rational arithmetic; the margin is $4.1 \cdot 10^{-11}$. Since every quantity is homogeneous of degree 0 in the window, the formal proof may and does use the window $\frac{4}{5}v$, with v the bracket above: it satisfies the side conditions of the window class of the Lean artifact of [2] (which $\tilde{\psi}$, being non-monotone with total variation on $\mathbb{R}$, including the jumps at $\pm \frac{1}{2}$, of about 2.016, does not; for $\frac{4}{5}v$ it is about 1.614), and it has the same H , the same kernel and the same certificate. $\square$

For distinct zeros the same chain gives $\frac{1}{2}(1 + 0.6733736895) = 0.8366868$ (Remark 6.2), superseded by Theorem 1.1. Earlier certificates of the same kind gave 0.6732063 (window $\cos 1.48s$, $K = 7$) and 0.6733246847 (window $\tilde{\psi}$, $K = 7$).

Remark 4.6 (Credit). For this statistic the stability inequality with the term $\text{tr } \Psi(M)$, the K -point functional with uniform weights $\gamma_{s,i} = 2/(K - s)$ and the offset-averaged aggregation were published by ainta [1], position weights with span capacity 2 and non-cosine windows by trmdy [24], both on 11 August 2026, position-dependent penalties by sxuff [20] (11 August), the exact count of penalties over shifted blocks (our τ ; the identification is ours [M]) in AMTOPA's work [3] and in trmdy (12 August), eight-point certificates by AMTOPA (12 August), and the aggregation function Φ_m by tawanerguo-cn [21] and sxuff (11 August) and trmdy (12 August), all before our work on it began (§1.4). What is ours here, as far as we have found, is the polynomial window with exact rational constants as a window for ζ and the box bound LB3 of §7.2, which we have not compared in substance with the pruning of other verifiers; Lemma 3.3 was published earlier by tawanerguo-cn, sxuff and trmdy (§1.4). Several unrefered works claim higher values (§1.4), among them a Lean theorem announced by Zhaoyi-Tian [32].

5 Simple and distinct zeros

Throughout this section $\psi(s) = \cos(1.6s)$, a window in the sense of §2.1 ($\cos 0.8 \leq \psi \leq 1$), with $H = H(\psi) = 0.67198155100037074\dots$ and $v_\psi(s) = \cos(8s/5)/(\frac{5}{4} \sin \frac{4}{5})$. The kernel $k = k_\psi$ is given by (3) with $\alpha = 1.6$.

5.1 The all-marks local inequality

Definition 5.1 (All-marks local inequality). Let $K \geq 2$, let $\gamma_{s,i}$ and $\mu_i \geq 0$ be position weights and penalties as in Definition 4.1, $\nu := \sum_i \mu_i$, and let $b_i(1), b_i(2) \in \mathbb{R}$ ($1 \leq i \leq K$) be *site credits*, with no sign condition; put $a_j := \sum_i b_i(j)$. For a *mark pattern* $\mathbf{m} \in \{1, 2\}^K$ and $g \in [0, \infty)^{K-1}$ let

$$F_{\mathbf{m}}(g) := \sum_{s=1}^{K-1} \sum_{i=1}^{K-s} \gamma_{s,i} m_i m_{i+s} k(g_i + \cdots + g_{i+s-1})^2 + \sum_{i=1}^{K-1} \mu_i g_i.$$

We say that $(\text{LI}_{\mathbf{m}})$ holds for k with data (γ, μ, b) if $F_{\mathbf{m}}(g) \geq C(\mathbf{m}) := \sum_i b_i(m_i)$ for every $\mathbf{m}$ and every $g \in [0, \infty)^{K-1}$.

The marks are multiplicities capped at 2. With $\mathbf{m} \equiv \mathbf{1}$ this is Definition 4.1 with $c = a_1$.

5.2 The slack of an on-line set

Let $\mathcal{O}$ be a finite set of distinct on-line zeros with $\gamma \in I'$ and multiplicities m_ρ . Put $d_\rho := \|u_\rho\|^2 = 1 - \delta_\rho$,

$$G_{\mathcal{O}} := \sum_{\rho \in \mathcal{O}} m_\rho u_\rho u_\rho^\top, \quad M_{\mathcal{O}} := (\sqrt{m_\rho m_{\rho'}} \langle u_\rho, u_{\rho'} \rangle)_{\rho, \rho' \in \mathcal{O}}, \quad E(\mathcal{O}) := \sum_{\rho \neq \rho'} (M_{\mathcal{O}})_{\rho \rho'}^2$$

(ordered pairs), and $\text{Slack}(\mathcal{O}) := s_1(\mathcal{O}) - 2 \text{tr } G_{\mathcal{O}} + \|G_{\mathcal{O}}\|_{\text{HS}}^2$, where $s_j(\mathcal{O})$ is the number of sites of multiplicity j . The matrices $G_{\mathcal{O}}$ and $M_{\mathcal{O}}$ have the same nonzero spectrum.

Lemma 5.2 (On-line slack identity). *Exactly,*

$$\begin{aligned} \text{Slack}(\mathcal{O}) &= \sum_{\rho \in \mathcal{O}} \beta_\rho + E(\mathcal{O}), \\ \beta_\rho &:= \mathbf{1}_{m_\rho=1} - 2m_\rho d_\rho + m_\rho^2 d_\rho^2 \geq m_\rho(m_\rho - 2) \mathbf{1}_{m_\rho \geq 3} - 2m_\rho(m_\rho - 1) \delta_\rho. \end{aligned}$$

Proof. $\text{tr } G_{\mathcal{O}} = \sum m_\rho d_\rho$ and $\|G_{\mathcal{O}}\|_{\text{HS}}^2 = \|M_{\mathcal{O}}\|_{\text{HS}}^2 = \sum_\rho m_\rho^2 d_\rho^2 + E(\mathcal{O})$. For $m = 1$, $\beta = \delta^2 \geq 0$. For $m \geq 2$, $\beta = f(d)$ with $f(x) = m^2 x^2 - 2mx$, and $f(1) - f(d) = m(1-d)(m(1+d) - 2) \leq 2m(m-1)\delta$. $\square$

With every $d_\rho = 1$ and all zeros on the line, this says that the slack of the rank–trace inequality is exactly the Gram energy among all on-line zeros plus $m(m-2)$ for each zero of multiplicity $m \geq 3$; the stability term of Lemma 3.1 is the part of this energy among simple zeros.

Proposition 5.3 (Aggregation). *Assume $(\text{LI}_{\mathbf{m}})$ holds for k_ψ with data (γ, μ, b) . Let $\mathcal{O}$ be as above; its Gram entries $\langle u_\rho, u_{\rho'} \rangle$ are symmetric in ρ, ρ' . Let $s_{\geq 2}(\mathcal{O})$ be the number of its sites with $m_\rho \geq 2$ and $B_b := \sum_i (|b_i(1)| + |b_i(2)|)$. Then*

$$E(\mathcal{O}) \geq a_1 s_1(\mathcal{O}) + a_2 s_{\geq 2}(\mathcal{O}) - \nu \Lambda - 16K^2 \left(\varepsilon_T |\mathcal{O}| + \sum_{\rho \in \mathcal{O}} \delta_\rho \right) - 2KB_b.$$

Proof. Order $\mathcal{O}$ by ordinate, cap the marks, $\tilde{m} := \min(m, 2)$, and let W run over the runs of K consecutive sites, with gap vector g_W . Put $F^\varphi(W) := \sum_{s,i} \gamma_{s,i} \tilde{m}_{W_i} \tilde{m}_{W_{i+s}} \langle u_{W_i}, u_{W_{i+s}} \rangle^2 + \sum_i \mu_i g_{W,i}$, and let $F^\psi(W) = F_{\tilde{\mathbf{m}}_W}(g_W)$ be the same with $k_\psi(x_{W_i} - x_{W_{i+s}})^2$ in place of $\langle u, u' \rangle^2$. *Upper bound.* A pair of sites at index distance $s < K$ occurs in the runs containing it with total weight at most $\sum_i \gamma_{s,i} = 2$, and $2\tilde{m}\tilde{m}' \langle u, u' \rangle^2 \leq mm'(\langle u, u' \rangle^2 + \langle u', u \rangle^2)$ is the contribution of its two ordered pairs to $E(\mathcal{O})$ — this is where the symmetry is used; each gap receives at most

ν , and the gaps add up to at most Λ . So $\sum_W F^\varphi(W) \leq E(\mathcal{O}) + \nu\Lambda$. *Transfer.* By Lemma 2.2, $|\langle u, u' \rangle - k_\psi| \leq \varepsilon_T + (\delta\delta')^{1/2}$ and both are at most 1 in absolute value, so $|F^\varphi(W) - F^\psi(W)| \leq 8 \sum_{s,i} \gamma_{s,i}(\varepsilon_T + \frac{1}{2}(\delta_{W_i} + \delta_{W_{i+s}}))$. A site lies in at most K runs, and $\sum_{s,i} \gamma_{s,i} = 2(K-1)$; hence $\sum_W |F^\varphi - F^\psi| \leq 16K^2(\varepsilon_T|\mathcal{O}| + \sum \delta_\rho)$. *Lower bound.* $F^\psi(W) \geq C(\tilde{\mathbf{m}}_W)$ for every run, and $\sum_W C(\tilde{\mathbf{m}}_W) = a_1s_1 + a_2s_{\geq 2}$ up to the first and last $K-1$ sites, which miss positions; the defect is at most $2KB_b$. $\square$

The symmetry cannot be dropped: for a general matrix C in place of the Gram entries, and a kernel that is not even, the inequality fails (with $K=2$, $\gamma_{1,1}=2$, $\mu_1=1$, $b \equiv 1$, the kernel $k(t) = \sqrt{1-t/2}$ on $[0,2]$ and 0 elsewhere, including $t < 0$, and 100 simple sites at spacing 1.5 with entries $C_{pq} = k(x_p - x_q)$, the right side of the inequality is $\frac{71}{2}$ while $E = \frac{99}{4}$; with C symmetrised, $E = \frac{99}{2}$).

5.3 Removing off-line pairs and heavy zeros

Split the zeros with $\operatorname{Re} \gamma \in I'$ into the *light* on-line zeros $\mathcal{O}_L$ (multiplicity 1 or 2; s_1 and s_2 of them), with $G_L := G_{\mathcal{O}_L}$ and $M_L := M_{\mathcal{O}_L}$; the *heavy* on-line zeros (multiplicity ≥ 3 ; s_H distinct ones, N_H with multiplicity), with $G_H := \sum m_\rho u_\rho u_\rho^\top$; and the off-line pairs (p of them, p_1 simple, N_{off} zeros with multiplicity), with $G_{\text{off}} := G - G_L - G_H$. Put $Q := G_H + G_{\text{off}}$, $b_Q := s_H + p$ and $N_Q := N_H + N_{\text{off}} \geq 3s_H + 2p$, so that $N(I') = s_1 + 2s_2 + N_Q$.

Lemma 5.4 (Pair and heavy removal). $n_+(Q) \leq b_Q$ and $n_-(Q) \leq p$. For a count X of zeros with $\operatorname{Re} \gamma \in I'$ let X_L be its part carried by $\mathcal{O}_L$, $\text{Slack}_X := X - 2\operatorname{tr} G + \|G\|_{\text{HS}}^2$ and $\text{Slack}_X^L := X_L - 2\operatorname{tr} G_L + \|G_L\|_{\text{HS}}^2$. Then

$$\begin{aligned} \text{Slack}_X &\geq \text{Slack}_X^L + (X - X_L) + 2(N_Q - 2b_Q) - 2E_T - \sum_{i \leq p} (\lambda_i(M_L) - 2)_+^2, \\ E_T &:= (N(I') - \operatorname{tr} G)_+ = o(N). \end{aligned}$$

Proof. Each pair block has $n_\pm \leq 1$ (AF2), $G_H \succeq 0$ has rank at most s_H , and $n_\pm$ are subadditive; G_H adds nothing to n_- . By (AF1), $\operatorname{tr} G_L \leq s_1 + 2s_2$, so $\operatorname{tr} Q = \operatorname{tr} G - \operatorname{tr} G_L \geq N_Q - E_T$, and $E_T = o(N)$ by (AF3). Apply Lemma 3.6 with $G = G_L$, whose nonzero eigenvalues are those of M_L ; its right side is increasing in $\operatorname{tr} Q$. $\square$

Only (AF1), (AF3) and the inertia of the pair blocks enter; no pair vector is ever evaluated, so no estimate at complex arguments is needed.

5.4 The on-line spectral condition

Given claims a_1, a_2 with $a_1 < 1$, let $c_* := 2 - 2a_1$, $\lambda_c := 2 + \sqrt{c_*}$, and let κ , ch be as in Theorem 3.7. Put

$$\text{mg}_L := \text{Slack}(\mathcal{O}_L) - a_1s_1 - a_2s_2 + \nu\Lambda.$$

The *on-line spectral condition* is the statement $\text{ch}(M_L) \leq \text{mg}_L + o(N)$. It is proved in Theorem 5.7 from three ingredients: room on separated sets, tight chains, and aggregation. The device of Lemmas 5.5 and 5.6, close pairs paid for locally and a separated set whose Gram norm is bounded by a Fourier majorant, was published earlier, for simple zeros on the line and with a sinc-square majorant, by gfreund123 in pull request #888 to GettysburgResearch/riemann (14 September 2026) [13] [V~]; the priority is theirs.

Lemma 5.5 (Separated sites have room). *Let $\mathcal{B}$ be an even, integrable function on $\mathbb{R}$ with $\mathcal{B} \geq 0$ on $\mathbb{R}$, $\mathcal{B} \geq v_\psi$ on $[-\frac{1}{2}, \frac{1}{2}]$, and $\widehat{\mathcal{B}}(t) := \int \mathcal{B}(s) \cos(2\pi ts) ds = 0$ for $|t| \geq 1$; let $\beta_* \geq \sup_{3/4 \leq t \leq 1} |\widehat{\mathcal{B}}(t)|$. If $\mathcal{R} \subset \mathcal{O}_L$ has normalised ordinates pairwise at least $\frac{3}{4}$ apart, then $\lambda_{\max}(M_{\mathcal{R}}) \leq (1 + \varepsilon'_T)(2\widehat{\mathcal{B}}(0) + 4\beta_*)$.*

Proof. For $y \in \mathbb{R}^{\mathcal{R}}$, by Lemma 2.2(i)–(ii) (dropping $-\widehat{E} \preceq 0$),

$$\begin{aligned} \langle y, M_{\mathcal{R}} y \rangle &\leq \int v_{\varphi} \left| \sum_a y_a \sqrt{m_a} e(x_a s) \right|^2 ds \\ &\leq (1 + \varepsilon'_T) \int_{\mathbb{R}} \mathcal{B} \left| \sum_a y_a \sqrt{m_a} e(x_a s) \right|^2 ds = (1 + \varepsilon'_T) \langle y, \mathcal{N} y \rangle, \end{aligned}$$

with $e(u) = e^{2\pi i u}$ and $\mathcal{N}_{ab} := \sqrt{m_a m_b} \widehat{\mathcal{B}}(x_a - x_b)$. Since $\widehat{\mathcal{B}}$ vanishes on $|t| \geq 1$ and the sites are pairwise $\geq \frac{3}{4}$ apart, each site has at most one other site within distance < 1 on each side, at distance in $[\frac{3}{4}, 1)$. So each row of $\mathcal{N}$ has diagonal entry $\leq 2\widehat{\mathcal{B}}(0)$ and at most two off-diagonal entries, each of absolute value $\leq 2\beta_*$; Gershgorin gives $\lambda_{\max}(\mathcal{N}) \leq 2\widehat{\mathcal{B}}(0) + 4\beta_*$. $\square$

The majorant. Let $\beta_0, \dots, \beta_{59}$ be sixty binary64 numbers, read as exact rationals (the certificate file of Table 3), and

$$\mathcal{B}(s) := \frac{1}{60} \operatorname{sinc}(\pi s/60)^2 P(s), \quad P(s) := \beta_0 + 2 \sum_{j=1}^{59} \beta_j \cos(\pi j s/30).$$

From $\int \operatorname{sinc}(\pi u)^2 \cos(2\pi \omega u) du = (1 - |\omega|)_+$ one gets $\widehat{\mathcal{B}}(t) = \sum_{|j| \leq 59} \beta_{|j|} (1 - |60t - j|)_+$, the piecewise-linear interpolant of the β_j at the nodes $j/60$, which vanishes for $|t| \geq 1$; so $\int \mathcal{B} = \widehat{\mathcal{B}}(0) = \beta_0 = 1.5556524027186915$, and, since $\frac{3}{4} = \frac{45}{60}$, $\sup_{[3/4, 1]} |\widehat{\mathcal{B}}| = \max_{45 \leq j \leq 59} |\beta_j| = |\beta_{52}| = 0.0237647699514 \dots =: \beta_*$. In exact rational arithmetic $2\beta_0 + 4\beta_* = 3.2063638852 \dots < \frac{33}{10}$. The two remaining conditions, $P \geq 0$ on $[0, 30]$ (P is even and 60-periodic; its minimum is about $1.12 \cdot 10^{-3}$, near $s = 3.50$) and $\mathcal{B} \geq v_{\psi}$ on $[0, \frac{1}{2}]$ (margin about $1.95 \cdot 10^{-5}$, near $s = 0.4994$), are finite computations, certified in Section 7. Hence, as soon as $\varepsilon'_T \leq \frac{1}{40}$,

$$\lambda_{\max}(M_{\mathcal{R}}) \leq \frac{41}{40} \cdot \frac{33}{10} < 3.4 \leq \lambda_c \quad (8)$$

for the claims used below ($c_* \geq 1.96$).

Lemma 5.6 (Tight chains pay for themselves). *Let $\mathcal{T} \subset \mathcal{O}_L$ be a union of n_{ch} maximal tight chains: maximal runs of consecutive sites with consecutive gaps $< \frac{3}{4}$, each with at least two sites. Let $k_* \geq 0$ with $k_{\varphi}(t) \geq k_*$ for $0 \leq t < \frac{3}{4}$, and put*

$$\varpi(k) := \min \left\{ 2k^2 - 2a_1, 4k^2 - 2a_2, \left(\frac{1}{2} + \left(\frac{1}{4} + 2k^2 \right)^{1/2} \right)^2 - 1 - a_1 - a_2 \right\}.$$

Assume $a_1, a_2 \geq 0$ and $\varpi(k_*) \geq 0$. Then

$$\Delta_{\mathcal{T}} := \operatorname{tr}(2I - M_{\mathcal{T}})_+^2 - (1 + a_1)s_1(\mathcal{T}) - a_2s_2(\mathcal{T}) \geq n_{\text{ch}}(\varpi(k_*) - \max(a_1, a_2)).$$

Proof. $f(x) := (2 - x)_+^2$ is convex and nonincreasing. Split each chain greedily into consecutive pairs and at most one singleton. By Lemma 3.4, $\operatorname{tr} f(M_{\mathcal{T}}) \geq \sum_w \operatorname{tr} f(M_w)$ over the blocks. By Lemma 2.2(i), $M_w = D_w^{1/2} (K_w^{\varphi} - \widehat{E}_w) D_w^{1/2} \preceq M'_w := D_w^{1/2} K_w^{\varphi} D_w^{1/2}$ with $D_w = \operatorname{diag}(m)$ and K_w^{φ} the kernel matrix; as f is nonincreasing, Weyl gives $\operatorname{tr} f(M_w) \geq \operatorname{tr} f(M'_w)$. For a pair at gap $< \frac{3}{4}$ with $k := k_{\varphi}(\text{gap}) \in [k_*, 1]$, the eigenvalues of M'_w are $1 \pm k$, $2 \pm 2k$, $\frac{3}{2} \pm (\frac{1}{4} + 2k^2)^{1/2}$ for marks $(1, 1)$, $(2, 2)$, $(1, 2)$, so $\operatorname{tr} f(M'_w)$ minus the claims of the block is $2k^2 - 2a_1$, $4k^2 - 2a_2$, $(\frac{1}{2} + (\frac{1}{4} + 2k^2)^{1/2})^2 - 1 - a_1 - a_2$ respectively, each increasing in k^2 , hence $\geq \varpi(k_*)$. A singleton costs at least $-a_1$ (simple: $f(d) \geq 1$) or $-a_2$ (double: $f(2d) \geq 0$). A chain with $q \geq 1$ pairs therefore contributes at least $q\varpi(k_*) - \max(a_1, a_2) \geq \varpi(k_*) - \max(a_1, a_2)$ if it has a singleton (using $\varpi(k_*) \geq 0$), and $q\varpi(k_*) \geq \varpi(k_*) - \max(a_1, a_2)$ if not (using $a_1, a_2 \geq 0$). $\square$

Both hypotheses are needed. Take $k_* = 0$, the kernel equal to 1 at 0 and 0 elsewhere, and $\widehat{E} = 0$. With the $K = 7$ claims (so $a_1, a_2 > 0$ but $\varpi(0) < 0$) and one chain of six simple sites at gaps $\frac{1}{2}$, the left side is $-6a_1 = -0.10949\dots$ and the right side $-3a_2 = -0.07008\dots$. With $a_1 = a_2 = -\frac{1}{10}$ (so $\varpi(0) = \frac{1}{5} \geq 0$) and one chain of two simple sites, the left side is $\frac{1}{5}$ and the right side $\frac{3}{10}$. Since k_ψ is nonincreasing on $[0, 1]$ ($k'_\psi(t) = -\int 2\pi\sigma v_\psi(\sigma) \sin(2\pi t\sigma) d\sigma \leq 0$ there), Lemma 2.2(ii) allows $k_* = k_\psi(\frac{3}{4}) - \varepsilon_T$, with $k_\psi(\frac{3}{4}) = 0.3556941746\dots$. At $\varepsilon_T = 0$ the three entries of ϖ are 0.2165399517, 0.4593506235 and 0.4206777378 for the $K = 7$ claims, and $\varpi - \max(a_1, a_2) = 0.1931785717$ (0.2118330717 for $K = 5$).

Theorem 5.7 (On-line spectral condition). *Assume (LI_m) for k_ψ with data (γ, μ, b) , with $a_1, a_2 \geq 0$, $\lambda_c > \frac{33}{10}$ and $\max(a_1, a_2) < \varpi(k_\psi(\frac{3}{4}))$. Let $\mathcal{T}$ be the set of light sites that have another light site within normalised distance $< \frac{3}{4}$, and $\mathcal{R} := \mathcal{O}_L \setminus \mathcal{T}$. Then, exactly,*

$$\text{mg}_L - \left[\|M_{\mathcal{T}} - 2I\|_{\text{HS}}^2 + 2\|M_{\mathcal{T}\mathcal{R}}\|_{\text{HS}}^2 - \text{tr}(2I - M_{\mathcal{T}})^2_+ \right] = \text{mg}(\mathcal{R}) + \Delta_{\mathcal{T}} - 2 \sum_{\rho \in \mathcal{T}} m_\rho \delta_\rho, \quad (9)$$

with $\text{mg}(\mathcal{R}) := \text{Slack}(\mathcal{R}) - a_1 s_1(\mathcal{R}) - a_2 s_2(\mathcal{R}) + \nu \Lambda$; and for $T \geq T_0$,

$$\text{ch}(M_L) \leq \text{mg}_L + (16K^2 + 4) \sum_{\rho \in \mathcal{O}_L} m_\rho \delta_\rho + 16K^2 \varepsilon_T |\mathcal{O}_L| + 2KB_b = \text{mg}_L + o(N).$$

Proof. Structure. Any two sites of $\mathcal{R}$ are at least $\frac{3}{4}$ apart. Every site of $\mathcal{T}$ has a site of $\mathcal{R}$ within $\frac{3}{4}$ that is adjacent to it in the order of $\mathcal{T}$ (a closer partner lies in between), so $\mathcal{T}$ is the disjoint union of its maximal tight chains. *Localisation.* By Lemma 5.5 and the majorant, $\lambda_{\max}(M_{\mathcal{R}}) \leq (1 + \varepsilon'_T) \frac{33}{10} < \lambda_c$ once ε'_T is small. After a permutation of indices, under which $\text{tr } f$ is invariant, $M_L = \begin{pmatrix} M_{\mathcal{T}} & M_{\mathcal{T}\mathcal{R}} \\ M_{\mathcal{R}\mathcal{T}} & M_{\mathcal{R}} \end{pmatrix}$, and Theorem 3.7 bounds $\text{ch}(M_L)$ by the bracket in (9).

The identity. $E(\mathcal{O}_L) = E(\mathcal{T}) + E(\mathcal{R}) + 2\|M_{\mathcal{T}\mathcal{R}}\|_{\text{HS}}^2$ (ordered pairs), the claims are additive, $\|M_{\mathcal{T}} - 2I\|_{\text{HS}}^2 = \sum_{\rho \in \mathcal{T}} (m_\rho d_\rho - 2)^2 + E(\mathcal{T})$, and $(md - 2)^2 - \beta_\rho = \mathbf{1}_{m=1} + 2m\delta$. Insert Lemma 5.2 for $\text{Slack}(\mathcal{O}_L)$ and $\text{Slack}(\mathcal{R})$. *The bound.* Take $k_* := k_\psi(\frac{3}{4}) - \varepsilon_T$ in Lemma 5.6 (valid since k_ψ is nonincreasing on $[0, 1]$); for ε_T small, $\varpi(k_*) \geq \max(a_1, a_2) \geq 0$ by the hypothesis and continuity, so $\Delta_{\mathcal{T}} \geq 0$. By Lemma 5.2 ($\beta_\rho \geq -4\delta_\rho$ for $m_\rho \leq 2$) and Proposition 5.3 applied to $\mathcal{O} = \mathcal{R}$ (removing sites only lengthens gaps, and (LI_m) holds for all $g \geq 0$), $\text{mg}(\mathcal{R}) \geq -4 \sum_{\mathcal{R}} \delta_\rho - 16K^2(\varepsilon_T |\mathcal{R}| + \sum_{\mathcal{R}} \delta_\rho) - 2KB_b$. Finally $2 \sum_{\mathcal{T}} m_\rho \delta_\rho \leq 4 \sum_{\mathcal{T}} \delta_\rho$, and the error is $o(N)$ by Lemma 2.2(iii)–(iv) and $\varepsilon_T \rightarrow 0$. $\square$

The truncation matrix $\widehat{E}$ enters only with the favourable sign and through $|\widehat{E}_{\rho\rho'}| \leq (\delta_\rho \delta_{\rho'})^{1/2}$, and the charge is computed on the true matrix, so no perturbation bound for M_L is needed. For the certificates used, the formal proof needs of the numerical constants only $k_\psi(\frac{3}{4}) \geq \frac{1}{4}$ (proved by the fundamental theorem of calculus and $3.14 < \pi < 3.15$; already $\varpi(\frac{1}{4}) > \max(a_1, a_2)$ for these claims), $\lambda_c \geq 3.4$ and $2\beta_0 + 4\beta_* \leq \frac{33}{10}$, not the more precise values quoted above.

5.5 Proof of Theorem 1.1

Theorem 5.8. *Let $\psi = \cos(1.6s)$ and suppose (LI_m) holds for k_ψ with data (γ, μ, b) whose claims a_1, a_2, ν satisfy: $a_1, a_2, \nu \geq 0$; $\lambda_c = 2 + \sqrt{2 - 2a_1} > \frac{33}{10}$; $\max(a_1, a_2) < \varpi(k_\psi(\frac{3}{4}))$; $1 - a_1 + \frac{a_2}{2} > 0$; and, for all integers $m \geq 2$ and $m' \geq 3$,*

$$\begin{aligned} (4 - a_2)m - 4 - c_* &\geq 0, & 2m(1 - a_2 + a_1) - 4a_1 + 2a_2 - c_* &\geq 0, \\ (2 - \frac{a_2}{2})m' - 4 &\geq 0, & m'(1 - a_2 + a_1) - 2 - 2a_1 + a_2 &\geq 0. \end{aligned}$$

Then

$$\liminf_{T \rightarrow \infty} \frac{N^s(T, 2T)}{N(T, 2T)} \geq \Sigma := \frac{H + \frac{a_2}{2} - \nu}{1 - a_1 + \frac{a_2}{2}},$$

$$\liminf_{T \rightarrow \infty} \frac{N^d(T, 2T)}{N(T, 2T)} \geq D := \frac{1 + H + a_2 - a_1 - \nu}{2 - 2a_1 + a_2} = \frac{1 + \Sigma}{2}.$$

Proof. Fix $\lambda \in (0, 1)$ and let X be $s_1 + 2p_1$ (the simple zeros with $\operatorname{Re} \gamma \in I'$: a simple off-line zero and its partner are both simple) or $s_1 + s_2 + s_H + 2p$ (the distinct ones); they differ from $N^s(T, 2T)$, $N^d(T, 2T)$ by $O(\sqrt{T} \log T)$.

Step 1. By definition $X = 2 \operatorname{tr} G - \|G\|_{\text{HS}}^2 + \text{Slack}_X$, so (AF3) and (AF4) give $X \geq (H_\lambda - o(1))N + \text{Slack}_X$.

Step 2. By Lemma 5.4, the inequality $(\lambda - 2)_+^2 \leq c_* + \kappa(\lambda)$, and Theorem 5.7, $\text{Slack}_X \geq (\text{Slack}_X^L - \text{Slack}(\mathcal{O}_L)) + a_1 s_1 + a_2 s_2 - \nu \Lambda + (X - X_L) + 2(N_Q - 2b_Q) - c_* p - o(N)$, where $\text{Slack}_X^L - \text{Slack}(\mathcal{O}_L) = X_L - s_1$ is 0 for simple zeros and s_2 for distinct zeros.

Step 3. For simple zeros compare the right side with $a_1 X + \frac{a_2}{2}(N(I') - X) - \nu \Lambda$; for distinct zeros with $(2a_1 - 1 - a_2)X + (1 + a_2 - a_1)N(I') - \nu \Lambda$. On the light zeros the two sides agree. Using $N(I') = s_1 + 2s_2 + N_Q$, the remaining difference (available minus required) is a sum over objects:

- a simple off-line pair: $2 - 2a_1 - c_* = 0$, for both counts;
- an off-line pair of multiplicity $m \geq 2$: $(4 - a_2)m - 4 - c_*$, resp. $2m(1 - a_2 + a_1) - 4a_1 + 2a_2 - c_*$;
- a heavy on-line zero of multiplicity $m \geq 3$: $(2 - \frac{a_2}{2})m - 4$, resp. $m(1 - a_2 + a_1) - 2 - 2a_1 + a_2$;

all nonnegative by hypothesis. So $\text{Slack}_X \geq a_1 X + \frac{a_2}{2}(N(I') - X) - \nu \Lambda - o(N)$, resp. $\text{Slack}_X \geq (2a_1 - 1 - a_2)X + (1 + a_2 - a_1)N(I') - \nu \Lambda - o(N)$.

Step 4. Insert into Step 1 with $N(I') = N + o(N)$ and $\Lambda \leq N(1 + o(1))$ (and $\nu \geq 0$): $X(1 - a_1 + \frac{a_2}{2}) \geq (H_\lambda + \frac{a_2}{2} - \nu)N - o(N)$, resp. $X(2 - 2a_1 + a_2) \geq (1 + H_\lambda + a_2 - a_1 - \nu)N - o(N)$. Divide by the positive denominators and let $\lambda \rightarrow 1^-$ (§2.5). $\square$

Proof of Theorem 1.1. The $K = 5$ and $K = 7$ certificates of Table 2 satisfy (LI_m) for k_ψ (Section 7) and the conditions of Theorem 5.8 with wide margins: $\lambda_c \geq 3.401$, $\varpi(k_\psi(\frac{3}{4})) - \max(a_1, a_2) \geq 0.193$, and every entry of the table in Step 3 for a multiple off-line pair or a heavy on-line zero is at least 0.97 (the entry for a simple off-line pair is exactly 0, by the choice of c_*). With $H = H(\cos 1.6s)$, enclosed in an interval of width $4 \cdot 10^{-21}$ (the formal proof uses an enclosure of width 10^{-16} from alternating Taylor series), the constants are those of Table 2. $\square$

Table 2: The all-marks certificates for $\psi = \cos 1.6s$ ($\nu = (K - 1)/500$). The digits of Σ and D are truncated. The $K = 6$ row is given for comparison; its certificate is not used and is not listed in Table 3.

K	a_1	a_2	ν	Σ	D	λ_c
5	$\frac{1280197}{10^8}$	$\frac{48749}{3125000}$	$\frac{1}{125}$	0.675158622199	0.837579311099	3.40513
6	$\frac{195379}{12500000}$	$\frac{976263}{50000000}$	$\frac{1}{100}$	0.675709032131	0.837854516065	3.40312
7	$\frac{1824837}{10^8}$	$\frac{1168069}{5 \cdot 10^7}$	$\frac{3}{250}$	0.676102666964	0.838051333482	3.40125

Remark 5.9 (What is not claimed). The proportion of zeros that are simple *and* on the line is not improved by this route. For it a simple off-line pair brings no bonus, the threshold in the charge would be 2 instead of λ_c , and Theorem 3.7 would need $\lambda_{\max}(M_{\mathcal{R}}) \leq 2$, which fails already for unit-spaced double zeros (two of them give $\lambda_{\max} = 2(1 + k_\psi(1)) = 2.1387$, and long runs approach $2/\int \psi = 2.2304$).

6 Zeros that are simple or on the line

Proof of Corollary 1.3. Use the window $\tilde{\psi}$ and the data of the proof of Theorem 1.2 throughout; the cancellation of R below needs the same window in (AF4) and in the packing term. Fix $\lambda \in (0, 1)$ and let $t := 2 + \sqrt{2}$, so that $t^2 = 4t - 2$ and $t \geq 2$. Let N_{on} be the number of on-line zeros with $\gamma \in I'$ counted with multiplicity, $r \leq N_{\text{on}}$ the number of distinct ones, and p_1, p_2 the numbers of simple and of multiple off-line pairs with $\text{Re } \gamma \in I'$, so that $p = p_1 + p_2$. The count to be bounded is $N^{\text{sc}}(I') := N_{\text{on}} + 2p_1$, which differs from $N^{\text{sc}}(T, 2T)$ by $O(\sqrt{T} \log T)$.

Apply Lemma 3.2 with $(s, t) = (1, 2 + \sqrt{2})$, $P := \sum_{\text{on-line}} m_\rho u_\rho u_\rho^\top = VV^\top$, where V has the r real columns $\sqrt{m_\rho} u_\rho$, $M := V^\top V$, and $Q := G - P$, the sum of the off-line pair blocks, so that $n_+(Q) \leq b := p$ by (AF2). By (AF1), $\text{tr } P \leq N_{\text{on}}$, hence $2 \text{tr } P + 2t \text{tr } Q = 2 \text{tr } G + 2(t-1) \text{tr } Q \geq 2 \text{tr } G + 2(t-1)(\text{tr } G - N_{\text{on}})$, and (AF3)–(AF4) give

$$r + t^2 p \geq (2t - R_\lambda(\tilde{\psi}))N - 2(t-1)N_{\text{on}} + \text{tr } \Psi_{1,t}(M) - o(N).$$

Since $N(I') - N_{\text{on}} \geq 2p_1 + 4p_2$, we have $p \leq \frac{1}{4}(N(I') - N_{\text{on}} + 2p_1)$; also $r \leq N_{\text{on}}$. Moving the N_{on} terms to the left, their coefficient is $2t - 1 - t^2/4 = t^2/4$, so

$$\frac{t^2}{4} N^{\text{sc}}(I') \geq (2t - R_\lambda(\tilde{\psi}) - \frac{t^2}{4})N + \text{tr } \Psi_{1,t}(M) - o(N).$$

Pinch M to the block M° of the simple on-line zeros (Lemma 3.4; $\Psi_{1,t}$ is convex) and use $\Psi_{1,t} \geq \Psi \geq 0$: $\text{tr } \Psi_{1,t}(M) \geq \text{tr } \Psi(M^\circ)$. For $c' < c$ let A', B' be as in Lemma 4.4, $\beta' := B'/m$, $\tau'' := (B'/A')\tau$ and $S_{c'} := (H_\lambda(\tilde{\psi}) - \tau'')/(1 - \beta')$. Lemma 4.4 gives $\text{tr } \Psi(M^\circ) \geq \beta' s_1 - \tau'' N - o(N)$, and the proof of Theorem 4.5 gives $s_1 \geq (S_{c'} - o(1))N$. Since $\beta' > 0$ and $\beta' S_{c'} - \tau'' = S_{c'} - H_\lambda(\tilde{\psi})$,

$$\text{tr } \Psi(M^\circ) \geq (S_{c'} - H_\lambda(\tilde{\psi}))N - o(N), \quad H_\lambda(\tilde{\psi}) = 2 - R_\lambda(\tilde{\psi}).$$

Substituting, $R_\lambda(\tilde{\psi})$ cancels:

$$N^{\text{sc}}(I') \geq \frac{8t - t^2 - 8 + 4S_{c'}}{t^2} N - o(N) = \frac{2t - 3 + 2S_{c'}}{2t - 1} N - o(N).$$

Let $T \rightarrow \infty$, then $c' \uparrow c$ and $\lambda \rightarrow 1^-$; $S_{c'}$ tends to the constant $S^* = 0.673373689541190\dots$ of Theorem 1.2, and $(1 + 2\sqrt{2} + 2S^*)/(3 + 2\sqrt{2}) = 0.887919569562\dots > 0.8879195$. $\square$

Remark 6.1. (i) The corollary follows from the *proof* of Theorem 1.2 (the linear packing bound, the fixed-point identity and the limit $c' \uparrow c$), not from the number S^* alone; it uses no input beyond those of Theorem 1.2 and Lemma 3.2.

(ii) The formula $(1 + 2\sqrt{2} + 2S)/(3 + 2\sqrt{2})$ is Lamzouri's [15], with Lamzouri's constant replaced by S^* . The value moves only through S , with slope $2/(3 + 2\sqrt{2}) = 0.343146$: 0.8876200 at $S = H(\cos \sqrt{2}s)$ (Lamzouri's value) and 0.8879195 at S^* . The whole gain, $2.9 \cdot 10^{-4}$, comes from Theorem 1.2, so Corollary 1.3 is not an independent advance. At $S = 0.67349239$, the largest unrefereed bandwidth-one value of §1.4, the formula gives 0.8879603.

(iii) The variant that counts a multiple on-line zero once is not claimed. Within the two-parameter family of Lemma 3.2, the choice $s = 1, t = 2 + \sqrt{2}$ is optimal for this argument, by a computation that we do not reproduce; nothing above depends on it.

Remark 6.2 (Distinct zeros from the chain of Theorem 1.2). The same argument gives

$$\liminf_{T \rightarrow \infty} \frac{N^d(T, 2T)}{N(T, 2T)} \geq \frac{1}{2}(1 + S^*) = 0.8366868\dots,$$

which Theorem 1.1 supersedes. We include it because it is the value that a bound of the kind of Theorem 1.2 gives for distinct zeros. The remark is not formalised. *Proof.* Fix $\lambda \in (0, 1)$ and $c' < c$, with $\beta', \tau'', S_{c'}$ as above. Apply Lemma 3.1 with V the matrix of the r real columns $\sqrt{m_\rho} u_\rho$ over the distinct on-line zeros with $\gamma_\rho \in I'$, $P := VV^\top$, $M := V^\top V$, $Q := G - P$ and $b := p$ (AF2). Let $N_{\text{off}} := N(I') - N_{\text{on}} \geq 2p$. By (AF1) and (AF3), $\text{tr } Q = \text{tr } G - \text{tr } P \geq N_{\text{off}} - o(N)$, so $2 \text{tr } P + 4 \text{tr } Q = 2 \text{tr } G + 2 \text{tr } Q$ and (AF4) give

$$r \geq H_\lambda(\tilde{\psi})N + 2N_{\text{off}} - 4p + \text{tr } \Psi(M) - o(N).$$

Pinch M into the block M° of the simple on-line zeros and the 1×1 blocks $m_\rho \|u_\rho\|^2 = m_\rho(1 - \delta_\rho)$ of the multiple ones (Lemma 3.4). Since Ψ is 2-Lipschitz on $[0, \infty)$ and $\Psi(m) = 2m - 3$ for $m \geq 2$, $\Psi(m_\rho(1 - \delta_\rho)) \geq 2m_\rho - 3 - 2m_\rho\delta_\rho$, and $\sum_\rho m_\rho\delta_\rho = o(N)$; and since $\text{tr } \Psi(M^\circ) \geq (S_{c'} - H_\lambda(\tilde{\psi}))N - o(N)$ as shown above,

$$N^d(I') := r + 2p \geq S_{c'}N + 2N_{\text{off}} - 2p + \sum_{m_\rho \geq 2} (2m_\rho - 3) - o(N),$$

the sum running over the multiple on-line zeros. Add the identity $N^d(I') = s_1 + n_{\geq 2} + 2p$, where $n_{\geq 2}$ is the number of multiple on-line zeros, and subtract $N(I') = s_1 + \sum_{m_\rho \geq 2} m_\rho + N_{\text{off}}$:

$$2N^d(I') - N(I') \geq S_{c'}N + N_{\text{off}} + \sum_{m_\rho \geq 2} (m_\rho - 2) - o(N) \geq S_{c'}N - o(N).$$

Since $N(I') = N + o(N)$ and $N^d(T, 2T) \geq N^d(I') - O(\sqrt{T} \log T)$, letting $T \rightarrow \infty$, then $c' \uparrow c$ and $\lambda \rightarrow 1^-$ gives the claim. With $\text{tr } \Psi(M^\circ) \geq 0$ in place of the packing bound the same computation gives $\frac{1}{2}(1 + H(\psi))$, the value of [2, Theorem A(ii)].

7 The certificates

7.1 What is certified

Four finite statements enter the proofs. Three are local inequalities for all gap vectors in an orthant, one is the majorant.

- **CertS8:** (LI) of Definition 4.1 for $k_{\tilde{\psi}}$, $K = 8$, with position weights and penalties summing to $\nu = \frac{7}{1700}$, at $c = \frac{199}{25000}$ (Theorem 1.2, Corollary 1.3).
- **CertAM5** and **CertAM7:** (LI_m) of Definition 5.1 for $k_{\cos 1.6s}$ with the claims of Table 2 for $K = 5$ and $K = 7$ (Theorem 1.1); a $K = 6$ certificate, shown in Table 2 for comparison, is not used.
- **CertMajV2:** the majorant of §5: $P \geq 0$ on $[0, 30]$ and $\mathcal{B} \geq v_\psi$ on $[0, \frac{1}{2}]$, with $2\beta_0 + 4\beta_* < \frac{33}{10}$. (The library also contains a first version, **CertMaj**, with the threshold 3.2063638853, only $5.7 \cdot 10^{-11}$ above the actual value; it is proved as well, and no theorem uses it.)

The data of the local inequalities (weights, penalties, credits) were chosen by linear programmes over finite sets of gap vectors found by local minimisation, and the claims were then set slightly below the programme's values: for $K = 5$ each claim lies $1.5 \cdot 10^{-5}$ below the programme's value, and for $K = 7$ every claim was lowered by 10^{-5} after the certifier had refuted the programme's own claims (§7.2). The certificates, not the programmes, carry the proof. In the formalisation each is a single **Prop** (Appendix A). The local ones are existential in the weights: they assert *some* data with the stated ν and credit sums a_1, a_2 , since the constants depend on the data only through these numbers; the files of Table 3 supply the witnesses. Gap vectors range over the whole orthant $[0, \infty)^{K-1}$; the reductions used by the certifiers (to a box, and to classes of patterns) are theorems, not part of the statements.

Table 3: Certificates, their data and their verification. Paths are relative to the folder `supplementary/certificates/` of the repository (§7.4). The seven sha256 values printed here are those of the files in `supplementary/`; those of the three local certificates are also recorded in the doc-strings of their Props (Appendix A.2).

name	data and sha256	verification
CertS8	CertS8/w.poly8A_K8_mu1700.json 257b340d9da4f343600a21589f16d891 5f3e082499d475e153c5453980671e78 the window, win.poly8A.json of the same folder (the unnormalised v): c4005f57a7920f78aa2df33b093c9c50 30efac2b8e39fc6cf7371b0c0f32dc14	interval branch-and-bound with LB3, about $4.5 \cdot 10^8$ boxes in 8 shards; re-run from scratch ($4.51 \cdot 10^8$ boxes). In Lean: displayed hypothesis.
CertAM5	CertAM5/claims.d15e6_a1.6.K5_mu500.json cf055dd2ee9a9e469b6510fe70b5fbc4 16b8a40515cc6746bc390d40d88777f2 the 20 files xpat.K5_*.json of the same folder, concatenated in sorted order: a5947c01d34ddcbc7366861b35e83942 099fcbf7ac71e13fcd432796882dd34e	20 classes; $1.18 \cdot 10^7$ boxes (LB1–LB2), $3.7 \cdot 10^5$ (LB3); re-run with identical counts. In Lean: replayed by the kernel (46,047 leaves, 89,362 point records; §7.3), so the $K = 5$ pair has no hypothesis; the replay files are the library ZetaSReplay of the repository.
CertAM7	CertAM7/claims.a1.6.K7_mu500.L1e5.json d2cf393019407bdc797c73147ec8e3dc 66ec41fe745600c84173f74ac66bfc5c the 72 files xpat.K7_*.json of the same folder, concatenated in sorted order: 33c4c8c5762bcebdf8f544ab34a137a6 a7ba4065d08f5caee0a51188e1a67514	72 classes covering 128 patterns; about $4.8 \cdot 10^8$ boxes (LB3); complete run 72/72 at these claims: logs/cert_lb3_rerun.txt of the same folder, completeness checked in logs/k7_verify_pexact2.txt (the log of the refutation described in §7.2 is not part of supplementary/). In Lean: displayed hypothesis.
CertMajV2	CertMajV2/majorant.delta1.beta.npy (60 binary64 numbers) fd909f79669e2fc1ba5fd4e44d01d600 7ccf5303b7ce62d4c5ee08f7ef024845	Arb at 106 bits. In Lean: kernel-checked, 614 cells, about 88 s; the hypothesis is discharged.

7.2 Interval branch-and-bound

Reductions. Since $k^2 \geq 0$, $F_{\mathbf{m}}(g) \geq \sum_i \mu_i g_i$, so a box on which $\sum_i \mu_i (c_i - h_i)$ already exceeds the claim is discarded, and only a bounded box needs to be covered. The weights and credits of the all-marks certificates are reversal-symmetric, so $F_{\mathbf{m}}(g)$ is unchanged when both the pattern and the gap vector are reversed; this reduces the 2^K patterns to 20, 36 and 72 classes for $K = 5, 6, 7$. Coincident points ($g_i = 0$) belong to the domain; there $k(0) = 1$ makes F large.

Kernel bounds. For a window $\psi \geq 0$, $|k_{\psi}^{(j)}| \leq (2\pi)^j \int v_{\psi} |s|^j$; for $\tilde{\psi}$ this gives $|k'''| \leq M_3 = \frac{1747823}{7775424} \pi^3 < 6.969843$. These are proved analytically and enter as constants. Values of k, k', k'' at box centres are enclosed in Arb (or in `mpmath.iv` with explicit Taylor remainders).

Lower bounds on a box $\mathbf{c} \pm \mathbf{h}$. For a span s (a run of consecutive gaps) let sc_s, hs_s be the sums of the c_i and of the h_i over the span.

- (LB1) On the span interval, $|k| \geq |k(sc)| - |k'(sc)| hs - \frac{1}{2} \sup |k''| hs^2$; clip at 0, square, sum with the weights, and add $\sum \mu_i (c_i - h_i)$.
- (LB2) $F(\mathbf{c}) - \sum_i |\partial_i F(\mathbf{c})| h_i - \frac{1}{2} \sum_s \gamma_s W_{2,s} hs_s^2$, where $W_{2,s}$ bounds $|w''| = 2|k'^2 + kk''|$ on the span interval; valid because the Hessian of a span term is $w'' \mathbf{v}_s \mathbf{v}_s^T$.

(LB3) By Taylor’s theorem with integral remainder, $F(\mathbf{c}+\boldsymbol{\delta}) \geq F(\mathbf{c})+\nabla F(\mathbf{c})\cdot\boldsymbol{\delta}+\frac{1}{2}\sum_s\gamma_s\underline{w}_s''(\mathbf{v}_s\cdot\boldsymbol{\delta})^2$ on the box, with $\underline{w}_s''$ a lower bound for w'' on the span interval from a Taylor model of k, k', k'' with $|k'''| \leq M_3$. Spans with $\underline{w}_s'' \leq 0$ are bounded below by $-\frac{1}{2}\gamma_s|\underline{w}_s''|\text{hs}_s^2$. The remaining quadratic q is convex, so $q(\boldsymbol{\delta}) \geq q(\hat{\mathbf{d}}) + \nabla q(\hat{\mathbf{d}}) \cdot (\boldsymbol{\delta} - \hat{\mathbf{d}})$ for *every* $\hat{\mathbf{d}}$ (a floating-point quadratic programme chooses $\hat{\mathbf{d}}$, which affects only tightness), and this affine function is minimised over the box exactly.

A box is accepted when a lower bound, evaluated in interval arithmetic, exceeds the claim; otherwise it is bisected. The certificate is the statement that the stack empties with every box domain-pruned or accepted. LB3 cuts the trees by factors of 30 to 200. Tangent planes after a convexity certificate are used in the verifiers of sxuff, AMTOPA and npip99 (§1.4); LB3 is our form of this step, and we have not compared it with theirs in substance.

Arithmetic and trust. Box edges are exact dyadics. The default engine uses float64 intervals widened outward by one unit in the last place at every operation, which assumes IEEE-754 round-to-nearest; a scalar engine in Arb balls avoids this assumption at about fifteen times the cost. The correctness of Arb is trusted. The closing arithmetic (H , Φ_m , the constants) was done in Arb from exact rationals and is also checked in Lean.

Validation and negative controls. Claims set just above the located minimum must fail, and do, with a rigorous counterexample: for $K = 5$, class 11111 at the claim raised by $4.5 \cdot 10^{-5}$; for $K = 6$, class 211112 at 0.0166770 ($F = 0.01667680$); for $K = 7$ the certifier refuted class 1121212 at the claims first proposed by the linear programme that chose the weights ($F = 0.02054932 < 0.02054941$), in a basin that the programme’s search had missed, after which every claim was lowered by 10^{-5} and all 72 classes certified; at the lowered claim plus 10^{-5} the certifier refutes class 2112112. For the polynomial window, the same engine at $K = 7$ refutes the claim 0.00736 ($F = 0.00735942$), and its certificate at the claim 0.00732 was reproduced box for box independently (18,215,893 boxes). The $K = 8$ claim 0.00796 lies 0.56% below the floating-point minimum 0.0080050. An independent kernel stress test of LB3 on 2400 boxes found no violation. The $K = 8$ and $K = 7$ certificates were each re-run in full: independent runs, not independent code.

7.3 Kernel replay inside Lean

The majorant (done). The checker is written in core Lean. For each cell it proves one enclosure of $(\sin \theta, \cos \theta)$, $\theta = \pi m/30$ at the cell midpoint m , and generates the other 58 cosines by a fixed-point complex recurrence $z_{k+1} = \lfloor z_k z_1 / 2^{64} \rfloor$ whose error grows only linearly ($|z_k - e^{ik\theta}| \leq k \cdot 2^{-55}$, proved); the sums $P(m)$, $P'(m)$, $P''(m)$ are exact integer sums, and a third-order bound with $|P'''| \leq 645$ covers the cell. $P \geq 0$ on $[0, 30]$ needs 600 cells, $\mathcal{B} \geq v_\psi$ on $[0, \frac{1}{2}]$ needs 14. The replay takes about 88 seconds in 13 files, at most 1.64 GB each; the smallest margins are $4.1 \cdot 10^{-6}$ and $4.4 \cdot 10^{-6}$. Negative controls fail as they should, in a python mirror (62 cells) and in the kernel, which proves that the checker returns **false** on $[3.25, 3.75]$ and on $[\frac{1}{4}, \frac{1}{2}]$ for deliberately wrong data. The analytic part (the Fourier pair $\text{sinc}^2 \leftrightarrow \text{triangle}$, the transform of $\mathcal{B}$, its support, $\int \mathcal{B} = \beta_0$, and $2\beta_0 + 4|\beta_{52}| < \frac{33}{10}$ in exact rationals) is proved in Lean. So the hypothesis **CertMajV2** is discharged.

The $K = 5$ all-marks certificate (replayed by the kernel). The branch-and-bound trees above cannot be replayed box by box in the kernel. A second certificate with large leaves was built for the same data: 92,074 nodes and 46,047 leaves over the 20 classes, of three kinds — cap leaves (the penalty alone exceeds the claim), convex leaves (a tangent plane at a floating-point minimiser plus an exact positive semidefiniteness check) and linear-programming leaves

(five affine minorants per span with rational dual multipliers; 4,660, 2,343 and 39,044 leaves respectively). Every leaf is accepted by the exact rational mirror of the corrected checker (version 3, described below; 46,047 leaves, no failure). In a differential test on mutated leaves the mirror was never more permissive than the Lean checker; in any case soundness rests on the kernel replay, not on the mirror. The checker stores each kernel enclosure (the values k, k', k'' at a rational point, or bounds on w, w'' over a cell of width $\frac{1}{64}$) as a record proved once by kernel evaluation; leaves only combine stored rationals. Measured in the replay of class 11111 (1,725 leaves): about 0.24 s per point record, 0.025 s per cell and 0.74 s per leaf; with at most 40 leaves per file the leaf files peak at 3.43 GB in the one-class pilot. The full replay (1,478 files: 1,475 shards and three common files; 46,047 leaves, 89,362 point records; manifest sha256 360d9b22469b9dbf449899a2c6d0bd5be696380e826f137ef20da71ae96c1609) was estimated from the pilot at about 15.6 Lean-hours; it took 20.4. At the minimiser of class 11111 the margin is $1.49 \cdot 10^{-5}$; raising the claim by $1.5 \cdot 10^{-5}$, to $8.9 \cdot 10^{-8}$ above the true minimum, makes exactly one leaf fail in the kernel, the one that contains the minimiser, as it should. The soundness theorem of the checker is split into 28 nodes. An earlier version of one leaf check was unsound: the linear-programming leaf took a tangent line at a point computed from the floating-point minimiser without checking that this point lies in the span interval, where alone the second-derivative penalty is valid; a kernel-checked counterexample (one span, a claim of $\frac{9}{10}$ accepted where the functional is about 0.446) shows it. The majorant certificate uses only the interval and enclosure layer of the checker, which is proved. Version 3 of the checker corrects the defect with one guard, which admits the tangent line at the floating-point point only when that point lies in the span interval; no other definition used by the checker changed, and the 20 class certificates pass it unchanged (through the exact mirror). The whole soundness chain of version 3 is proved with the three standard axioms and is in the library (Appendix A.5): `ZetaS.CertV2.ChainV3.certAM5_of_checks` derives `CertAM5` from the acceptance of the 20 class certificates by the checker and the identity of their data with the data of `CertAM5`.

The replay files use the library’s own definitions. Built from an empty build folder, the one-class pilot compiles 56 files with no failure in 1,976 s, at a peak of 3.49 GB. The full replay, against the library at commit `bb7223b`, whose Lean files the release changes only in comments (Appendix A.9), compiled all 1,478 files, with no failure and no retry, in 2 h 50 min of wall time and 20.4 Lean-hours of processor time; the lowest free memory was 2.48 GB and the build products take 7.1 GB. It ran in two pools, 273 light files sixteen at once and 1,203 heavy files six at once, and then the two top files alone (21 s and 18 s, peak 4.65 and 4.67 GB). The kernel checked the 46,047 leaves and 89,362 point records in exact rational arithmetic. The final file applies `certAM5_of_checks` with the identity of the data by `rfl` and proves `ZetaS.CertV2.certAM5_replayed`; the two theorems of the $K = 5$ pair without hypothesis follow (Appendix A.6). These three theorems depend only on `propext`, `Classical.choice` and `Quot.sound` and rest on no open declaration (Appendix A.8). The 1,478 source files match their manifest file by file, and a scan of all of them finds no `sorry`, `native_decide`, `axiom`, `unsafe`, `implemented_by` or `extern`. The 1,478 files, byte-identical to those checked, are the library `ZetaSReplay` of the repository, and the two theorems are in its module `ReplayHeadlines`; `lake build ZetaSReplay` repeats the replay (§7.4).

The $K = 7$ and $K = 8$ certificates. Large leaves cut the $K = 7$ trees only by a factor of about four: an estimated $6 \cdot 10^7$ leaves and 15 GB of certificate, because the functional has a lattice of near-minimal basins whose number grows geometrically with K (the branch-and-bound trees grow by a factor of 20 to 30 per step in K). This is beyond the 32 GB machine used. These certificates therefore remain displayed hypotheses, identified by the sha256 of Table 3 and verified outside Lean as described in §7.2.

7.4 How to reproduce

Everything is in the repository <https://github.com/Oliverds321/simple-zeros>, at tag v1.0 (Appendix A.9). The Lean libraries are at its root, the replay files of §7.3 are the library ZetaSReplay, and the folder `supplementary/` holds the data of Table 3, the certifiers and their logs, and the generator of the replay files, with a guide (`supplementary/README.md`), the commands that run them (`supplementary/VERIFY.md`) and the sha256 of every file (`supplementary/CHECKSUMS.sha256`). Commands are run from the root of the repository.

- *The statements and their axioms.* The toolchain is pinned in `lean-toolchain`, and `lake exe cache get` fetches the prebuilt Mathlib. Then

```
lake build ZetaS
lake build Solution.FollowUpZeta
lake env lean comparator/PrintAxioms/FollowUpZeta.lean
```

Every line printed by the last command must read `... depends on axioms: [propext, Classical.choice, Quot.sound]`. This checks the fourteen statements of record (Appendix A.3), in which `CertAM5`, `CertAM7` and `CertS8` are hypotheses. The files `comparator/README_followup.md` and `comparator/config-followup.json` (the fourteen theorem names and the permitted axioms) describe the same check with the comparator tool used by the repository of [2].

- *The kernel replay of `CertAM5`* (about 20 Lean-hours, and several GB of memory per file, §7.3):

```
lake build ZetaSReplay
lake env lean comparator/PrintAxioms/FollowUpReplay.lean
```

The last command prints the axioms of `ZetaS.Top.zeta_simple_K5_uncond`, `ZetaS.Top.zeta_distinct_K5_uncond` and `ZetaS.CertV2.certAM5_replayed`; again every line must read as above.

- *The local certificates outside Lean.* In `supplementary/certificates/`: `CertS8/` (the certifier `bnb_lb3.py` with `polywin.py`; the driver logs of the run and of the re-run, which end with `ALL SHARDS OK`), `CertAM5/` (`bnb_interval_w.py`, the $K = 5$ run, with its driver `runcert.sh`) and `CertAM7/` (`bnb_lb3_w.py`, the LB3 engine with the all-marks check, with its driver `runlb3.sh`, and the exact constants `pexact2.py`), each with its data and logs; and `common/`, with the input checks `verify_inputs.py` (weights nonnegative, span sums 2, penalties summing to ν , claims exactly $\sum_i b_i(m_i)$, all patterns covered) and `d819_exact.py`.
- *The majorant:* `supplementary/certificates/CertMajV2/r4_majorant_arb.py` (Arb, 106 bits); in Lean the files `MajCheck.lean`, `MajSoundNPos.lean`, `MajSoundNMaj.lean` and the 13 replay shards, all in `ZetaS/Majorant/` and built with `ZetaS`.
- *The replay files from the certificate:* `supplementary/replay-K5/` holds the certificate with large leaves (`v2_K5_*.json`, one file per class, and `v2common_K5_4.json`) and the generator `gen_replay.py`, which writes the replay files from it; `ZetaSReplay/MANIFEST.sha256` gives the sha256 of each replay file.

8 Remarks on the limits of the method

Ceilings. Alpöge and Furman [2, §7.2] call a *bandwidth-one certificate* for N_0^s/N a function of the configuration that depends on it only through the first two trace moments against

test functions of Fourier support in $[-1, 1]$ and the on-line/off-line partition, and is valid configuration by configuration, and they describe a ceiling of about 0.682 for such certificates. What is proved there is a quantitative bound with qualifications. For a constant c_0 and a weight r on $[0, 1]$ ($r \in C^1$, with r' differentiable off a countable set and r'' integrable) such that $c_0 + \sum_{j=1}^{256} (S_j/256) r(j/256) \leq p$, where (S_j) is the sampled form factor of an explicit 256-periodic law, the Lean theorem `Zeta23.PairCeiling.ceiling_law256` gives $c_0 + \int_0^1 r(x) x dx \leq p + 0.824|r(1)| + 2.55 \cdot 10^{-6}(|r'(1)| + \int_0^1 |r''|)$, under the displayed hypothesis `Enc10K` that (S_j) lies in enclosures obtained outside Lean by interval arithmetic. For a certificate of this form that is valid configuration by configuration, averaging over the law (a step done on paper) gives the hypothesis with $p = p_0 \leq 0.6818287$, the proportion of simple points of the law, and the bound is then below 0.6819 if $r(1) = 0$ and $|r'(1)| + \int_0^1 |r''| < 8.2$, in particular for the two windows of [2, (2.6)]. By the 0.682 ceiling we mean this bound; a bound of 0.682 for every bandwidth-one certificate is not proved in [2] or here. Theorem 1.2 lies below 0.682. Within that class, the route of Section 4, which keeps only the energy among *simple* zeros, is capped much lower: periodic adversaries with gaps alternating near 1.03 and 1.96 carry an energy of about 0.0019 per simple zero, which caps cosine windows near 0.6737 (0.673666 at $\cos 1.48s$ with exact energies, 0.673697 for the best cosine; Devine [10] states that with the fixed optimal Montgomery–Taylor kernel the pure Gram/rank machinery cannot force a bound above 0.6736 [V~]; trmdy [24] states on 11 August 2026 that certificates reading only pair energies of simple zeros cannot exceed 0.674826, and Zhongshan-Big-Jun [33] give pressure-class ceilings near 0.6731 [V~]), and the same adversaries give at most 0.673776 for the window $\tilde{\psi}$; the caps of the chain as K grows level off near 0.67350 for $\tilde{\psi}$. These values are floating-point computations, not theorems; they explain why the bandwidth-one values of §1.4 and Theorem 1.2 are all close together. (The repository teal-sea/zeta-lab states that the pressure-certificate family built on the Montgomery–Taylor window saturates at $\sup \Phi_n \leq 0.675142509660254$ [V~]; we have not compared that family with the class of our computations.)

The route of Section 5 is not subject to the simple-only cap, because the binding adversaries of that cap use on-line double zeros, whose energy it keeps. Its own limits, again in floating point: the linear programme that chooses the weights gives 0.6761153 at $K = 7$ ($\mu = \frac{1}{500}$) and 0.6763110 at $K = 8$ ($\mu = \frac{1}{600}$), and the all-marks functional with a single window and exact energies is capped at 0.681447 (at $\psi = \cos 2s$). A $K = 8$ certificate was not attempted: the trees grow by a factor of 20 to 30 per step in K , an estimated 1 to $1.4 \cdot 10^{10}$ boxes over 136 classes. Flat-topped windows lowered the $K = 5$ value, so $\cos 1.6s$ was kept. Whether an analogue of the 0.682 ceiling holds for bandwidth-one certificates of N^s/N we have not determined.

What would be needed to go further. For N^s and N^d the constants move with the local certificate. The $K = 5$ certificate with large leaves was small enough for the kernel to replay (§7.3); larger K , better windows and a certificate at $K \geq 6$ small enough for kernel replay would need a better leaf (one whole basin per leaf) rather than more boxes. Beyond bandwidth one, or beyond the first two trace moments, the inputs of [2] are no longer enough.

For N_0^s the route of Section 5 would give the same constants, 0.675158622 with the $K = 5$ certificate and 0.676102666 with $K = 7$, if one inequality held for the zeros of ζ : that the slack of the rank–trace inequality is at least $a_1 s_1 + \frac{a_2}{2}(N - s_1) - \nu N - o(N)$. Lemma 5.2 and Proposition 5.3 prove it whenever all zeros with $\operatorname{Re} \gamma \in I'$ lie on the line, so the obstruction is the off-line pairs near the line, which for N_0^s bring no bonus (Remark 5.9; trmdy [24] named the pricing of off-line pairs as the gate beyond the simple-only route on 11 August 2026 [V~]). We have reduced this inequality to families of local block inequalities and proved several of them on explicit classes of configurations, but it is not proved, and none of that work is part of this paper.

A The Lean statements

The Lean text below is copied verbatim, by a script that reads the files, from the release of the formalisation: the repository <https://github.com/Oliverds321/simple-zeros> at tag v1.0 (§A.9). Only line breaks inside long lines are added by the typesetting.

What is claimed is stated in three files, which import nothing but Mathlib and one another: `comparator/ChallengeDeps.lean` of the Lean artifact of [2], which defines the zeros of `riemannZeta` and the counting functions (§A.1); the definition layer of this paper, `comparator/ChallengeDeps/FollowUpZeta.lean`, with the windows, the local inequalities and the three certificate Props (§A.2); and the fourteen statements of record, `comparator/Challenge/FollowUpZeta.lean` (§A.3). A reader who wants to know what is claimed needs only these three files. The statements of record carry `sorry` by design: they are the challenge side of the comparator set-up of the repository of [2]. The file `comparator/Solution/FollowUpZeta.lean` proves exactly these fourteen statements from the library `ZetaS`, whose definitions are copies of those of §A.2 in the namespace `ZetaS`; the conversions between the two are proofs, not trusted text. The library's theorems are in §A.4, the chain for the $K = 5$ certificate in §A.5, the replay's theorems in §A.6 and the majorant in §A.7.

The doc-strings are part of the Lean text and are reproduced unchanged. They are development notes, not part of what is claimed, and not all current: they refer to an earlier draft of this paper by its labels (for instance `thm:zeta-mainw2` is Theorem 1.2, `thm:sgd-Sigma` and `thm:sgd-D` are Theorem 1.1, `cor:oll-SC` is Corollary 1.3, `lem:sgd-env` is Lemma 5.5) and to its line numbers, they name work packages, folders and dates of the development, and “level A” in them means kernel-checked in the sense of §1.5. The data files that the doc-strings name are in `supplementary/certificates/`, with the same sha256 (Table 3). The header comment of §4 of the definition layer says that the certificates were established in Arb ball arithmetic; as §7.2 states, the box bounds were computed in outward-rounded float64 intervals, with the kernel values enclosed in Arb. Only the Lean code is claimed.

A.1 The counting functions

From `comparator/ChallengeDeps.lean` of the repository of [2], unchanged from the tree as shipped (the functions used below are `Ncount`, `Ndist`, `N0simple`, `Nsimple` and `zerosIn`). The phrase “Not used by the challenge statements” in two of these doc-strings refers to the challenge statements of [2]; `Nsimple` is used by those of §A.3.

```
#!/- ## 1. Nontrivial zeros of  $\zeta$  and the counting functions (Theorems A--D) -/

/--  $\rho$  is a nontrivial zero of the Riemann zeta function:  $\zeta(\rho) = 0$  with  $0 < \text{Re } \rho < 1$  (the
↪ open
critical strip). Every zero in the strip is "nontrivial" in the usual sense (it is neither a
↪ trivial
zero  $-2(n+1)$ , of real part  $\leq -2$ , nor the pole  $s = 1$ ); the converse --- that every zero
↪ other than the
trivial ones lies in the open strip --- is classical and is not needed to state the
↪ theorems. -/
def IsNontrivialZero ( $\rho : \mathbb{C}$ ) : Prop := riemannZeta  $\rho$  = 0  $\wedge$  0 <  $\rho$ .re  $\wedge$   $\rho$ .re < 1

/--  $m_\rho$ , the multiplicity of  $\rho$ : the order of vanishing of  $\zeta$  at  $\rho$ , via Mathlib's
↪ `analyticOrderAt`
( $\mathbb{N}\infty$ -valued; `toNat` sends  $\top$  to 0, so  $\text{`1} \leq \text{zeroMult } \rho\text{'}$  encodes both " $\zeta$  is not locally
↪ identically
zero at  $\rho$ " and " $\rho$  is a zero"). -/
def zeroMult ( $\rho : \mathbb{C}$ ) :  $\mathbb{N}$  := (analyticOrderAt riemannZeta  $\rho$ ).toNat
```

```

/-- The nontrivial zeros with ordinate in the window  $(T_1, T_2]$ :  $\{\rho \mid \zeta(\rho) = 0, 0 < \text{Re } \rho < 1, T_1 < \text{Im } \rho \leq T_2\}$ . (Positive-ordinate window, not  $|\text{Im } \rho|$ .) -/
def zerosIn (T1 T2 : ℝ) : Set ℂ := {ρ | IsNontrivialZero ρ ∧ T1 < ρ.im ∧ ρ.im ≤ T2}

/-- N(T1,T2): the number of nontrivial zeros with  $T_1 < \text{Im } \rho \leq T_2$ , counted WITH multiplicity
(`Σf` is Mathlib's finite sum; that the window contains finitely many zeros is proved on the
solution side, not assumed). -/
def Ncount (T1 T2 : ℝ) : ℕ := Σf ρ ∈ zerosIn T1 T2, zeroMult ρ

/-- Nd(T1,T2): the number of nontrivial zeros with  $T_1 < \text{Im } \rho \leq T_2$ , each distinct zero
↪ counted
once. -/
def Ndist (T1 T2 : ℝ) : ℕ := (zerosIn T1 T2).ncard

/-- N0(T1,T2): the zeros ON the critical line  $\text{Re } \rho = 1/2$  with  $T_1 < \text{Im } \rho \leq T_2$ , with
↪ multiplicity.
(Not used by the challenge statements; kept so that the block matches
↪ Zeta23/Statement.lean.) -/
def N0 (T1 T2 : ℝ) : ℕ := Σf ρ ∈ zerosIn T1 T2 ∩ {ρ | ρ.re = 1 / 2}, zeroMult ρ

/-- N0*(T1,T2): the number of DISTINCT zeros ON the critical line  $\text{Re } \rho = 1/2$  with
 $T_1 < \text{Im } \rho \leq T_2$ . This is what Theorems A and D bound from below. -/
def N0star (T1 T2 : ℝ) : ℕ := (zerosIn T1 T2 ∩ {ρ | ρ.re = 1 / 2}).ncard

/-- N0s(T1,T2): the number of SIMPLE zeros (multiplicity exactly 1) ON the critical line with
 $T_1 < \text{Im } \rho \leq T_2$ . -/
def N0simple (T1 T2 : ℝ) : ℕ := (zerosIn T1 T2 ∩ {ρ | ρ.re = 1 / 2} ∩ {ρ | zeroMult ρ =
↪ 1}).ncard

/-- Ns(T1,T2): the number of simple zeros with  $T_1 < \text{Im } \rho \leq T_2$  (anywhere in the strip).
(Not used by the challenge statements; kept so that the block matches
↪ Zeta23/Statement.lean.) -/
def Nsimple (T1 T2 : ℝ) : ℕ := (zerosIn T1 T2 ∩ {ρ | zeroMult ρ = 1}).ncard

```

A.2 The definitions of this paper

The definition layer `comparator/ChallengeDeps/FollowUpZeta.lean` (namespace `FollowUpZeta`), from its first line of code to its end; the file's header comment, which lists its contents, is omitted.

```

import ChallengeDeps

open scoped BigOperators
open Finset

noncomputable section

namespace FollowUpZeta

/-! ## §1 Zeros that are simple or on the critical line -/

/-- Nsc(T1,T2) (cor:oll-SC): the nontrivial zeros with  $T_1 < \text{Im } \rho \leq T_2$  that lie ON the
↪ critical line, counted WITH
multiplicity, plus those OFF the line that are simple (both members  $\rho, 1 - \rho$  of each simple
↪ off-line pair; multiple
off-line zeros are not counted). An off-line zero in the set is simple, so its weight
↪ `zeroMult ρ` is 1. -/
def Nsc (T1 T2 : ℝ) : ℕ :=
  Σf ρ ∈ zerosIn T1 T2 ∩ ({ρ | ρ.re = 1 / 2} ∪ {ρ | zeroMult ρ = 1}), zeroMult ρ

/-! ## §2 Windows and the normalised kernel -/

```

```

/-- The polynomial window  $\psi^*$  of thm:zeta-mainw2:
 $\psi^*(s) = (50000/50037)(1 + (127/2500)u - (8997/10000)u^2 + (6187/10000)u^3 + (3/1250)u^4)$ ,  $u = \hookrightarrow (2s)^2$ ,
on  $[-1/2, 1/2]$  (values outside are never read). -/
def psiPoly8A (s : ℝ) : ℝ :=
  50000 / 50037 * (1 + 127 / 2500 * (2 * s) ^ 2 - 8997 / 10000 * ((2 * s) ^ 2) ^ 2
    + 6187 / 10000 * ((2 * s) ^ 2) ^ 3 + 3 / 1250 * ((2 * s) ^ 2) ^ 4)

/-- The window  $\psi(s) = \cos(1.6 s)$  of thm:sgd-Sigma and thm:sgd-D. -/
def psiCos16 (s : ℝ) : ℝ := Real.cos (8 / 5 * s)

/-- The normalised Fourier transform of a window:  $k_\psi(x) = \int_{-1/2}^{1/2} \psi(s) \cos(2\pi xs)$ 
 $\hookrightarrow ds / \int_{-1/2}^{1/2} \psi$ . -/
def kPsi ( $\psi : \mathbb{R} \rightarrow \mathbb{R}$ ) (x : ℝ) : ℝ :=
  ( $\int s \text{ in } (-(1 / 2 : \mathbb{R}))..(1 / 2)$ ,  $\psi s * \text{Real.cos } (2 * \text{Real.pi} * x * s)$ )
    /  $\int s \text{ in } (-(1 / 2 : \mathbb{R}))..(1 / 2)$ ,  $\psi s$ 

/-! ## §3 The local inequalities -/

/-- Position weights for K consecutive points (lem:zeta-Dprime), 0-based: pair weights  $\gamma$  s
 $\hookrightarrow i$  for spans
 $1 \leq s \leq K - 1$  and positions  $0 \leq i < K - s$ , nonnegative and summing to 2 for every span;
 $\hookrightarrow$  gap penalties  $\mu \ 1 \geq 0$ ,
l : Fin (K - 1). Values of  $\gamma$  outside the index range are never read. -/
structure LocalWeights (K : ℕ) where
   $\gamma : \mathbb{N} \rightarrow \mathbb{N} \rightarrow \mathbb{R}$ 
   $\mu : \text{Fin } (K - 1) \rightarrow \mathbb{R}$ 
  two_le :  $2 \leq K$ 
   $\gamma_{\text{nonneg}} : \forall s \in \text{Icc } 1 \ (K - 1), \forall i \in \text{range } (K - s), 0 \leq \gamma \ s \ i$ 
   $\gamma_{\text{sum}} : \forall s \in \text{Icc } 1 \ (K - 1), \sum i \in \text{range } (K - s), \gamma \ s \ i = 2$ 
   $\mu_{\text{nonneg}} : \forall l, 0 \leq \mu \ l$ 

/--  $\nu := \sum_i \mu_i$ . -/
def LocalWeights.nu {K : ℕ} (W : LocalWeights K) : ℝ :=  $\sum l, W.\mu \ l$ 

/-- The span  $g_i + \dots + g_{\{i+s-1\}}$  of a gap vector (0-based; indices  $\geq K - 1$  contribute
 $\hookrightarrow$  nothing). -/
def gapSpan {K : ℕ} (g : Fin (K - 1) → ℝ) (i s : ℕ) : ℝ :=
   $\sum l \in \text{univ.filter } (\text{fun } l : \text{Fin } (K - 1) \Rightarrow i \leq l.\text{val} \wedge l.\text{val} < i + s), g \ l$ 

/--  $F_{\{\gamma, \mu\}}(g) = \sum_i \mu_i g_i + \sum_{\{s=1\}^{\wedge}\{K-1\}} \sum_i \gamma_{\{s, i\}} k(g_i + \dots + g_{\{i+s-1\}})^2$ 
 $\hookrightarrow$  (eq:zeta-LIw). -/
def localF (k : ℝ → ℝ) {K : ℕ} (W : LocalWeights K) (g : Fin (K - 1) → ℝ) : ℝ :=
  ( $\sum l, W.\mu \ l * g \ l$ )
    +  $\sum s \in \text{Icc } 1 \ (K - 1), \sum i \in \text{range } (K - s), W.\gamma \ s \ i * k \ (\text{gapSpan } g \ i \ s) ^ 2$ 

/-- (LI') of lem:zeta-Dprime:  $F_{\{\gamma, \mu\}}(g) \geq c$  for every gap vector  $g \in [0, \infty)^{\wedge}\{K-1\}$ . -/
def LocalCert (k : ℝ → ℝ) {K : ℕ} (W : LocalWeights K) (c : ℝ) : Prop :=
   $\forall g : \text{Fin } (K - 1) \rightarrow \mathbb{R}, (\forall l, 0 \leq g \ l) \rightarrow c \leq \text{localF } k \ W \ g$ 

/-- All-marks data (def:zeta-LIm): position weights as in `LocalWeights` and site credits `b
 $\hookrightarrow i \ j$  for site  $i : \text{Fin } K$  and
capped mark  $j : \text{Fin } 2$  ( $j = 0 \mapsto \text{mark } 1, j = 1 \mapsto \text{mark } 2$ ); no sign condition. -/
structure MarkWeights (K : ℕ) extends LocalWeights K where
  b : Fin K → Fin 2 → ℝ

/--  $a_j := \sum_i b_i(j)$ ;  $a_0 = a_1$ ,  $a_1 = a_2$ . -/
def MarkWeights.a {K : ℕ} (W : MarkWeights K) (j : Fin 2) : ℝ :=  $\sum i, W.b \ i \ j$ 

/-- The numerical mark  $m_i \in \{1, 2\}$  at 0-based position  $i$  of a pattern (0 outside  $[0, K)$ ).
 $\hookrightarrow$  -/
def markVal {K : ℕ} (m : Fin K → Fin 2) (i : ℕ) : ℝ :=

```

```

    if h : i < K then ((m ⟨i, h⟩ : ℕ) : ℝ) + 1 else 0

/-- F_m(g) =  $\sum_{\{s,i\}} \gamma_{\{s,i\}} m_i m_{\{i+s\}} k(g_i + \dots + g_{\{i+s-1\}})^2 + \sum_i \mu_i g_i$ 
↪ (def:zeta-LIm). -/
def localFm (k : ℝ → ℝ) {K : ℕ} (W : MarkWeights K) (m : Fin K → Fin 2) (g : Fin (K - 1)
↪ → ℝ) : ℝ :=
  (∑ s ∈ Icc 1 (K - 1), ∑ i ∈ range (K - s),
    W.γ s i * markVal m i * markVal m (i + s) * k (gapSpan g i s) ^ 2)
  + ∑ l, W.μ l * g l

/-- (LI_m) of def:zeta-LIm: for EVERY mark pattern  $m \in \{1,2\}^K$  and every  $g \in [0,\infty)^{K-1}$ ,
↪  $F_m(g) \geq \sum_i b_i(m_i)$ . -/
def LocalCertAM (k : ℝ → ℝ) {K : ℕ} (W : MarkWeights K) : Prop :=
  ∀ m : Fin K → Fin 2, ∀ g : Fin (K - 1) → ℝ, (∀ l, 0 ≤ g l) → ∑ i, W.b i (m i) ≤
  ↪ localFm k W m g

/-! ## §4 The certificate hypotheses

Each Prop asserts that SOME weights with the stated constants satisfy the local inequality
↪ on the whole orthant. The
weights that realise it are exact rationals in the data files named below (identified by
↪ sha256); the inequality for those
weights was established outside Lean by interval branch-and-bound (Arb ball arithmetic,
↪ every box accepted only when a
rigorous lower bound exceeds the claim; the unbounded part of the orthant is discarded
↪ exactly, since
 $F \geq (\min \mu) \cdot \sum g$  there). Paths are relative to the development's project folder; the files
↪ named here are in
`supplementary/certificates/` of the repository, some logs under new names (map:
↪ `supplementary/README.md`). -/

/-- **Certificate S8** (thm:zeta-mainw2, rem:zeta-P8, lem:zeta-Dprime):  $K = 8$ , window  $\psi$ ,
↪ weights with  $\nu = 7/1700$  for
which (LI') holds at the claim  $c = 0.00796$ .
Data: `round2/d6_5_numerics/w_poly8A_K8_mu1700.json`, sha256
`257b340d9da4f343600a21589f16d8915f3e082499d475e153c5453980671e78`; window
↪ `round2/d6_5_numerics/win_poly8A.json`,
sha256 `c4005f57a7920f78aa2df33b093c9c5030efac2b8e39fc6cf7371b0c0f32dc14` (the unnormalised
↪  $v = (50037/50000)\psi$ ; the kernel
is homogeneous of degree 0). Certifier `round2/d6_5_numerics/bnb_lb3.py` with `polywin.py`;
↪ runs
`round2/d6_5_numerics/runs/P8/DRIVER.log` and the independent re-run `runs/P8b/DRIVER.log`,
↪ both ending
`ALL SHARDS OK: certificate claim=0.00796 holds` (about  $4.5 \cdot 10^8$  boxes each). -/
def CertS8 : Prop :=
  ∃ W : LocalWeights 8, W.nu = 7 / 1700 ∧ LocalCert (kPsi psiPoly8A) W (199 / 25000)

/-- **Certificate AM5** (thm:zeta-allmarks, rem:sigd-cert, the  $K = 5$  row): all-marks data
↪ with  $a_1 = 1280197/10^8$ ,
 $a_2 = 48749/3125000$ ,  $\nu = 1/125$  for which (LI_m) holds for  $k_\psi$ ,  $\psi = \cos 1.6s$ .
Data: `round1/X2_numerics/claims_a1.6_K5_mu500.json` (byte-identical to
↪ `claims_d15e6_a1.6_K5_mu500.json`), sha256
`cf055dd2ee9a9e469b6510fe70b5fbc416b8a40515cc6746bc390d40d88777f2`; the 20 class files
↪ `xpat_K5_*.json` of the same
folder, sha256 of their concatenation in sorted file-name order
`a5947c01d34ddcbcb7366861b35e83942099fcbf7ac71e13fcd432796882dd34e`. Certifier
↪ `round1/X2_numerics/bnb_interval_w.py`;
logs `cert_K5_w_a16_mu500_d15e6.txt` and `cert_K5_w_a16_mu500_d15e6_b.txt` (20/20 reversal
↪ classes ok, covering the 32
patterns; the weights are reversal-symmetric). A replay of this certificate by the Lean
↪ kernel was completed on
28 September 2026: `ZetaS.CertV2.certAM5_replayed` proves the library's copy
↪ `ZetaS.CertAM5`, and the hypothesis-free

```

```

K = 5 theorems are in the module `ReplayHeadlines` of the library `ZetaSReplay`. In the
↪ statements of this topic the
Prop remains a displayed hypothesis, like the other two. -/
def CertAM5 : Prop :=
  ↪ ∃ W : MarkWeights 5, W.a 0 = 1280197 / 10 ^ 8 ∧ W.a 1 = 48749 / 3125000 ∧ W.nu = 1 / 125
  ↪ ∧
  ↪ LocalCertAM (kPsi psiCos16) W

/-- **Certificate AM7** (thm:sgd-Sigma, thm:sgd-D, rem:sgd-cert, the K = 7 row):
↪ all-marks data with
a1 = 1824837/108, a2 = 1168069/(5·107), ν = 3/250 for which (LIm) holds for kψ, ψ = cos
↪ 1.6s.
Data: `round2/d8_9_numerics/k7m500/rerun/claims_a1.6_K7_mu500_L1e5.json` (the final claims,
↪ every claim of the LP lowered by
exactly 10-5), sha256 `d2cf393019407bdc797c73147ec8e3dc66ec41fe745600c84173f74ac66bfc5c`;
↪ the 72 class files
`xpat_K7_*.json` of the same folder, sha256 of their concatenation in sorted file-name order
`33c4c8c5762bcebdf8f544ab34a137a6a7ba4065d08f5caee0a51188e1a67514`. Certifier
↪ `round2/d8_9_numerics/bnb_lb3_w.py`
(driver `runlb3.sh`); certificate log `round2/d8_9_numerics/k7m500/rerun/cert_lb3_rerun.txt`
↪ (72/72 reversal classes ok,
each at exactly the final claim); completeness (the 72 classes cover all 128 patterns, each
↪ certified at a claim ≥ the
required one) `round2/d8_21_numerics/k7_verify_pexact2.txt`. The weights are
↪ reversal-symmetric. -/
def CertAM7 : Prop :=
  ↪ ∃ W : MarkWeights 7, W.a 0 = 1824837 / 10 ^ 8 ∧ W.a 1 = 1168069 / (5 * 10 ^ 7) ∧ W.nu = 3
  ↪ / 250 ∧
  ↪ LocalCertAM (kPsi psiCos16) W

end FollowUpZeta

```

A.3 The fourteen statements of record

The statement file `comparator/Challenge/FollowUpZeta.lean`, from its first line of code to its end (the header comment is omitted). The six headline statements are those of Theorems 1.1 and 1.2 and Corollary 1.3, each with its single certificate hypothesis; the six `_cumulative` forms are the same bounds on the windows $0 < \text{Im } \rho \leq T$; the last two statements show that the counts are genuine (every window holds finitely many zeros) and that the denominator tends to infinity.

```

import ChallengeDeps.FollowUpZeta

open FollowUpZeta

noncomputable section

/-- **Simple zeros (K = 5 certificate)**, thm:sgd-Sigma with the K = 5 row of
↪ rem:sgd-cert: at least 0.675158622 of
the nontrivial zeros with  $T < \text{Im } \rho \leq 2T$  (counted with multiplicity) are simple, for all
↪ large T. -/
theorem zeta_simple_K5 (hcert : CertAM5) :
  ↪ ∀ ε > 0, ∃ T0 : ℝ, ∀ T ≥ T0,
  ↪ ((675158622 : ℝ) / 10 ^ 9 - ε) * (Ncount T (2 * T) : ℝ) ≤ Nsimple T (2 * T) := by
  ↪ sorry

/-- `zeta_simple_K5` on the cumulative windows  $0 < \text{Im } \rho \leq T$  (same constant). -/
theorem zeta_simple_K5_cumulative (hcert : CertAM5) :
  ↪ ∀ ε > 0, ∃ T0 : ℝ, ∀ T ≥ T0,
  ↪ ((675158622 : ℝ) / 10 ^ 9 - ε) * (Ncount 0 T : ℝ) ≤ Nsimple 0 T := by

```

```

sorry

/-- **Distinct zeros (K = 5 certificate)**, thm:sgd-D with the K = 5 row of rem:sgd-cert:
↪ at least 0.837579311. -/
theorem zeta_distinct_K5 (hcert : CertAM5) :
   $\forall \varepsilon > 0, \exists T_0 : \mathbb{R}, \forall T \geq T_0,$ 
   $((837579311 : \mathbb{R}) / 10^{9 - \varepsilon}) * (\text{Ncount } T (2 * T) : \mathbb{R}) \leq \text{Ndist } T (2 * T) := \text{by}$ 
  sorry

/-- `zeta_distinct_K5` on the cumulative windows  $0 < \text{Im } \rho \leq T$  (same constant). -/
theorem zeta_distinct_K5_cumulative (hcert : CertAM5) :
   $\forall \varepsilon > 0, \exists T_0 : \mathbb{R}, \forall T \geq T_0,$ 
   $((837579311 : \mathbb{R}) / 10^{9 - \varepsilon}) * (\text{Ncount } 0 T : \mathbb{R}) \leq \text{Ndist } 0 T := \text{by}$ 
  sorry

/-- **Simple zeros**, thm:sgd-Sigma (K = 7 certificate): at least 0.676102666. -/
theorem zeta_simple_K7 (hcert : CertAM7) :
   $\forall \varepsilon > 0, \exists T_0 : \mathbb{R}, \forall T \geq T_0,$ 
   $((676102666 : \mathbb{R}) / 10^{9 - \varepsilon}) * (\text{Ncount } T (2 * T) : \mathbb{R}) \leq \text{Nsimple } T (2 * T) := \text{by}$ 
  sorry

/-- `zeta_simple_K7` on the cumulative windows  $0 < \text{Im } \rho \leq T$  (same constant). -/
theorem zeta_simple_K7_cumulative (hcert : CertAM7) :
   $\forall \varepsilon > 0, \exists T_0 : \mathbb{R}, \forall T \geq T_0,$ 
   $((676102666 : \mathbb{R}) / 10^{9 - \varepsilon}) * (\text{Ncount } 0 T : \mathbb{R}) \leq \text{Nsimple } 0 T := \text{by}$ 
  sorry

/-- **Distinct zeros**, thm:sgd-D (K = 7 certificate): at least 0.838051333. -/
theorem zeta_distinct_K7 (hcert : CertAM7) :
   $\forall \varepsilon > 0, \exists T_0 : \mathbb{R}, \forall T \geq T_0,$ 
   $((838051333 : \mathbb{R}) / 10^{9 - \varepsilon}) * (\text{Ncount } T (2 * T) : \mathbb{R}) \leq \text{Ndist } T (2 * T) := \text{by}$ 
  sorry

/-- `zeta_distinct_K7` on the cumulative windows  $0 < \text{Im } \rho \leq T$  (same constant). -/
theorem zeta_distinct_K7_cumulative (hcert : CertAM7) :
   $\forall \varepsilon > 0, \exists T_0 : \mathbb{R}, \forall T \geq T_0,$ 
   $((838051333 : \mathbb{R}) / 10^{9 - \varepsilon}) * (\text{Ncount } 0 T : \mathbb{R}) \leq \text{Ndist } 0 T := \text{by}$ 
  sorry

/-- **Simple zeros on the critical line**, thm:zeta-mainw2 (window  $\psi^*$ , K = 8): at least
↪ 0.6733736895. -/
theorem zeta_simple_on_line (hcert : CertS8) :
   $\forall \varepsilon > 0, \exists T_0 : \mathbb{R}, \forall T \geq T_0,$ 
   $((6733736895 : \mathbb{R}) / 10^{10 - \varepsilon}) * (\text{Ncount } T (2 * T) : \mathbb{R}) \leq \text{N0simple } T (2 * T) := \text{by}$ 
  sorry

/-- `zeta_simple_on_line` on the cumulative windows  $0 < \text{Im } \rho \leq T$  (same constant). -/
theorem zeta_simple_on_line_cumulative (hcert : CertS8) :
   $\forall \varepsilon > 0, \exists T_0 : \mathbb{R}, \forall T \geq T_0,$ 
   $((6733736895 : \mathbb{R}) / 10^{10 - \varepsilon}) * (\text{Ncount } 0 T : \mathbb{R}) \leq \text{N0simple } 0 T := \text{by}$ 
  sorry

/-- **Zeros that are simple or on the critical line**, cor:oll-SC (on-line zeros with
↪ multiplicity plus simple
off-line zeros): at least 0.8879195. -/
theorem zeta_simple_or_critical (hcert : CertS8) :
   $\forall \varepsilon > 0, \exists T_0 : \mathbb{R}, \forall T \geq T_0,$ 
   $((8879195 : \mathbb{R}) / 10^{7 - \varepsilon}) * (\text{Ncount } T (2 * T) : \mathbb{R}) \leq \text{Nsc } T (2 * T) := \text{by}$ 
  sorry

/-- `zeta_simple_or_critical` on the cumulative windows  $0 < \text{Im } \rho \leq T$  (same constant). -/
theorem zeta_simple_or_critical_cumulative (hcert : CertS8) :
   $\forall \varepsilon > 0, \exists T_0 : \mathbb{R}, \forall T \geq T_0,$ 

```

```

      ((8879195 : ℝ) / 10 ^ 7 - ε) * (Ncount 0 T : ℝ) ≤ Nsc 0 T := by
    sorry

/-- Non-vacuity companion: every window  $T_1 < \text{Im } \rho \leq T_2$  holds finitely many nontrivial zeros,
↪ so `Ncount`, `Nsc`
(finite sums) and `Nsimple`, `Ndist`, `N0simple` (`Set.ncard`) are genuine counts (Lean's
↪ `finsum` and `ncard` would
return 0 on an infinite set). -/
theorem zeta_zerosIn_finite (T1 T2 : ℝ) :
  (zerosIn T1 T2).Finite := by
  sorry

/-- Non-vacuity companion: the denominator  $N(T, 2T)$  tends to infinity. -/
theorem zeta_tendsto_Ncount :
  Filter.Tendsto (fun T : ℝ => (Ncount T (2 * T) : ℝ)) Filter.atTop Filter.atTop := by
  sorry

end

```

A.4 The library's headline theorems

The four certificate-only theorems for simple and distinct zeros (`ZetaS/Top/TopFinal.lean`). They apply the forms with the majorant threshold $\frac{33}{10}$ to the proved majorant certificate `MajSound.certMajV2.proved` (§A.7), so that their only hypothesis is the local certificate:

```

namespace ZetaS
namespace Top

/-- `zeta_simple_K5_V2` with the majorant certificate proved (`MajSound.certMajV2.proved`).
↪ -/
theorem zeta_simple_K5_final (hcert : CertAM5) :
  ∀ ε > 0, ∃ T0 : ℝ, ∀ T ≥ T0,
    ((675158622 : ℝ) / 10 ^ 9 - ε) * (Ncount T (2 * T) : ℝ) ≤ Nsimple T (2 * T) :=
  zeta_simple_K5_V2 hcert MajSound.certMajV2.proved

/-- `zeta_distinct_K5_V2` with the majorant certificate proved
↪ (`MajSound.certMajV2.proved`). -/
theorem zeta_distinct_K5_final (hcert : CertAM5) :
  ∀ ε > 0, ∃ T0 : ℝ, ∀ T ≥ T0,
    ((837579311 : ℝ) / 10 ^ 9 - ε) * (Ncount T (2 * T) : ℝ) ≤ Ndist T (2 * T) :=
  zeta_distinct_K5_V2 hcert MajSound.certMajV2.proved

/-- `zeta_simple_K7_V2` with the majorant certificate proved (`MajSound.certMajV2.proved`).
↪ -/
theorem zeta_simple_K7_final (hcert : CertAM7) :
  ∀ ε > 0, ∃ T0 : ℝ, ∀ T ≥ T0,
    ((676102666 : ℝ) / 10 ^ 9 - ε) * (Ncount T (2 * T) : ℝ) ≤ Nsimple T (2 * T) :=
  zeta_simple_K7_V2 hcert MajSound.certMajV2.proved

/-- `zeta_distinct_K7_V2` with the majorant certificate proved
↪ (`MajSound.certMajV2.proved`). -/
theorem zeta_distinct_K7_final (hcert : CertAM7) :
  ∀ ε > 0, ∃ T0 : ℝ, ∀ T ≥ T0,
    ((838051333 : ℝ) / 10 ^ 9 - ε) * (Ncount T (2 * T) : ℝ) ≤ Ndist T (2 * T) :=
  zeta_distinct_K7_V2 hcert MajSound.certMajV2.proved

end Top
end ZetaS

```

The two theorems for simple zeros on the line and for zeros that are simple or on the line

(ZetaS/Top/TopSSC.lean, namespace ZetaS.Top); statements only, each followed in the file by a tactic proof of about ten lines:

```

/-- **S --- simple zeros on the critical line** (the architect's `zeta_simple_on_line`,
↪ thm:zeta-mainw2). -/
theorem zeta_simple_on_line (hcert : CertS8) :
  ∀ ε > 0, ∃ T₀ : ℝ, ∀ T ≥ T₀,
    ((6733736895 : ℝ) / 10 ^ 10 - ε) * (Ncount T (2 * T) : ℝ) ≤ N0simple T (2 * T) := by

```

```

/-- **SC --- simple or critical zeros** (the architect's `zeta_simple_or_critical`,
↪ cor:oll-SC). -/
theorem zeta_simple_or_critical (hcert : CertS8) :
  ∀ ε > 0, ∃ T₀ : ℝ, ∀ T ≥ T₀,
    ((8879195 : ℝ) / 10 ^ 7 - ε) * (Ncount T (2 * T) : ℝ) ≤ Nsc T (2 * T) := by

```

The statements of these six theorems are, character for character, those of the six dyadic statements of record in §A.3, with the definitions of the namespace ZetaS in place of those of FollowUpZeta; comparator/Solution/FollowUpZeta.lean derives the fourteen statements of record from them. The library’s definition of CertAM5 (ZetaS/Top/TopDefs.lean) is:

```

/-- **Certificate AM5**: all-marks data with `a₁ = 1280197/10⁸`, `a₂ = 48749/3125000`, `ν =
↪ 1/125` for which (LI_m)
holds for `k_ψ`, `ψ = cos 1.6s` (X2's `round1/X2_numerics/xpat_K5_*.json`,
↪ `claims_d15e6_a1.6_K5_mu500.json`). -/
def CertAM5 : Prop :=
  ∃ W : MarkWeights 5, W.a 0 = 1280197 / 10 ^ 8 ∧ W.a 1 = 48749 / 3125000 ∧ W.nu = 1 / 125
  ↪ ∧
    LocalCertAM (kPsi psiCos16) W

```

A.5 The chain for the $K = 5$ certificate

From ZetaS/Cert/ChainV3.lean: the soundness theorem of the corrected checker and the theorem that twenty checked certificates, one per canonical class, give CertAM5 (the file ends with a diagnostic #eval, omitted):

```

namespace ZetaS.CertV2.ChainV3

open ZetaS.CertV2

/-- **The certificate theorem** (node C21, L5_1's proof), against the V3 checker, all
↪ dependencies proved. -/
theorem cert_check_sound (C : Cert) (h : C.check = true) : C.data.Holds := Cert.check_sound'
↪ C h

/-- **Twenty checked certificates, one per canonical class, give `CertAM5`** (V3 checker,
↪ level A). -/
theorem certAM5_of_checks (C : List Cert) (hdata : C.map Cert.data = classes5.map classData)
  (hchk : ∀ c ∈ C, c.check = true) : CertAM5 :=
  certAM5_of_checks_of_sound cert_check_sound C hdata hchk

end ZetaS.CertV2.ChainV3

```

A.6 The replay’s theorems: the $K = 5$ pair without hypothesis

These are in the library ZetaSReplay of the repository, which holds the 1,478 files of the replay (§7.3) in one flat folder, byte-identical to the files that the kernel checked, and is not

among the default build targets; `lake build ZetaSReplay` repeats the replay (§7.4). The replay compiled all 1,478 files with no failure, in 2 h 50 min of wall time and 20.4 Lean-hours, against the library at commit `bb7223b`, whose Lean files the release changes only in comments (§A.9). From `ZetaSReplay/ReplayAM5.lean`, the statement of `certAM5_replayed` and the first line of its proof (the proof term continues with the twenty theorems `R<class>.cert_ok` proved by the kernel; omitted):

```
namespace ZetaS.CertV2

theorem certAM5_replayed : ZetaS.CertAM5 :=
  ChainV3.certAM5_of_checks [R11111.cert, R11112.cert, R11121.cert, R11122.cert,
    ↪ R11211.cert, R11212.cert, R11221.cert, R11222.cert, R12112.cert, R12121.cert,
    ↪ R12122.cert, R12212.cert, R12221.cert, R12222.cert, R21112.cert, R21122.cert,
    ↪ R21212.cert, R21222.cert, R22122.cert, R22222.cert] rfl
```

From `ZetaSReplay/ReplayHeadlines.lean`, the two theorems of the $K = 5$ pair without hypothesis; their statements are those of `ZetaS.Top.zeta_simple_K5_final` and `zeta_distinct_K5_final` (§A.4) without the hypothesis `hcert`:

```
namespace ZetaS
namespace Top

/-- **Simple zeros, K = 5, hypothesis-free** (thm:sgd-Sigma with the K = 5 certificate). -/
theorem zeta_simple_K5_uncond :
  ∀ ε > 0, ∃ T₀ : ℝ, ∀ T ≥ T₀,
    ((675158622 : ℝ) / 10 ^ 9 - ε) * (Ncount T (2 * T) : ℝ) ≤ Nsimple T (2 * T) :=
  zeta_simple_K5_final ZetaS.CertV2.certAM5_replayed

/-- **Distinct zeros, K = 5, hypothesis-free** (thm:sgd-D with the K = 5 certificate). -/
theorem zeta_distinct_K5_uncond :
  ∀ ε > 0, ∃ T₀ : ℝ, ∀ T ≥ T₀,
    ((837579311 : ℝ) / 10 ^ 9 - ε) * (Ncount T (2 * T) : ℝ) ≤ Ndist T (2 * T) :=
  zeta_distinct_K5_final ZetaS.CertV2.certAM5_replayed

end Top
end ZetaS
```

A.7 The majorant

The majorant is not a hypothesis of any headline theorem: it is proved. Its `Prop` (`ZetaS/ChallengeZetaS.lean`, namespace `ZetaS`), the instance used (`ZetaS/Majorant/CertMajV2.lean`) and the theorem (`ZetaS/Majorant/CertMajV2Proof.lean`, namespace `ZetaS.MajSound`):

```
/-- **`MajorantCert ψ λ`** --- the certified input of lem:sgd-env (1.1283--1296), in the
  ↪ form the proof uses:
  an even integrable B ≥ 0 on ℝ with B ≥ v_ψ := ψ/∫ψ on [-1/2,1/2], whose Fourier
  ↪ transform
  B^(t) = ∫ B(s) cos(2πts) ds vanishes for |t| ≥ 1, and with 2 B^(0) + 4 sup_{[3/4,1]}
  ↪ |B| < λ.
  The draft's instance: B` piecewise linear with nodes j/60, β₀ = 1.5556524027186915,
  ↪ β* = |β₅₂|,
  and λ = λ_c = 2 + √(2 - 2a₁). -/
def MajorantCert (ψ : ℝ → ℝ) (lam : ℝ) : Prop :=
  ∃ B : ℝ → ℝ, MeasureTheory.Integrable B ∧ (∀ s, B (-s) = B s) ∧ (∀ s, 0 ≤ B s) ∧
    (∀ s ∈ Set.Icc (-(1 / 2) : ℝ) (1 / 2), ψ s / (∫ t in (-(1 / 2) : ℝ)..(1 / 2), ψ t) ≤
      ↪ B s) ∧
```

$$\begin{aligned}
& (\forall t : \mathbb{R}, 1 \leq |t| \rightarrow \int s, B s * \text{Real.cos } (2 * \text{Real.pi} * t * s) = 0) \wedge \\
& \exists \beta\text{star} : \mathbb{R}, (\forall t \in \text{Set.Icc } (3 / 4 : \mathbb{R}) 1, |\int s, B s * \text{Real.cos } (2 * \text{Real.pi} * t * s)| \\
& \rightarrow \leq \beta\text{star}) \wedge \\
& 2 * (\int s, B s) + 4 * \beta\text{star} < \text{lam}
\end{aligned}$$

```

/-- **CertMajV2** --- `MajorantCert psiCos16 (33/10)` -/
def CertMajV2 : Prop := MajorantCert psiCos16 (33 / 10)

```

```

/-- **CertMajV2**, instantiated with the C05/C06 node statements. -/
theorem certMajV2_proved : CertMajV2 := certMajV2_of_CNodes cnodes_of_statements

```

ZetaS/Majorant/CertMajProof.lean also proves, as MajSound.certMaj_proved, the first version CertMaj of ZetaS/ChallengeZetaS.lean, with the threshold 3.2063638853; no theorem uses it.

A.8 Axiom census and sorry

At commit bb7223b, `#print axioms` gives exactly [propept, Classical.choice, Quot.sound] for each of the six theorems of §A.4, for ZetaS.CertV2.ChainV3.certAM5_of_checks and ZetaS.CertV2.ChainV3.cert_check_sound, and for all fourteen theorems of comparator/Solution/FollowUpZeta.lean. For the replay’s three theorems of §A.6 it gives the same three axioms, and no open declaration lies beneath them. In the repository, comparator/PrintAxioms/FollowUpZeta.lean prints the axioms of the fourteen theorems of record and comparator/PrintAxioms/FollowUpReplay.lean those of the three theorems of §A.6 (§7.4); the logs and the `#print axioms` outputs of the builds of the release are in audit/followup/.

The library ZetaS has 206 files and 22,751 lines. It contains 3 `sorry`, all in the open node ZetaS/Skeleton/A8g-SigmaDistGeneral.lean: Theorems 5.7 and 5.8 for arbitrary certificate data. That node is used only by the forms with a general majorant threshold, ZetaS.Top.zeta_simple_K5_of_majorant and its three companions (ZetaS/Top/TopSolutionGeneral.lean), and by no headline theorem; the import closure of the fourteen theorems of record contains no `sorry`. There is no `axiom`, `admit` or `native_decide` in the library.

A.9 Repository layout

The formalisation is published in a modified copy of the Lean repository of [2] (zeta-23-lean): <https://github.com/Oliverds321/simple-zeros>, at tag v1.0. The commit of the tag is not printed in this paper; it is recorded in the file RELEASE_NOTES.md of the repository and in the Zenodo record of the release. The repository is a snapshot of the release without its development history: the commits named in this paper (such as bb7223b) are commits of that history. The parts that concern this paper:

- **ZetaS/** with its root module ZetaS.lean: the library of this paper, outside the default build targets (its folders are listed below); ZetaSCertShims/: three forwarding modules through which the checker files are imported by their bare names, so that they are kept byte for byte as checked;
- **ZetaSReplay/**: the 1,478 files of the kernel replay of CertAM5 (1,475 generated shards and the three common modules ReplayCommon, ReplayAM5 and ReplayHeadlines), in one flat folder, as a library outside the default build targets (§7.3, §A.6);

- `comparator/`: the comparator set-up of [2], with seven new files for this paper: `ChallengeDeps/FollowUpZeta.lean` and `Challenge/FollowUpZeta.lean` (quoted above), `Solution/FollowUpZeta.lean`, `PrintAxioms/FollowUpZeta.lean` and `PrintAxioms/FollowUpReplay.lean`, `config-followup.json` and `README_followup.md`;
- `supplementary/`: the data, certifiers and logs of Table 3 and the generator of the replay files (§7.4);
- `papers/zeta/`: the sources and the PDF of this paper; `audit/followup/`: the logs and `#print axioms` outputs of the builds of the release.

The repository also contains the library `ZetaShell/` of a companion paper on families of Dirichlet L -functions, whose sources are in `papers/family/`, and the library `ZetaQ/` added to the repository earlier (see NOTICE); `ZetaS` imports neither. The statement files of [2] (`Zeta23/Statement.lean` and the existing files of `comparator/`) are unchanged. Further, the repository contains changes to 25 files of the library `Zeta23` (3,122 lines added, 121 deleted, between the upstream base commit and `bb7223b`): a weaker parameter condition (`Params.ValidQ` in place of `Params.Valid`), with the proofs adapted, and new modules for Gevrey ramps and a free buffer (`Tail/GevreyTail`, `Taper/GevreyProduct`, `WindowD`). The axiom census above is for the tree with these changes. The Lean files of `Zeta23`, `ZetaS`, `ZetaSCertShims` and `comparator/` in the release are those of commit `bb7223b`, at which the census of §A.8 was taken and against which the replay ran, except for comments: for the release, the copyright headers of the four new comparator files of commit `bb7223b` and stale doc-strings in `comparator/ChallengeDeps/FollowUpZeta.lean` and `ZetaS/ChallengeZetaS.lean` were corrected. Two files were added to `comparator/PrintAxioms/`, `FollowUpReplay.lean` (§A.8) and `FollowUpShell.lean`, for the companion paper, and `comparator/README_followup.md` was brought up to date. The comparator package `FollowUpFamily` of the companion paper (`ChallengeDeps/`, `Challenge/`, `Solution/` and `PrintAxioms/FollowUpFamily.lean`, `config-followup-family.json` and `README_followup_family.md`) was also added. The folders of `ZetaS`:

- `ZetaS/Interfaces.lean`, `InterfacesV2.lean`: the interfaces between the zero side and the local arguments (Gram data, frame families at bandwidth $\lambda < 1$); `ZetaS/ChallengeZetaS.lean`: the library's copies of the definitions of §A.2, together with `MajorantCert` and `CertMaj`;
- `LinAlg/`: Section 3 (stability inequality and its two-parameter form, Φ_m , pinching, monotonicity, pair removal, localisation, inertia);
- `Window/`: the two windows, their exact moments and constants, and the window-general pipeline at bandwidth $\lambda < 1$; `Bandwidth/`: the limit $\lambda \rightarrow 1^-$;
- `ZeroSide/`: the Poisson and Gram identities, positive definiteness, truncation defect, inertia of pair blocks, and the assembly from the zeros of `riemannZeta`;
- `SigmaDist/`: Section 5; `KSide/`: Sections 4 and 6; `Top/`: the numerics of the constants, the robustness conditions, the headline theorems (`TopFinal.lean`, `TopSSC.lean`) and the forms with a general majorant threshold (`TopSolutionGeneral.lean`);
- `Majorant/`: the majorant, its analytic part, its cell checker, the 13 replay shards and the two majorant theorems; `Cert/`: the corrected checker for the $K = 5$ certificate (`CheckerBase.lean`, `CheckerLeaves.lean`, `CheckerCoreV3.lean`, with kernel-checked negative controls in `CheckerTestsV3.lean`), the specification `CertSpecAM5.lean`, the soundness nodes and the chain `ChainV3.lean`;

- `Skeleton/`: re-export files for the proved nodes, and the one open node `A8g_SigmaDistGeneral.lean`.

The new Lean files of `comparator/` carry the author’s copyright notice under the Apache License 2.0, the licence of the repository of [2]; the files that come from [2] keep their own notices (see NOTICE).

B How the work was done

Disclosure of generative-AI use. This paper and its Lean development were produced substantially by a generative AI system, Claude (Anthropic), working under the direction of the named author, who takes full responsibility for its contents. No AI system is listed as an author.

The local inequalities of Section 7 were found by linear programming over gap vectors located by floating-point search and verified by interval branch-and-bound; the certificates were re-run, some box for box, and their inputs and closing constants were checked again with independently written code. What is checked by the Lean kernel, and under which hypotheses, is stated in §1.5. The Lean libraries, the certificate data with their certifiers and logs, the generator of the replay files, and the build logs and `#print axioms` outputs of the release are in the repository <https://github.com/Oliverds321/simple-zeros> at tag v1.0 (Appendix A.9).

Scope of the verification. Theorems 1.1 and 1.2 and Corollary 1.3 are Lean theorems that depend only on the three standard axioms. For the $K = 5$ pair of Theorem 1.1 the local certificate is replayed in the Lean kernel, so these two statements have no hypothesis. For the $K = 7$ pair, Theorem 1.2 and Corollary 1.3 the local certificate is a displayed hypothesis: a finite inequality checked outside Lean by interval arithmetic. The paper has not been refereed, and the Lean development has not been rebuilt outside this project.

References

- [1] ainta. zeta-simple-zeros: research draft (docs/proof.md), 2026. GitHub repository ainta/zeta-simple-zeros, 11 Aug 2026; unrefereed.
- [2] Levent Alpöge and Ralph Furman. More than two thirds of the zeros of the Riemann zeta function are simple and on the critical line. arXiv:2608.13637, 2026. v2, 19 Aug 2026. Lean formalisation: github.com/anthropics/zeta-23-lean, tag v1.0 (the address now serves the repository anthropics/formal-math, folder zeta23/).
- [3] AMTOPA. zeta-exact-pressure. GitHub repository AMTOPA/zeta-exact-pressure, 2026. 12 Aug 2026 (commit dates); unrefereed research draft; claims 0.673416490971.
- [4] AxiomMath. ZetaZeros. GitHub repository AxiomMath/ZetaZeros, 2026. Commits of 2 Sep 2026; a Lean formalisation of Lamzouri’s proof (arXiv:2609.02882), with the Riemann–von Mangoldt formula and a pair-correlation asymptotic as stated hypotheses.
- [5] H. M. Bui and D. R. Heath-Brown. On simple zeros of the Riemann zeta-function. *Bull. London Math. Soc.*, 45:953–961, 2013.
- [6] Emanuel Carneiro, Vorrapan Chandee, Friedrich Littmann, and Micah B. Milinovich. Hilbert spaces and the pair correlation of zeros of the Riemann zeta-function. *J. reine angew. Math.*, 725:143–182, 2017.
- [7] Andrés Chirre, Felipe Gonçalves, and David de Laat. Pair correlation estimates for the zeros of the zeta function via semidefinite programming. *Adv. Math.*, 361:Paper No. 106926, 2020. arXiv:1810.08843.

- [8] Leonardo de Moura and Sebastian Ullrich. The Lean 4 theorem prover and programming language. In *Automated Deduction – CADE 28*, volume 12699 of *Lecture Notes in Computer Science*, pages 625–635. Springer, 2021. https://doi.org/10.1007/978-3-030-79876-5_37.
- [9] Michael Devine. A computation-free spectral improvement to the Montgomery–Taylor bound for simple critical zeros. Zenodo record 22287432, 2026. v1.0.0, 3 Sep 2026; unrefereed.
- [10] Michael Devine. An unconditional 67.3399% bound and conditional advances beyond 67.92% for simple critical zeros of the Riemann zeta function. Zenodo record 22066689, 2026. v1.0.3, 23 Aug 2026; unrefereed; claims $673399/10^6$.
- [11] David W. Farmer. Counting distinct zeros of the Riemann zeta-function. *Electron. J. Combin.*, 2, 1995. Research Paper R1.
- [12] Jonas Jakob Gebendorfer. An expanded window and joint position pressures for simple critical-line zeros. Zenodo record 22643957, 2026. v1, 7 Sep 2026; v2 (record 22646541) of the same day with the same constant; described by the author as an internally checked research draft; claims 0.6734760237.
- [13] gfreund123. Pull request #888 (draft) to the repository GettysburgResearch/riemann. GitHub, 2026. Draft of 14 Sep 2026; claims 0.67349239 as a proposed bound; the pull request states that independent analytic review is required and that “no independently reviewed world-record claim is made”.
- [14] Haseo Ki and Yoonbok Lee. Zeros of the derivatives of the Riemann zeta-function. *Funct. Approx. Comment. Math.*, 47(1):79–87, 2012. Cited after Lamzouri; not consulted by us.
- [15] Youness Lamzouri. A new proof that more than $2/3$ of the zeros of the Riemann zeta function are simple and on the critical line. arXiv:2609.02882, 2026. v2, 8 Sep 2026.
- [16] MichaelMobiis. `simple_zeros_of_the_riemann_zeta_function`. GitHub repository MichaelMobiis/simple_zeros_of_the_riemann_zeta_function, 2026. Latest commits 21–22 Sep 2026; unrefereed; claims 0.6731175265, announced with a hypothesis-free Lean proof; the README names Michael Hurtado as author [V~].
- [17] npip99. `zeta-zeros`. GitHub repository npip99/zeta-zeros, 2026. Commits of 11–12 Aug 2026; unrefereed; claims 0.673195198901, with a Lean reduction to one local-certificate premise.
- [18] RakitinD. `zeta-xi-zero-results`. GitHub repository RakitinD/zeta-xi-zero-results, 2026. Commits of 14–15 Sep 2026; unrefereed; claims 0.6734775921119; the README names Dmitry Rakitin as author [V~].
- [19] Santibáñez-Leal. Zenodo record 22940291. Zenodo, 2026. v0.01, 24 Sep 2026; unrefereed; claims 0.6725007995946757558 (simple, on the line) and 0.8362503997973378779 (distinct).
- [20] sxuff. `zeta-positioned-pressure`. GitHub repository sxuff/zeta-positioned-pressure, 2026. 11 Aug 2026; unrefereed; claims 0.6732059784228011963.
- [21] tawanerguo-cn. `zeta-simple-zeros`. GitHub repository tawanerguo-cn/zeta-simple-zeros, 2026. 11 Aug 2026 (commit dates); unrefereed; claims 0.6731929114731422.
- [22] teal-sea. `zeta-lab`. GitHub repository teal-sea/zeta-lab, 2026. 2026; unrefereed; claims bounds 0.6727373345 and 0.6728470197 with certificates proved in Lean.

- [23] The mathlib Community. The Lean mathematical library. In *Proceedings of the 9th ACM SIGPLAN International Conference on Certified Programs and Proofs (CPP 2020)*, pages 367–381. ACM, 2020. <https://doi.org/10.1145/3372885.3373824>.
- [24] trmdy. zeta-simple-zeros-673137. GitHub repository trmdy/zeta-simple-zeros-673137, 2026. First public 11 Aug 2026 (0.673137630699); unrefereed; claims 0.673312742272 since 12 Aug 2026.
- [25] uwe-schwarz. zeta-simple-zeros-673026. GitHub repository uwe-schwarz/zeta-simple-zeros-673026, 2026. Date not recorded; unrefereed; claims 0.6730266625438475.
- [26] vstaln. riemann: multi-agent research ledger, 2026. GitHub repository vstaln/riemann, 2026; unrefereed.
- [27] Biao Wang. Proportions of the non-trivial zeros of the Riemann zeta function. arXiv:2609.24167, 2026. v1, 21 Sep 2026.
- [28] Xiaosheng Wu. Distinct zeros of the Riemann zeta-function. arXiv:1206.3737, 2012. v2, 18 Sep 2012.
- [29] Hongyi Yang and Shihua Yang. More than 79.62% of the zeros of the Riemann zeta function are simple and on the critical line. Zenodo record 21975237; reproduction repository JoshuaHKU/zeta-0.7947-reproduction on GitHub, 2026. v1.0, 17 Aug 2026; unrefereed; claims 0.7962 for simple zeros on the critical line and 0.8981 for distinct zeros, with ineffective constants; declared by the authors “at the certified-candidate level, not as an established theorem”.
- [30] Shihua Yang and Hongyi Yang. Almost all zeros of the Riemann zeta function are simple and on the critical line. Zenodo record 22065921, 2026. v1.0, 22 Aug 2026; unrefereed; claims that almost all zeros are simple and on the critical line, with ineffective constants, and the explicit bounds 0.910460 (simple, on the line) and 0.955230 (distinct); the record describes the analytic chain as “a certified candidate pending external review”; files under embargo until 31 Dec 2045, the record giving as the reason that the version contains errors and needs revision.
- [31] yuhangshi888. zeta-simple-zeros-673316977. GitHub repository yuhangshi888/zeta-simple-zeros-673316977, 2026. v0.1.0, 29 Aug 2026; unrefereed; claims 0.6733169771424713; the README names Yuhang Shi as author [V~].
- [32] Zhaoyi-Tian. zeta-simple-zeros-lean. GitHub repository Zhaoyi-Tian/zeta-simple-zeros-lean, 2026. Commits of 31 Aug and 3 Sep 2026; unrefereed; announces a Lean theorem at 6734/10000 (`dyadic_target`).
- [33] Zhongshan-Big-Jun. Riemann-Conjecture-Research, 2026. GitHub repository Zhongshan-Big-Jun/Riemann-Conjecture-Research, 2026; unrefereed.