

Highlights

The numerical layer of the two-thirds theorem, checked without Lean: a certificate that waits on a file, and the ξ' constants raised by an exact window and a longer tail*

Carlos Toledo

- The formalisation's one numerical certificate holds over its printed integers by a route that shares nothing with Lean; its hypothesis `EnclOK` rests on an unpublished file, which we ask for.
- The Montgomery–Taylor and ξ' constants are decided by exact rational arithmetic, no interval kernel; the ξ' rationals re-derive from their definitions.
- The ξ' window optimised exactly: the formalisation's constants rise from 0.86864017 and 0.93432008 to 0.86864150 and 0.93432075 by a degree-six window and fifteen series terms, with a certificate that no even polynomial window of degree up to twenty does better.
- With the ζ kernel the same machinery returns the Montgomery–Taylor constant to 18 digits, and exactly 0.6725 at degree two: the calibration against a closed form.
- A truncated print is not a rounded one: the abstract's 0.8362 is true as a lower bound and false as a rounding.
- Of 65 sentences the story's sources print, 10 reduce to decidable claims; every theorem is refused, with the rule.

The numerical layer of the two-thirds theorem, checked without Lean: a certificate that waits on a file, and the ξ' constants raised by an exact window and a longer tail

Carlos Toledo^{a,*}

^aIndependent researcher, the cert-machine project, Brazil

Abstract

In 2026 a proof that more than two thirds of the nontrivial zeros of the Riemann zeta function are simple and on the critical line was found by a language model, verified by Alpöge and Furman, and formalised in Lean. The formalisation has one numerical input: 256 integer enclosures of a form factor at scale 2^{140} , kernel-checked downstream of a hypothesis, EnclOK, that the kernel does not check and that rests on a law file the release describes as available from the authors. We re-check the certificate from its integers with two programs that share no code with Lean, state exactly what that establishes and what it does not, and ask for the file. On the way, the Montgomery–Taylor constant and the four ξ' constants of the paper are decided by exact rational arithmetic alone, with no interval kernel. The window of the paper’s Remark 7.1 is then optimised exactly: the constants depend on the window through a Rayleigh quotient, the paper’s quartic sits $8.85 \cdot 10^{-7}$ below its own family’s optimum, and a degree-six window with the certificate’s series truncated at fifteen terms instead of nine raises the formalisation’s eight-digit constants from 0.86864017 and 0.93432008 to 0.86864150 and 0.93432075, with an exact certificate that no even polynomial window of degree up to twenty does better at $\lambda = 1$; the same machinery with the ζ kernel recovers the Montgomery–Taylor constant to 18 digits, which calibrates it. The paper’s abstract prints one constant truncated where a rounding would print another digit, true as a lower bound; and a survey of everything the story’s sources print—the manuscripts, the Lean files, the research post, a video’s captions—finds that 10 of 65 sentences are finite claims at all. No theorem is touched. Every number here is read from a

*Draft v0.3 of 2026-10-22, a companion to the cert-machine method paper, built from repository commit 4df5fe2f7b06. Every number and count is read at build time from the records named in the Data availability statement; a record that no longer supports a sentence stops the build. Nothing here is a theorem, and nothing here bears on the truth of the theorems the story is about.

*Corresponding author.

Email address: `carlos@carlostoledo.co` (Carlos Toledo)

public record at build time and re-derived by a third implementation before it is printed.

Keywords: Riemann zeta function, simple zeros, window optimisation, Lean, formal verification, independent verification, exact rational arithmetic, machine-generated mathematics

1. The certificate, and what it rests on

Alpöge and Furman (Alpöge and Furman, 2026) prove that at least two thirds of the nontrivial zeros of ζ , counted with multiplicity, are simple and on the critical line, and that at least five sixths are distinct; with the Montgomery–Taylor window (Montgomery, 1973; Carneiro et al., 2017) the constants become $2 - c_{\text{MT}}^{-1} = 0.67250\dots$ and $\frac{1}{2}(3 - c_{\text{MT}}^{-1}) = 0.83625\dots$. The argument was discovered and written by Claude, an AI developed by Anthropic; the authors verified it and take responsibility for it, as the paper’s own dedication says. Anthropic formalised the result in Lean (Anthropic, 2026b) and described the work in a research post (Anthropic, 2026a); a video explained the checking to a wide audience (Adem B10, 2026). The formalisation’s audit prints the axioms of every main statement; they are Lean’s three standard ones.

One part of the formalisation is numerical. The bandwidth-one ceiling, `Zeta23/PairCeiling/`, is the stability inequality behind the paper’s remark on the optimality of its method: a bound on the simple-point fraction of any law of 256-periodic marked configurations whose form factor is near the CUE values on the band. Its input is the generated file `Zeta23/PairCeiling/LawN256.lean` (26,084 bytes, sha256 `11ae9f7dcc17`), whose header states that it was computed by 70-digit interval arithmetic at scale $K = 2^{140}$ and holds enclosures $\text{lo}_j \leq K \cdot S(j) \leq \text{hi}_j$, $j = 1, \dots, 256$, of the grid form factor of a law recorded in a certificate file, `cert_N256_blk_b128m.json` (sha256 `cc3de9917db4`), “available from the authors”. The enclosures are 43-digit integers, every one exactly 1 unit wide. The row certificate they satisfy is

$$|256 S(j) - j| \leq \tau := 3 \cdot 10^{-40} \quad (0 < j < 256), \quad |D(1)| \leq 0.82395317, \quad (1)$$

with the closed-band row $j = 256$ free; the law's simple-point fraction is the exact rational $p_0 = 10909258999421303588095230195816054408197/1600000000000000000000000000000000000000$, that is $0.68182868746\dots$, which the paper prints as $p_0 \leq 0.6818287$.

The Lean side is as the audit prints it. `LawN256_check`, the kernel check of (1) over the enclosures, depends on the axiom `propext` alone. `ceiling_law256` and the other ceiling theorems depend on the three standard axioms and, as the audit says of all of them except the two kernel

checks, carry the displayed hypothesis `Enc10K`: that the true form factor of the law lies in the enclosures. That hypothesis is about the law itself, the weights and positions in `cert_N256_blk_b128m.json`, and the release does not publish that file. The audit is exemplary in saying so. But it means that the one place where the formalisation’s numerical certificate meets the real numbers is checked by interval arithmetic outside Lean, from bytes no one outside can read.

2. What was checked without Lean

We decoded the integer literals of `LawN256.lean` into a JSON file beside it, naming the Lean file’s digest, and decided three claims over those integers with two programs, one in Python and one in JavaScript, that share no code with each other or with Lean and evaluate in exact rational arithmetic (Table 1; the programs are described in Section 6):

- the band rows, $|256 S(j) - j| \leq \tau$ over the 255 rows, with both ends of each enclosure tested: the worst row sits at $1.836 \cdot 10^{-40}$ from its target against the tolerance $3 \cdot 10^{-40}$;
- the edge bound, $|D(1)|$ with $D(1) = \sum_{j \leq 256} S(j)/256 - \frac{1}{2}$ evaluated at both ends of the enclosures, which is exact because $D(1)$ is monotone in each $S(j)$: $|D(1)|$ at the ends is $0.82395316071\dots$, under the printed bound by $9.28 \cdot 10^{-9}$;
- p_0 , the file’s rational against the printed bar.

All three hold, in milliseconds. Beside them, three pieces of the paper’s own arithmetic on the ceiling were decided from the file’s p_0 and τ : $p_0 - \frac{2}{3} = 0.0151620207\dots$ is within the paper’s 0.016; the decimal constant the Lean prints beside the ceiling, $1/(6 \cdot 256^2) + \tau/512 = 2.5431315104 \cdot 10^{-6}$, is under the paper’s rounded $2.55 \cdot 10^{-6}$; and the paper’s bound for any window with $|r'(1)| + \int |r''| < 8.2$, $p_0 + 2.55 \cdot 10^{-6} \cdot 8.2 = 0.68184959746\dots$, is under its printed 0.6819.

A property of the file. Every band enclosure contains $j \cdot K/256$ exactly: 124 rows at the upper end, 131 at the lower, counted from the decoded integers and decided as $|2jK/256 - \text{lo}_j - \text{hi}_j| \leq 1$ over the 255 rows. The law’s form factor is thus consistent with $S(j) = j/256$ on every band frequency to 2^{-140} : the enclosures do not merely lie within $3 \cdot 10^{-40}$ of the CUE values, they sit on them, and $D(1) = S(256)/256 - 1/512$ to the same precision, with $256 S(256) = 54126.5943404763\dots$. This is most likely so by construction, an optimal law of a linear programme whose constraints are the CUE band values being expected to meet them exactly, and it is reported as a property of the file, not as a finding about the method.

claim, as printed	source	decided by	word	s
$ 256 \cdot S(j) - j \leq 3 \cdot 10^{-40} \ (0 < j < 256)$	zeta23/ README.md	the 255 band rows, each end against the tolerance	CERTIFIED	0.014
$ D(1) \leq 0.82395317$	zeta23/ PairCeiling/ CeilingLaw256.lean	$ D(1) $ at both ends of the enclosures	CERTIFIED	0.003
$p_0 \leq 0.6818287$	arXiv/paper- v5-draft18.tex	the exact rational the file prints against the bar	CERTIFIED	0.001
$c_1^* = 0.7532960 \dots$	zeta23/ ChallengeDeps.lean	the series certificate, $K = 8$	CERTIFIED	0.003
$2 - 1/c_1^* = 0.67250 \dots$	zeta23/ formalization.yaml	the series certificate	CERTIFIED	0.002
$(3 - 1/c_1^*)/2 = 0.83625 \dots$	zeta23/ formalization.yaml	the series certificate	CERTIFIED	0.002
$2 - c_{MT}^{-1} = 0.67250 \dots$	arXiv/paper- v5-draft18.tex	kept on the closed form: transcendental head, admitted, not decided	undecided	0.0
EnclOK	zeta23/ PairCeiling/ LawN256.lean	no bytes, no claim	skipped	—

Table 1: The certificate run: 8 claims, 6 decided and holding (CERTIFIED), 1 admitted and undecided (the paper’s own closed form, kept on that route on purpose; see Section 3), 1 with no bytes to decide from.

What this establishes, and what it does not. The three checks are decided from the same integers the Lean kernel accepted; the independence is in the checking code, not in the data. They establish that the row certificate (1) holds over the printed enclosures by a route that does not pass through Lean, its elaborator or its `decide` tactic, a no-shared-code concordance with `LawN256_check`. They do not establish that the enclosures enclose anything. That is `EnclOK`, and it waits on the law file: there is no reading of any published byte that decides it, because the bytes it is about are not there.

The request. The file had been asked for on the formalisation’s repository twice before, in August 2026 (its issues 8 and 14), without an answer; on 2026-10-22 we joined that request as a comment on issue 8, with the re-check above and the ask for the interval-arithmetic script as well. With the file, the enclosures can be re-derived by an independent interval arithmetic from the law’s weights and positions, and the certificate becomes checkable end to end by exactly the pipeline described here with one more claim. Until then it is a conditional certificate, for the kernel and for us alike.

3. Two families of constants, decided by rational arithmetic

The Montgomery–Taylor constant. The formalisation defines $c_1^* = \sqrt{2} \tan \vartheta / (1 + \vartheta \tan \vartheta)$ at $\vartheta = 1/\sqrt{2}$, the reciprocal of the paper’s

$c_{\text{MT}}^{-1} = \frac{1}{2} + \frac{1}{\sqrt{2}} \cot \frac{1}{\sqrt{2}}$. An exact checker over rationals cannot evaluate a tangent, and an interval kernel would be the usual answer. None is needed. With $\vartheta^2 = \frac{1}{2}$ the Maclaurin series give $\sin \vartheta = \vartheta s$ and $\cos \vartheta = c$ for

$$s = 1 + \sum_{k \geq 1} \frac{(-1)^k}{2^k(2k+1)!}, \quad c = 1 + \sum_{k \geq 1} \frac{(-1)^k}{2^k(2k)!}, \quad (2)$$

alternating series of rationals whose terms fall by a factor of at least twelve at every step; hence $\tan \vartheta = \vartheta s/c$ and, since $\sqrt{2} \vartheta = 1$,

$$c_1^* = \frac{2s}{2c+s}, \quad 2 - \frac{1}{c_1^*} = \frac{3}{2} - \frac{c}{s}, \quad \frac{1}{2} \left(3 - \frac{1}{c_1^*} \right) = \frac{5}{4} - \frac{c}{2s}. \quad (3)$$

For an even truncation order K the partial sums bracket the limits, $s_{K+1} \leq s \leq s_K$ and $c_{K+1} \leq c \leq c_K$, and $2s/(2c+s)$ increases in s and decreases in c on positive values, so $2s_{K+1}/(2c_K + s_{K+1}) \leq c_1^* \leq 2s_K/(2c_{K+1} + s_K)$. At $K = 8$ this is the enclosure $[0.75329606785607067702, 0.75329606785607067722]$, of width $1.96 \cdot 10^{-19}$, in sums, products and powers over integer ranges. The three prints of the sources, $c_1^* = 0.7532960\dots$, $2 - 1/c_1^* = 0.67250\dots$ and $(3 - 1/c_1^*)/2 = 0.83625\dots$, hold as truncation boxes, and so do the two bounds one unit of the eighth digit either side, $c_1^* \leq 0.75329607$ and $c_1^* \geq 0.75329606$. This is elementary; the point is only that it removes the kernel, and that the paper's own print of $2 - c_{\text{MT}}^{-1}$, kept on the closed-form route on purpose, stands admitted and undecided in Table 1 to show what the checker cannot do without it.

The ξ' constants. The paper's Remark 7.1 states that at least 0.85838 of the zeros of ξ' are simple and on the critical line and at least 0.92919 are distinct (flat window), 0.86864 and 0.93432 with the quartic window $v(s) = 1 - \frac{7}{100}(2s)^2 - \frac{51}{200}(2s)^4$; the formalisation proves eight-digit versions, 0.85838371, 0.92919185, 0.86864017, 0.93432008, from a certificate every number of which is rational: the diagonal density $D_1(s) = s - 4s^2 + \sum_{k \geq 0} a_k s^{2k+3}$ with $a_k = 2 \cdot 4^{k+1} k! / (2k+2)!$, truncated at nine terms ($k = 0, \dots, 8$) with the certified tail $\varepsilon_9 = a_9 \cdot 231/211 = 3.424 \cdot 10^{-7}$ (the coefficient ratio is at most $20/231$ from $k = 9$ on); the window integral $J(v) = 2 \int_0^1 D_1^{(9)}(r) (v*v)(r) dr$ with $v*v$ the window's autocorrelation; $\kappa_9(v) = (\int v^2 + J(v)) / (\int v)^2$; and the proportions $2 - \kappa_9 - \varepsilon_9$ and $\frac{3}{2} - (\kappa_9 + \varepsilon_9)/2$. We re-derived every number from those definitions in exact fractions, the quartic window's degree-nine autocorrelation polynomial included, twice and independently (once in the translator's session, once in the generator of this paper's numbers, which stops the build if any differs), and each equals the Lean's computer-algebra-

generated closed form:

$$\begin{aligned}
J_{\text{flat}} &= 12517210259/88388425125, & \varepsilon_9 &= 1024/2990212875, \\
\kappa_9^{\text{flat}} &= 100905635384/88388425125, \\
J_{\text{quartic}} &= 29965280330934234270211/285991090937677125000000, \\
\kappa_9^{\text{quartic}} &= 277244140547469154168336/245053976636191319722125.
\end{aligned}$$

The four eight-digit constants hold with margins $3.59 \cdot 10^{-9}$, $6.79 \cdot 10^{-9}$, $3.59 \cdot 10^{-9}$ and $6.79 \cdot 10^{-9}$; the paper’s five-digit prints hold; $P_{\text{flat}} = 16008819950038/18649957701375 = 0.8583837135 \dots$ exactly, and a control one unit of the eighth digit above it, $P_{\text{flat}} \geq 0.85838372$, is refuted as it must be. (The ξ' certificate’s Lean files—**Statement**, **Defs**, **Certificate** and its **D1**, **Poly** and **AtOne**—are pinned by digest with the other sources; the identity claims quote the derivation note as the line that prints them.)

Two corrections the machine made. Both re-derivations began wrong, and both times the exact decision was the correction. The grammar’s first reading gave a print cut off with an ellipsis, $0.7532960 \dots$, the half-unit rounding box; the checker then refuted the true print and certified a false $0.7532961 \dots$, and it was right, since $c_1^* = 0.75329606785 \dots$ lies in $[0.7532960, 0.7532961)$ and not within half a unit of 0.7532961 . An ellipsis is a truncation, and the grammar now reads it so. And a first pass took “**Ditrunc 9**” as ten terms and reproduced none of the printed rationals; nine terms reproduce all of them, and the generator checks that a ten-term reading still matches none. We report both because they are the kind of error an exact decision catches and a careful reader does not.

4. The ξ' window, optimised exactly, and the tail recovered

The constants are a Rayleigh quotient of the window. For an even polynomial window $v(s) = \sum_{i=0}^{m-1} c_i s^{2i}$ on $[-\frac{1}{2}, \frac{1}{2}]$, every quantity in $\kappa_9(v) = (\int v^2 + J(v))/(\int v)^2$ is a rational form in the coefficients: $\int v = b^\top c$ with $b_i = \int s^{2i} = 1/(4^i(2i+1))$, and $\int v^2 + J(v) = c^\top M c$ with $M_{ij} = \int s^{2i+2j} + \frac{1}{2}(H_{ij} + H_{ji})$, $H_{ij} = 2 \int_0^1 D_1^{(K)}(r) (s^{2i} \star s^{2j})(r) dr$, where the pair autocorrelation $(s^{2i} \star s^{2j})(r) = \int_{-1/2}^{1/2-r} t^{2i} (t+r)^{2j} dt$ is a polynomial in r . So $\kappa_9 = c^\top M c / (b^\top c)^2$, minimised over each degree exactly at $c \propto M^{-1} b$ with minimum $\kappa^* = 1/(b^\top M^{-1} b)$: rational linear algebra, no numerics. The formalisation’s own comment on its quartic window says it “saturates the numerically optimal window to 10^{-6} ”; this section makes that exact, and then takes the two cheap steps the exact picture shows.

The optima by degree and truncation order. Table 2 gives the minimum of $2 - \kappa^* - \varepsilon_K$ over even polynomial windows of each degree, for the certificate’s

nine-term truncation of D_1 and for three longer ones. At nine terms the paper’s quartic, $0.868640173597\dots$, is $8.85 \cdot 10^{-7}$ below the optimum of its own degree, and the optimum saturates from degree six on: twelve digits agree up to degree twenty, $9.853 \cdot 10^{-7}$ above the paper’s window. The larger lever is in the rows: the nine-term truncation costs $\varepsilon_9 = 3.42 \cdot 10^{-7}$, more than a third of the whole window gain, and the formalisation’s truncation lemma is stated for every $K \geq 9$ with the tail $\varepsilon_K = a_K \cdot \frac{(K+2)(2K+3)}{(K+2)(2K+3)-2(K+1)}$, the geometric comparison from the coefficient ratio at $k = K$, which decreases. At fifteen terms $\varepsilon_{15} = 4.52 \cdot 10^{-14}$ ($= 8192/180952685768570625$) and the constants move by the whole of ε_9 ; at twenty terms nothing moves in twelve digits.

terms	ε_K	2	4	6	20
K = 9	$3.42 \cdot 10^{-7}$	0.868200285550	0.868641059186	0.868641158941	0.868641158956
K = 12	$1.71 \cdot 10^{-10}$	0.868200626893	0.868641400587	0.868641500343	0.868641500358
K = 15	$4.52 \cdot 10^{-14}$	0.868200627064	0.868641400758	0.868641500514	0.868641500529
K = 20	$1.59 \cdot 10^{-20}$	0.868200627064	0.868641400758	0.868641500514	0.868641500529

Table 2: The most the window method gives for simple zeros of ξ' at $\lambda = 1, 2 - \kappa^* - \varepsilon_K$, within even polynomial windows of degree 2, 4, 6 and 20 (the columns; degrees 8 to 20 agree to twelve places), by the truncation order of D_1 (the rows); every entry an exact rational printed to twelve places. The paper’s quartic gives 0.868640173597 .

The certified window. The degree-six optimum rounded to thousandths,

$$v_6(s) = 1 - \frac{157}{500}s^2 - \frac{3763}{1000}s^4 - \frac{122}{125}s^6,$$

satisfies the formalisation’s window profile by inspection of its coefficients: every coefficient after the first is negative, so $v_6 \leq 1$ on $\mathbb{R}$ and v_6 decreases in $|s|$, and $v_6(\frac{1}{2}) = 10737/16000 > 0$; both facts are decided from the certificate’s bytes. With the nine-term certificate it gives $P \geq 0.86864115$ and $D \geq 0.93432057$ (0.868641158927 and 0.934320579463 exactly), one unit of the sixth digit above the Lean’s 0.86864017 and 0.93432008 ; with fifteen terms, $P \geq 0.86864150$ and $D \geq 0.93432075$ (0.868641500500 and 0.934320750250), a gain of $1.330 \cdot 10^{-6}$ over the formalisation’s constants, by the same window shape and three more terms of a series it already holds. The controls one unit of the eighth digit beyond are refuted, as they must be. The paper’s five printed digits, 0.86864 and 0.93432 , do not move.

The optimality, decided. That no window of a degree does better is the statement that $N := M - \kappa^*bb^\top$ is positive semidefinite, with exactly one zero eigenvalue at c^* . The certificate prints the pair autocorrelations as coefficient rows, from which the checker recomputes M exactly, together with L , D and c^* for $N = LDL^\top$; the claims are that the factorisation’s residual is zero, that every pivot is at least zero, that the least pivot is not

positive (the one zero), that $\kappa_9(c^*) - \kappa^*$ is zero and not negative (attainment), and that P^* lies in an eleven-digit bracket. All hold, and the two controls that must fail do, at degrees six and twenty, at nine and at fifteen terms. So “no even polynomial window of degree at most twenty gives the method more than 0.868641500529 at $\lambda = 1$ ” is a decision of the checker, not a sentence. The generator of this paper’s numbers re-derives every one of these certificates from the definitions alone, in a third implementation over arbitrary-precision rationals: the autocorrelation rows, M and b , the solve, κ^* and the pivots, and it stops the build if any differs from the record.

The parameter. The paper’s Theorem 8.2 holds at every fixed $\lambda \in (0, 1)$ with $\kappa_1(\lambda, v) = 1/c_\lambda(v; D_1)$, $c_\lambda(v; D) = \lambda(\int v)^2/(\int v^2 + \lambda\mathcal{J})$ and $\mathcal{J} = 2 \int_0^1 D(\lambda r)(v \star v)(r) dr$, the limit form at $\lambda = 1$ stated separately; the functional at λ is a Rayleigh quotient too, with $M = G/\lambda + \text{sym } H(\lambda)$. For every window and degree tried the constant falls as λ leaves 1 (Table 3); on record, v_6 at $\lambda = 19/20$ gives $P \leq 0.83686$ and at $9/10$ $P \leq 0.79415$, with the control $P \geq 0.84$ refuted. The limit form is the right place, and there is nothing to gain in that direction. This is closed by evaluation at the sampled values, not by a monotonicity proof.

λ	degree 6	degree 20
1	0.8686415005	0.8686415005
0.99	0.8631965314	0.8631965314
0.95	0.8377283095	0.8377283095
0.9	0.7970315955	0.7970315955
0.8	0.6825441347	0.6825441347
0.5	-0.0520572078	-0.0520572078

Table 3: The most the method gives at fixed λ , fifteen terms, within even polynomial windows of degrees six and twenty.

The calibration. With the ζ kernel $D(s) = s$ in place of D_1 the same functional is the Montgomery–Taylor problem, whose optimum over all windows is the paper’s closed form $2 - c_{\text{MT}}^{-1} = 0.6725007036794116\dots$ (the series bracket of Section 3), attained at the window $\cos(\sqrt{2}s)$: the Euler–Lagrange equation of the quotient, $v'' + 2v = 0$. Polynomial windows must stay at or below it and approach it. They do (Table 4): degree two gives exactly $2 - 531/400 = 269/400 = 0.6725$, the abstract’s four digits to the digit; degree four gives $0.6725007036667\dots$; from degree six on, $0.672500703679411\dots$, inside the series bracket of width $2 \cdot 10^{-19}$, the closed form to 18 digits, and the distinct proportion $0.8362503518397058\dots$ likewise. The optimality certificates for ζ at degrees two, six and twenty are decided by the same route, controls refuted. This is the calibration of the whole machinery against a known answer, and it says what the saturation

in Table 2 means: for ζ the polynomial optima reach the true optimum because the optimal window is analytic and the quotient’s error is quadratic in the approximation error; for ξ' the saturation from degree six on is the same behaviour, and the L^2 optimum is not certified here.

degree	$2 - \kappa^*$ (simple, on the line)	$\frac{3}{2} - \kappa^*/2$ (distinct)
2	0.672500000000000	0.836250000000000
4	0.672500703666743	0.836250351833371
6	0.672500703679411	0.836250351839705
10	0.672500703679411	0.836250351839705
20	0.672500703679411	0.836250351839705

Table 4: The ζ kernel: the polynomial optima of the Montgomery–Taylor functional by degree, against the closed form 0.6725007036794116... and 0.8362503518397058....

What this is, and what it is not. It is a certified improvement of the formalisation’s ξ' constants by $1.330 \cdot 10^{-6}$, within the paper’s own method and hypotheses, with an exact certificate of optimality over polynomial windows of degree up to twenty and the tail taken three terms further on the formalisation’s own route; and it is a calibration of the machinery that produced it. It is not a new digit in the paper’s print, not an optimum over all windows, and not a Lean proof: a Lean instance needs the truncation lemma at $k \geq 15$ (the coefficient ratio is at most $32/561$ there, beside the $20/231$ of $k \geq 9$) and the four lemma instances the file `AtOne.lean` proves for the quartic—its two integrals, the closed form of $v \star v$, positivity—with v_6 ’s rationals. Everything a formaliser needs is printed in the certificates.

5. A print, and a rule

The paper’s abstract prints: “With the Montgomery–Taylor window the constants become 0.6725 and 0.8362.” Its theorem prints 0.67250... and 0.83625.... The exact second constant is 0.836250351839705822..., which lies above the box the four printed digits allow as a rounding, $[0.83615, 0.83625]$: 0.8362 is a truncation, and a rounding would print 0.8363. As a lower bound on a proportion the truncated print is still true, a weaker true statement, and the paper says nothing else about it. The rule it illustrates is general: a print cut off with an ellipsis and a print rounded to four digits are different claims, and a reader with exact arithmetic will tell them apart. In our grammar the two are read differently, and Table A.10 in the appendix shows the claim “ $(3 - 1/c_1^*)/2$, published as 0.8362” refuted under the rounding reading and its companion 0.6725 holding.

6. How the claims were read, and how much of the story is finite

The claims above were not typed by hand. They were extracted from the sources by a translator with a closed grammar, `claimbox` (Toledo, 2026b), and decided by an exact checker, `cert-machine`, whose grammar of three verdicts with witnesses is described in the method paper this is a companion to (Toledo, 2026a, the method paper). The translator is a standard-library pattern matcher, with no model and no parser framework. A sentence is read only if it matches one of a short list of shapes: a trailing finite range; `|EXPR| REL NUMBER`; `EXPR = NUMBER...`, a value printed to its digits and cut off; `NAME REL NUMBER`; a constant “published as” a number; an exact identity `NAME = RATIONAL`; a few others for other corpora. The number is kept as the source spells it and read as the exact rational it means. Everything else is refused with a rule: an asymptotic, an unbounded quantifier, an approximation sign, two relations in one sentence, no printed quantity, prose. The object the translator emits carries the line of the source that prints the claim with the file’s digest, the pin of the bytes and a sentence saying how they were decoded, the predicate, the hypotheses under which the predicate is the claim, and the shape a witness against it would take; the checker resolves the pin, reads the data as stated, evaluates the predicate over exact rationals and prints the word with a witness. Two checkers, in Python and in JavaScript, share no code and agree on every object here. Neither runs the claimant’s code, reads a proof, or evaluates a transcendental function: a closed form over `exp`, `log`, `sin` or `cos` is admitted and left undecided.

Put to the whole story, the grammar is a measuring instrument. The 65 sentences the sources print were extracted byte for byte, tagged by author, and read: 10 match a shape and 55 are refused (Table 5), 43 as prose with no claim shape, 6 with no printed quantity, 3 as ambiguous, 2 as asymptotic, 1 as an approximation. Every theorem is among the refused: a statement about the limit inferior of a ratio of counting functions is not a finite claim, and the grammar says so rather than inventing one. The 10 read are the certificate’s inequalities and the digits prints of Sections 2 and 3, and nothing else. The 11/12 companion manuscript of the `openai/math` release (OpenAI, 2026) (784,503 bytes, sha256 `f919b57829b1`; 3,073 lines by `pdftotext` 26.07.0) states 34 numbered theorems, lemmas, propositions and corollaries and prints no decimal constant at all; it is a proof for human readers, and nothing in it is a claim this pipeline can touch. The same is true of eighty phrasings of 8 of the story’s claims written by a fresh AI agent in the voices of a paper, an abstract, a README, a Lean docstring, a blog, a tweet, a referee report and a chat assistant: 3 read, 77 refused, because a claim is read at the head of its sentence or not at all, and nothing in the grammar was extended for a phrasing no pinned source prints.

That the finite layer of a celebrated result is this thin is, we think, the

rule	printed sentences	AI phrasings
prose: no claim shape	43	66
no printed quantity	6	1
ambiguous: two readings	3	3
asymptotic ($o(1)$, as $T \rightarrow \infty$)	2	0
an approximation sign	1	4
an unbounded quantifier	0	3

Table 5: The refusals by rule, over the 55 refused printed sentences and the 77 refused AI phrasings. The 10 sentences read are listed in Table A.7.

most useful number in this paper. It is not a criticism of the result; a proof is not a list of inequalities. It is a measurement of what independent machine checking can reach in a story of this kind without reading proofs, and the answer, here, is one certificate and a handful of printed constants. The battery that holds the instrument to account, planted claims with their expected words, the author’s own claims as a mathematician with expectations recorded before each run, and 782 planted defects at scale on the certificate with 511 controls undisturbed, is in Appendix A.

7. Conclusion

The formalisation’s numerical certificate holds over its printed integers by a route that shares nothing with Lean; its hypothesis EnclOK rests on a file that only its authors can read, and we have asked for it. The constants that looked as if they needed an interval kernel needed only an alternating series and the certificate’s own definitions, and once they were rational the window and the tail could be optimised exactly: the formalisation’s ξ' constants rise by $1.330 \cdot 10^{-6}$, with an optimality certificate and a calibration that returns the Montgomery–Taylor constant to 18 digits. One print in the abstract is a truncation. Of everything the story’s sources print, 10 sentences in 65 are finite claims. None of this bears on the theorems. What it offers their authors is small and concrete: publish the law file, and the one conditional step in an otherwise kernel-checked certificate becomes a step anyone can check; and take the window v_6 and fifteen terms, which raise two formalised constants by the same proofs with new rationals.

Declaration of competing interest

The author declares no competing financial or personal interests. The authors of the paper, the formalisation and the post have not been consulted on this draft; the request for the law file was posted on the formalisation’s repository on 2026-10-22, as a comment on its open issue 8.

CRediT authorship contribution statement

Carlos Toledo: conceptualization, methodology, software, validation, investigation, data curation, writing—original draft, writing—review and editing. The author list and the roles are to be agreed with the prospective co-authors.

Data availability

Every record, script and table source is public in the cert-machine repository (<https://github.com/carlostoledo1891/cert-machine>, commit 4df5fe2f7b06; archived as Toledo, 2026a): the copies of the translator’s records and pinned corpus `corpus/claimbox-riemann/` with their provenance `PROVENANCE-claimbox.json` (claimbox commit afd776c380318bf7b6a7949ec18d130dc64e62a1, 166 files, each with its sha256 on both ends); the manifest `LIFT-claimbox.json`; the battery `tools/check-claimbox-lift.js`; and the generator of this paper’s numbers `tools/paper-numbers/riemann-finite-layer.js`. The translator, `claimbox` (Toledo, 2026b), holds the grammar, the readers, the runs and the recall at the commit cited. The sources are cited at their origin and pinned by digest in `corpus/claimbox-riemann/corpus/riemann/PINS.json`; the `zeta-23-lean` repository’s address redirected to `anthropics/formal-math` when checked on 2026-10-22, and is cited as pinned. The 11/12 manuscript is cited by digest and not redistributed.

Declaration of generative AI and AI-assisted technologies

During the preparation of this work the author used large language models (Claude, Anthropic) to draft code, exposition and this manuscript in a certified-verification setting: constructions and prose were proposed with the help of the models, and every quantitative claim is decided by outward-rounded interval arithmetic or exact rational arithmetic, with adversarial controls that must fire for a record to be accepted. After using the tool the author reviewed and edited the content and takes full responsibility for the content of the publication. The records, not the models, are the evidence offered; they are public and re-runnable. For this paper specifically: the translator, its grammar and readers, the certificate files, the planted claims, the mathematician’s claims with their derivations, the recall at scale and this prose were drafted by AI agents (Claude) under written briefs and consolidated into the public record; every verdict here is a computation by the checker over pinned bytes, and no reading by an agent enters any verdict. The two corrections of Section 3 are the machine correcting the agent.

Funding

This work received no external funding.

Appendix A. The battery

The sources. Twenty-five files were fetched byte for byte on 2026-10-22 and pinned by sha256 (Table A.6): `openai/math` at commit `1581d4cbdd69`, `anthropics/zeta-23-lean` at `e1a4e6508154`, the paper’s e-print arXiv 2608.13637v2 (the L^AT_EX source, sha256 `ae369aaffb66`), the post as served (`840ff1f13cda`), the video’s captions as pasted (9,538 bytes, `b86fd8794cfb`). The sources were read for the claims they print and never run. Every decision took at most 0.014 seconds on an 8-core laptop.

source	files	bytes
openai/math (the index, the scope note, the four Comparator statements, three READMEs, the history)	11	660,334
anthropics/zeta-23-lean (README, formalization.yaml, the audit, the Challenge files, the five PairCeiling files)	16	233,062
arXiv 2608.13637v2 (the e-print’s members)	2	87,018
Anthropic’s research post, as served	1	202,748
the video’s captions, as pasted	1	9,538

Table A.6: The 31 pinned sources of the story, by origin.

The sentences read. Table A.7 lists the 10 of 65 sentences the grammar reads, and Table A.8 the count by author.

The planted claims. Fifteen claims about the certificate were written by Claude in the translator’s session, as a claimant summarising the certificate might write them, some right and some wrong on purpose, each with the word expected of it recorded before the run (Table A.9). They are a planted battery, not a sample of what any published AI text says. 6 refuted, 4 certified, 5 refused with the rule, every one as expected. The one that measures a limit is the wrong last digit of the transcendental, $c_1^* = 0.7532961\dots$: through the series certificate it is refuted; through the closed form it would be admitted and undecided.

The mathematician’s claims. The author of the translator’s session then wrote claims of its own, derived from the sources and the bytes, each with the word expected of it before the run: 99 in all, 81 holding and 18 refuted, every refutation a planted control. The first 22 (Table A.10): the two eighth-digit bounds on c_1^* ; the two “published as” prints; the CUE containment; the three pieces of the paper’s ceiling arithmetic; the five ξ' identities, the four eight-digit and the four five-digit constants; and a control one unit above

sentence, as printed	author	file	shape
$2 - c_{\text{MT}}^{-1} = 0.67250 \dots$	Claude: the paper, arXiv 2608.13637v2	arXiv/paper-v5-draft18.tex	digits
$\frac{1}{2}(3 - c_{\text{MT}}^{-1}) = 0.83625 \dots$	Claude: the paper, arXiv 2608.13637v2	arXiv/paper-v5-draft18.tex	digits
$p_0 \leq 0.6818287$	Claude: the paper, arXiv 2608.13637v2	arXiv/paper-v5-draft18.tex	bound
$ 256 S(j) - j \leq 3 \cdot 10^{-40}$	Claude: the paper, arXiv 2608.13637v2	arXiv/paper-v5-draft18.tex	bound
$ D(1) \leq 0.82395317$	Claude: the paper, arXiv 2608.13637v2	arXiv/paper-v5-draft18.tex	bound
$ 256 \cdot S(j) - j \leq 3 \cdot 10^{-40} \ (0 < j < 256)$	Claude: the Lean formalization zeta23	zeta23/README.md	bound
$2 - 1/c_1^* = 0.67250 \dots$	Claude: the Lean formalization zeta23	zeta23/formalization.yaml	digits
$(3 - 1/c_1^*)/2 = 0.83625 \dots$	Claude: the Lean formalization zeta23	zeta23/formalization.yaml	digits
$c_1^* = 0.7532960 \dots$	Claude: the Lean formalization zeta23	zeta23/ChallengeDeps.lean	digits
$ D(1) \leq 82395317/10^8$	Claude: the Lean formalization zeta23	zeta23/PairCeiling/LawN256.lean	bound

Table A.7: The 10 sentences the grammar reads.

author, source	sentences	read	refused
OpenAI: the manuscript’s abstract, as the index prints it	6	0	6
OpenAI: the release (index, scope note, Lean statements, history)	7	0	7
Claude: the paper, arXiv 2608.13637v2	18	5	13
Claude: the Lean formalization zeta23	21	5	16
Anthropic: the research post	4	0	4
the video’s captions	9	0	9

Table A.8: The 65 printed sentences by author.

P_{flat} . The other 77 (Table A.11) are the window search of Section 4: the windows, the optimality certificates at nine and fifteen terms, the tail, the parameter and the ζ calibration, with 16 controls refuted.

Table A.11: The 77 claims of the window search: the words as decided (REFUTED is a planted control in every case).

claim	word
$P_{\text{quartic}} \leq 0.86864018$	CERTIFIED
$P_{\text{q4}} \geq 0.86864105$	CERTIFIED
$D_{\text{q4}} \geq 0.93432052$	CERTIFIED
$P_{\text{w6}} \geq 0.86864115$	CERTIFIED
$D_{\text{w6}} \geq 0.93432057$	CERTIFIED
$P_{\text{w6}} \geq 0.86864116$	REFUTED

claim	word
$max_{w6} \leq 0$	CERTIFIED
$w6(1/2) = 10737/16000$	CERTIFIED
$max_{q4} \leq 0$	CERTIFIED
$q4(1/2) = 26893/40000$	CERTIFIED
$\kappa_9^{w6} = 1811057962010301991562419298932/$ $1600781683450427687806446724875$	CERTIFIED
$P_{o20} \geq 0.86864115895$	CERTIFIED
$P_{o20} \leq 0.86864115896$	CERTIFIED
$D_{o20} \geq 0.93432057947$	CERTIFIED
$max_{o20} \leq 0$	CERTIFIED
$res_{opt20} \leq 0$	CERTIFIED
$dmin_{opt20} \geq 0$	CERTIFIED
$dmin_{opt20} > 0$	REFUTED
$gap_{opt20} \leq 0$	CERTIFIED
$gap_{opt20} < 0$	REFUTED
$Pstar_{opt20} \geq 0.86864115895$	CERTIFIED
$Pstar_{opt20} \leq 0.86864115896$	CERTIFIED
$res_{opt6} \leq 0$	CERTIFIED
$dmin_{opt6} \geq 0$	CERTIFIED
$dmin_{opt6} > 0$	REFUTED
$gap_{opt6} \leq 0$	CERTIFIED
$gap_{opt6} < 0$	REFUTED
$Pstar_{opt6} \geq 0.86864115894$	CERTIFIED
$Pstar_{opt6} \leq 0.86864115895$	CERTIFIED
$\varepsilon_{15} = 8192/180952685768570625$	CERTIFIED
$P_{w6k15} \geq 0.86864150$	CERTIFIED
$D_{w6k15} \geq 0.93432075$	CERTIFIED
$P_{w6k15} \geq 0.86864151$	REFUTED
$P_{w6k15} \geq 0.8686415005$	CERTIFIED
$P_{w6k15} \geq 0.8686415006$	REFUTED
$max_{w6k15} \leq 0$	CERTIFIED
$w6k15(1/2) = 10737/16000$	CERTIFIED
$P_{o20k15} \geq 0.86864150052$	CERTIFIED
$P_{o20k15} \leq 0.86864150053$	CERTIFIED
$D_{o20k15} \geq 0.93432075026$	CERTIFIED
$max_{o20k15} \leq 0$	CERTIFIED
$res_{opt20k15} \leq 0$	CERTIFIED
$dmin_{opt20k15} \geq 0$	CERTIFIED
$dmin_{opt20k15} > 0$	REFUTED
$gap_{opt20k15} \leq 0$	CERTIFIED
$gap_{opt20k15} < 0$	REFUTED
$Pstar_{opt20k15} \geq 0.86864150052$	CERTIFIED

claim	word
$Pstar_{\text{opt20k15}} \leq 0.86864150053$	CERTIFIED
$res_{\text{opt6k15}} \leq 0$	CERTIFIED
$dmin_{\text{opt6k15}} \geq 0$	CERTIFIED
$dmin_{\text{opt6k15}} > 0$	REFUTED
$gap_{\text{opt6k15}} \leq 0$	CERTIFIED
$gap_{\text{opt6k15}} < 0$	REFUTED
$Pstar_{\text{opt6k15}} \geq 0.86864150051$	CERTIFIED
$Pstar_{\text{opt6k15}} \leq 0.86864150052$	CERTIFIED
$Dstar_{\text{opt20k15}} \geq 0.93432075026$	CERTIFIED
$P_{\text{w6l95}} \leq 0.83686$	CERTIFIED
$P_{\text{w6l95}} \geq 0.84$	REFUTED
$P_{\text{w6l90}} \leq 0.79415$	CERTIFIED
$res_{\text{zeta20}} \leq 0$	CERTIFIED
$dmin_{\text{zeta20}} \geq 0$	CERTIFIED
$dmin_{\text{zeta20}} > 0$	REFUTED
$gap_{\text{zeta20}} \leq 0$	CERTIFIED
$gap_{\text{zeta20}} < 0$	REFUTED
$Pstar_{\text{zeta20}} \geq 0.672500703679$	CERTIFIED
$Pstar_{\text{zeta20}} \leq 0.672500703680$	CERTIFIED
$res_{\text{zeta6}} \leq 0$	CERTIFIED
$dmin_{\text{zeta6}} \geq 0$	CERTIFIED
$dmin_{\text{zeta6}} > 0$	REFUTED
$gap_{\text{zeta6}} \leq 0$	CERTIFIED
$gap_{\text{zeta6}} < 0$	REFUTED
$Pstar_{\text{zeta6}} \geq 0.672500703679$	CERTIFIED
$Pstar_{\text{zeta6}} \leq 0.672500703680$	CERTIFIED
$Dstar_{\text{zeta20}} \geq 0.836250351839$	CERTIFIED
$Pstar_{\text{zeta2}} = 269/400$	CERTIFIED
$res_{\text{zeta2}} \leq 0$	CERTIFIED
$dmin_{\text{zeta2}} \geq 0$	CERTIFIED

Recall at scale. A defect planted in the bytes, the sentence or the page must come out as a refusal with its rule or as a changed decision, and an edit that is not a defect must leave the decision as it was. Every enclosure was moved by one unit either way and widened by one and two units; rows were dropped and exchanged; K and N were edited; the free row was moved; p_0 was edited in the file; the bar was moved in the sentence and in a copy of the page with the sentence following; another claimant's file was given as the bytes; the last digit of c_1^* was changed in a copy of its file (Table A.12). The expectation of every row edit was computed from the edited integers by hand: an enclosure sits on one side of $j \cdot K/256$, which side varies by row, and τ is 1.63 units of $256/K$, so a move away from j crosses the tolerance and a

claim	what it is	word
$ D(1) \leq 0.8239$	exact $ D(1) $ at the enclosures' ends is 0.82395316..., above 0.8239	REFUTED
$ D(1) \leq 0.824$	a rounding up: true	CERTIFIED
$ D(1) < 0.82395316$	cut short by $7 \cdot 10^{-10}$	REFUTED
$ 256 \cdot S(j) - j \leq 2 \cdot 10^{-40}$ ($0 < j < 256$)	worst row is $1.84 \cdot 10^{-40}$	CERTIFIED
$ 256 \cdot S(j) - j \leq 1 \cdot 10^{-40}$ ($0 < j < 256$)	below the worst row	REFUTED
$ 256 \cdot S(j) - j \leq 3 \cdot 10^{-40}$ ($0 < j < 100$)	not the certificate's rows	REFUSED (count)
$ 256 \cdot S(j) - j \leq 3 \cdot 10^{-40}$ ($0 < j < 257$)	one row too many: the closed-band row is free	REFUSED (count)
$p_0 \leq 0.68182$	$p_0 = 0.68182868...$	REFUTED
$p_0 \leq 0.6819$	a rounding up: true	CERTIFIED
$p_0 \geq 0.6818287$	relation the wrong way	REFUTED
$c_1^* = 0.7532961...$	a wrong last digit of a transcendental: decided from the series certificate (mathematician.md §1), the enclosure 0.75329606785... lies below the box $[0.75329605, 0.75329615]$ — the gap of the closed-form route closed by the reduction	REFUTED
$2 - 1/c_1^* = 0.6725...$	true to four digits; decided from the series certificate	CERTIFIED
$D(1) \leq 0.82395317$	a different quantity: no functional is written for $D(1)$ without the bars	REFUSED (decider)
we checked that $ D(1) \leq 0.82395317$	a prefix: the grammar reads a claim at the head of its sentence	REFUSED (prose)
$1/(6 \cdot 256^2) + \tau/512 < 2.5431316 \cdot 10^{-6}$	no shape: an expression on the left of a relation is not a named quantity	REFUSED (prose)

Table A.9: The 15 planted claims and their words.

move towards it does not. The first draft of the recall assumed every one-unit shift flips the decision; the first sample showed half do not, which is what the controls are for. 782 plants, 782 caught; 511 controls, 511 undisturbed; 1 undecided as recorded, the digits plant on the paper's closed-form claim (digits:af-mt-simple:0.67251). On the translator's earlier corpora the same test caught 56 of 56.

The translator's own battery, and its earlier corpora. The translator's test suite at the pinned commit prints 288 checks passing with 48 of 48 red controls firing, observed once from this repository on 2026-10-22 and kept with the copies. The stronger evidence for the pipeline is on its earlier corpora, in the method paper's register: EinsteinArena's 26-circle packing, decided as 6 predicates holding and one refuted on 3 pairs that overlap by 10^{-18} , float64 tangencies the platform admits and exact arithmetic does not; an earlier C_{1c} certificate of the optimization-constants registry, $C \geq 0.41030211724$, refuted by $3 \cdot 10^{-14}$, a lower bound rounded up; and a 605-point kissing record read two ways from one byte string, 9,510 against 9,480 violating pairs, the lesson that a reading is part of a verdict (Toledo, 2026c).

claim	why the word was expected	word
$c_1^* \leq 0.75329607$	the upper end of the enclosure, 0.75329606785607067722..., is below the bar	CERTIFIED
$c_1^* \geq 0.75329606$	the lower end, 0.75329606785607067702..., is above the bar	CERTIFIED
$(3 - 1/c_1^*)/2$, published as 0.8362	the paper's abstract prints 0.8362; the exact value 0.836250351839705822... lies above the box [0.83615, 0.83625] its digits allow — a truncation (a rounding prints 0.8363); as a lower bound the print is still true	REFUTED
$2 - 1/c_1^*$, published as 0.6725	0.672500703... lies in [0.67245, 0.67255]	CERTIFIED
$ 2 \cdot j \cdot K/256 - l_{0j} - h_{ij} \leq 1$ ($0 < j < 256$)	every band enclosure contains the exact CUE value $j \cdot K/256$ (124 at the upper end, 131 at the lower end)	CERTIFIED
$ p_0 - 2/3 \leq 0.016$	$p_0 - 2/3 = 0.0151620207...$	CERTIFIED
$B(8.2) \leq 0.6819$	$p_0 + 2.55 \cdot 10^{-6} \cdot 8.2 = 0.68184954...$	CERTIFIED
$c_{256} \leq 2.55 \cdot 10^{-6}$	$1/(6 \cdot 256^2) + \tau/512 = 2.5431315104... \cdot 10^{-6}$	CERTIFIED
$J_{\text{flat}} = 12517210259/88388425125$	the flat window integral, re-derived exactly from a_k and $\int s^{\wedge} d(1 - s)$	CERTIFIED
$\kappa_9^{\text{flat}} = 100905635384/88388425125$	$1 + J_{\text{flat}}$	CERTIFIED
$\varepsilon_9 = 1024/2990212875$	$a_9 \cdot 231/211$ with $a_9 = 2 \cdot 4^{10} \cdot 9!/20!$	CERTIFIED
$J_{\text{quartic}} = 299652803309$	the quartic window integral over my own	CERTIFIED
$34234270211/285991090937$	autocorrelation coefficients; $b_0 = \int v^2$ and $\sum$	
677125000000	$b_1 = 0$ beside it	
$\kappa_9^{\text{quartic}} = 277244140547$	$(\int v^2 + J_{\text{quartic}})/(\int v)^2$	CERTIFIED
$469154168336/245053976636$		
191319722125		
$P_{\text{flat}} \geq 0.85838371$	the Lean constant; margin $3.6 \cdot 10^{-9}$	CERTIFIED
$D_{\text{flat}} \geq 0.92919185$	the Lean constant; margin $6.8 \cdot 10^{-9}$	CERTIFIED
$P_{\text{quartic}} \geq 0.86864017$	the Lean constant	CERTIFIED
$D_{\text{quartic}} \geq 0.93432008$	the Lean constant	CERTIFIED
$P_{\text{flat}} \geq 0.85838$	the paper's print	CERTIFIED
$D_{\text{flat}} \geq 0.92919$	the paper's print	CERTIFIED
$P_{\text{quartic}} \geq 0.86864$	the paper's print	CERTIFIED
$D_{\text{quartic}} \geq 0.93432$	the paper's print	CERTIFIED
$P_{\text{flat}} \geq 0.85838372$	one unit of the eighth digit above the exact value 0.8583837135...	REFUTED

Table A.10: The mathematician's first 22 claims, with the expectation recorded before each run.

plant	must flip or refuse	must stay	undecided
one enclosure moved by ± 1 unit of K	255	255	0
one enclosure widened by one unit on one side	255	255	0
widened by two units	255	0	0
two enclosures exchanged	1	0	0
a row removed	2	0	0
the scale K edited	1	0	0
the count N edited	1	0	0
the free closed-band row moved by 10^{37}	1	1	0
the sentence edited: the bar moved, the range shortened, the relation flipped	5	0	0
the page copied with the bar moved and the sentence following	3	0	0
another claimant's file as the bytes	1	0	0
the exact rational p_0 edited in the file	1	0	0
the last printed digit of c_1^* changed in a copy of its file	1	0	1

Table A.12: The 1,294 rows of the recall at scale on the certificate.

Appendix B. Reproduction

Each result re-derives with the command beside it in Table B.1: in the translator’s repository at commit `afd776c38031`, with `cert-machine` beside it read-only, for the runs; in this repository for the copies and the numbers. Every run re-pins the sources before it reads them and refuses a file that does not match.

result	command (in <code>claimbox</code> unless marked)	expected
the pins	<code>make riemann && make drift</code>	31 files held
the certificate	<code>python3 -m claimbox batch corpus/ riemann-jobs.json --out records/ riemann --check --record records/ riemann-run.json</code>	Table 1
the planted claims	<code>python3 -m claimbox batch corpus/ riemann-ai-jobs.json --out records/ riemann-ai --check --record records/ riemann-ai-run.json</code>	Table A.9
the mathematician’s claims	<code>python3 -m claimbox batch corpus/ riemann-math-jobs.json --out records/ riemann-math --check --record records/riemann-math-run.json</code>	Table A.10
recall at scale	<code>CLAIMBOX_FULL=1 python3 tests/riemann_ recall.py</code>	782/782, 511/511
the window search	<code>python3 tools/xi_window.py --check; python3 tools/xi_window.py --sweep-terms; --sweep-lambda; --sweep-zeta</code>	every rational of the quartic reproduced; Tables 2, 3, 4
the translator’s battery	<code>make test</code>	288 pass, 48/48 reds
the copies (<code>cert-machine</code>)	<code>node tools/check-claimbox-lift.js</code>	166 files held; the numbers identical
this paper’s numbers (<code>cert-machine</code>)	<code>node tools/paper-numbers/ riemann-finite-layer.js</code>	the macro file rewritten identically

Table B.1: Each result, the command that re-derives it and what it must print.

References

- Adem B10, 2026. Mathematicians checked the AI’s Riemann breakthrough. Video; its captions as pasted on 2026-10-22. URL: <https://youtu.be/dsBXYAwQieI>.
- Alpöge, L., Furman, R., 2026. More than two thirds of the zeros of the Riemann zeta function are simple and on the critical line. arXiv preprint [arXiv:2608.13637](https://arxiv.org/abs/2608.13637).

- Anthropic, 2026a. Claude has improved on a longstanding lower bound for the fraction of zeros of the Riemann zeta function that satisfy the Riemann hypothesis. Research post, as served on 2026-10-22. URL: <https://www.anthropic.com/research/riemann-zeta>.
- Anthropic, 2026b. zeta-23-lean. GitHub repository, the Lean 4 formalization of Alpöge–Furman; commit e1a4e650. URL: <https://github.com/anthropics/zeta-23-lean>.
- Carneiro, E., Chandee, V., Littmann, F., Milinovich, M.B., 2017. Hilbert spaces and the pair correlation of zeros of the Riemann zeta-function. *Journal für die reine und angewandte Mathematik* 725, 143–182. doi:10.1515/crelle-2014-0078.
- Montgomery, H.L., 1973. The pair correlation of zeros of the zeta function, in: *Analytic Number Theory*. American Mathematical Society. volume 24 of *Proceedings of Symposia in Pure Mathematics*, pp. 181–193. doi:10.1090/pspum/024/9944.
- OpenAI, 2026. math: manuscripts and Lean challenges. URL: <https://github.com/openai/math>. Commit fd4aeeb2.
- Toledo, C., 2026a. cert-machine: independent exact certification of machine-generated mathematics. doi:10.5281/zenodo.22225860. The repository archive: code, records and reports. The concept DOI resolves to the latest version; the version this paper was built from is named on its title page.
- Toledo, C., 2026b. claimbox: the claim-machine. GitHub repository; a standard-library translator from a claim’s text to the predicate objects cert-machine decides; commit b37d980. URL: <https://github.com/carlostoledo1891/claimbox>.
- Toledo, C., 2026c. Independent exact certification of machine-generated mathematics: the first register. URL: <https://carlostoledo.co/paper/register.pdf>. Archived with the repository, doi:10.5281/zenodo.22225860. Draft.