ROBIN'S 1984 CRITERION FOR THE RIEMANN HYPOTHESIS: A FORMALLY VERIFIED PROOF

JONAS WHIDDEN

Abstract. Guy Robin proved in 1984 that the Riemann hypothesis is equivalent to the strict inequality

$$\sigma(n) < e^{\gamma} n \log\log n \qquad (n > 5040).$$

We give a detailed mathematical reconstruction paired with a formally verified Lean 4 proof. The argument reduces arbitrary integers to colossally abundant extrema, proves the bounded range with exact rational certificates, establishes the large-height implication under the Riemann hypothesis through a weighted explicit formula and a strict coefficient gap, and proves the converse by turning Nicolas–Landau oscillations into least-common-multiple counterexamples. The source-specific weighted formula and oscillation argument are proved in repository-owned modules, while substantial foundational analysis is imported from pinned public libraries: axiom-free here does not mean dependency-free or redeveloped from first principles.

A declaration-level map connects the exposition to the checked source. Comparator verifies that the solution proves the exact Mathlib-only challenge statements and respects the configured axiom boundary; Lean's kernel and the independent NanoDa kernel then replay the export. Public source, certificates, CI logs, release artifacts, and reproduction instructions make this assurance chain inspectable. The reconstruction used a human-directed automated context system to track sources, theorem state, dependencies, numerical evidence, and verification status. The paper documents its role through retained outputs and two concrete failure-and-revision cases, without claiming a measured productivity advantage. Reusable outputs include a generic positive-transform theorem, multiplicity-aware zero infrastructure, and a theorem-parametric finite-certificate checker. Robin's result remains an equivalence criterion; it does not prove the Riemann hypothesis.

1. Introduction

For a positive integer n, let

$$\sigma(n) = \sum_{d \mid n} d, \qquad \mathcal{A}(n) = \frac{\sigma(n)}{n},$$

and let γ be the Euler–Mascheroni constant. Robin's theorem is the following exact criterion. Throughout, RH denotes the standard zero-locus formulation: every zero s of the analytically continued Riemann zeta function which is neither a negative even integer nor the pole at 1 satisfies $\Re s = 1/2$. This is also Mathlib's definition of `RiemannHypothesis`; it is stronger in presentation, but equivalent in the usual way, to saying that zeta has no zero with $\Re s > 1/2$.

Theorem 1.1 (Robin [12])**.** *The Riemann hypothesis holds if and only if*

$$\sigma(n) < e^{\gamma} n \log\log n \tag{1.1}$$

Date: September 2026.

2020 *Mathematics Subject Classification.* 11M26, 11N37, 03B35.

Key words and phrases. Riemann hypothesis, sum of divisors, Robin's inequality, colossally abundant numbers, Nicolas criterion, formal verification, Lean.

for every integer $n > 5040$.

The cutoff is part of the theorem. In particular, the result is not merely an eventual estimate. It combines an analytic argument above an explicit height with a finite proof below that height.

Robin's original article [12] is in French. Although a scanned PDF is now available online, the paper remains comparatively difficult to discover through ordinary bibliographic channels. One purpose of the present paper is therefore expository: to give an accessible English, proof-oriented account of the argument while preserving the attribution and source trail of the original. Nicolas's related Euler-function criterion has a later English treatment [10]; the present paper concerns Robin's distinct sum-of-divisors argument.

The complete Lean source, certificate data, reproducibility instructions, and release artifacts are available at

https://github.com/kimihiro64/Robin1984.

Section 11 links the public theorem boundary and the principal modules for each mathematical stage, while the finite-range section links the exact certificate rows directly.

The second principal result makes Robin's extremal reduction explicit. For $\varepsilon > 0$, put

$$Q_\varepsilon(m) = \frac{\sigma(m)}{m^{1+\varepsilon}} = \frac{\mathcal{A}(m)}{m^\varepsilon}. \tag{1.2}$$

We say that $N > 1$ is colossally abundant with parameter ε if $Q_\varepsilon(m) \le Q_\varepsilon(N)$ for every $m > 1$.

Theorem 1.2. *The Riemann hypothesis holds if and only if (1.1) holds for every $N > 5040$ which is colossally abundant with some positive parameter.*

Colossally abundant numbers were introduced by Alaoglu and Erdős [1]. The criterion itself and its use of those extrema are due to Robin. The converse uses the oscillation theorem of Nicolas [9] together with Landau's positive-transform principle [5]. The explicit estimates for prime functions used below trace to Rosser and Schoenfeld [13] and to Costa Pereira [2, 3]. Our purpose is not to claim new priority for those ingredients, but to give a complete proof whose exact logical and numerical dependencies have been checked.

The new contribution is consequently one of reconstruction, reusable formal infrastructure, and verification. The analytic argument is carried through a multiplicity-counted zero divisor; Landau's positive-transform principle is isolated as a generic theorem about moment-generating functions; the finite range is reduced to a reusable soundness theorem plus exact data; and the public biconditional is assembled without a project axiom. This separation is useful beyond the present theorem: source mathematics, generic analytic lemmas, untrusted certificate discovery, kernel-checked certificate validation, and final theorem assembly remain distinct layers.

The development was assisted by an automated research and formalization context system. It maintained source provenance, exact theorem targets, definition and import dependencies, live proof obligations, failed routes, numerical evidence, and verification status while human direction selected the mathematical claims and accepted the final statements. Section 11 describes this architecture and its trust boundary. The claims made about that system concern the retained source, certificates, documented revisions, and validation records. No controlled comparison of human effort, elapsed development time, or error rate is claimed.

The proof has four parts. Section 5 reduces arbitrary integers to colossally abundant numbers, the global maximizers of Q_ε. Section 6 closes the bounded range. Section 7 proves the inequality under RH from a strict explicit coefficient gap. Sections 8 and 9 give the Nicolas–Landau converse: an off-critical zero forces recurring prime-product oscillations, which least common multiples turn into integer counterexamples. At the highest level the two directions are

$$\text{RH} \implies \text{explicit prime-function bounds} \implies \mathcal{M}(n) < 0,$$

and

$$\neg\mathrm{RH} \implies \text{negative Nicolas oscillations} \implies \mathcal{M}(L_P) > 0 \text{ for arbitrarily large } P.$$

The finite theorem joins the first chain to the exact cutoff 5040, and the LCM identity joins the second chain to literal integer counterexamples. The dedicated main-text account of the Lean organization and checking is confined to Section 11; the appendices contain supporting audits.

2. Roadmap and informal guide

This section gives the argument without its explicit constants or formalization details. It is intended as a map for readers who do not work regularly with analytic number theory. The later sections prove each step in the form actually needed for the theorem.

The proof and its checked endpoints can be read as the following directed graph. The declaration names are printed here so that a reader can move from the mathematical route to the exact Lean statement without searching the repository.

mathematical stage	paper endpoint and checked declaration
CA extremal reduction	Proposition 5.5; `nativeRobinInequalityAll_iff_colossallyAbundant` in `CAReduction`
finite range	Theorem 6.5; `nativeRobinInequality_finite_log` in `FiniteComplete`
large-height RH branch	Theorem 7.4; `nativeRobinInequality_of_RH_large_log` in `LargeHeightRobin`
false-RH converse	Theorem 9.3; `riemannHypothesis_of_nativeRobinInequalityAll` in `LCM bridge`
public biconditionals	Theorems 1.1 and 1.2; `Robin1984.robin_inequality_iff_riemannHypothesis` and `Robin1984.riemannHypothesis_iff_colossallyAbundant_robin` in `Solution`

2.1. What Robin's inequality measures.

The ratio $\mathcal{A}(n) = \sigma(n)/n$ measures how rich n is in divisors relative to its size. If

$$n = \prod_p p^{a_p},$$

then $\mathcal{A}(n)$ is a product of local factors, one for each prime dividing n. Small primes increase this ratio most efficiently, and raising a prime to a higher power gives progressively smaller additional gains. Integers built from many small prime powers are therefore the difficult cases for an upper bound on $\mathcal{A}(n)$.

Robin's inequality says that even these difficult integers satisfy

$$\mathcal{A}(n) < e^{\gamma} \log\log n$$

once $n > 5040$, exactly when RH is true. Taking logarithms turns the comparison into the sign of

$$\mathcal{M}(n) = \log \mathcal{A}(n) - \gamma - \log\log\log n.$$

The proof is consequently a search for structural reasons why this margin is negative, or, if RH fails, why it must sometimes be positive.

2.2. **Why the extremal integers suffice.** For a positive penalty ε, the quantity

$$Q_\varepsilon(m) = \mathcal{A}(m)m^{-\varepsilon}$$

rewards divisors but charges for size. Its maximizers are the colossally abundant integers. They are useful because the penalty separates into independent prime-power decisions: include a prime-power event when its marginal gain exceeds its size cost, exclude it when the cost is larger.

Given a general integer n, choose the penalty

$$\varepsilon = \frac{1}{\log n \, \log\log n}.$$

Let q maximize Q_ε. Maximality says that the logarithmic abundancy of n lies below the affine line determined by q and this penalty. The function appearing in Robin's bound is concave, and the chosen penalty is its tangent slope at n. The tangent inequality therefore transfers Robin's inequality from q back to n. A short exact calculation at the transition $5040 \leftrightarrow 55440$ ensures that the chosen maximizer lies above the exceptional cutoff whenever this reduction is used. The bounded interval not covered by that argument is settled separately. This explains why checking colossally abundant integers is both sufficient and necessary.

2.3. **Why the Riemann hypothesis implies the inequality.** Write $H = \log n$. The prime-power factorization of n expresses $\log \mathcal{A}(n)$ as a sum of marginal gains. Completing this sum to a convenient set of small primes produces two competing quantities. One is an analytic error term, controlled through weighted estimates for the Chebyshev functions ϑ and ψ. The other is a definite arithmetic deduction: an integer of height H cannot take every small-prime gain for free, so omitted primes or finite exponents impose a cost.

Under RH, the weighted explicit formula bounds the analytic error on the natural square-root scale

$$S(H) = H^{-1/2}(\log H)^{-1}.$$

For $H \geq 74500$, the proof shows that the total error is strictly less than $(113/50)S(H)$, while a complete block of prime squares guarantees a cost of at least $(113/50)S(H)$. Hence error minus cost is negative, which is exactly $\mathcal{M}(n) < 0$.

The square-root scale has a simple spectral interpretation. The explicit formula rewrites the weighted prime-counting discrepancy as a superposition of contributions resembling H^ρ, one for each nontrivial zero ρ of zeta. Under RH, every contribution has magnitude on the $H^{1/2}$ scale; the weight damps the oscillation and the inverse-square zero sum controls the aggregate amplitude. The displayed error term is therefore a rigorous envelope for all of those waves, not a numerical fit.

The prime cost is the complementary combinatorial picture. An integer with logarithmic size H has a finite budget of prime exponents. It cannot collect every favorable small-prime contribution without spending that budget. The minimum in (7.9) records whichever of two unavoidable losses is cheaper for a prime. The strict coefficient gap says that even this conservative total loss is larger than the worst analytic envelope.

Below that analytic height, the argument is finite but not a decimal computation. Exact factorizations handle the first interval. The remaining range is covered by rational log-height rows. Each row lists a finite box of possible profitable prime powers and proves, by exact integer products and rational logarithm bounds, that the Robin margin is negative throughout the whole interval. Thus the finite and analytic halves meet at $H = 74500$ without an unverified numerical gap.

2.4. **Why the inequality implies the Riemann hypothesis.** The converse is best read contrapositively. If RH is false, symmetry of the zeta function supplies a nontrivial zero to the right of the critical line. That zero becomes a pole of a Mellin transform built from a weighted prime-counting error. Landau's positivity principle says that this transform cannot remain analytic as far as that pole if the underlying error is eventually of one sign. Consequently the error has negative excursions of a fixed power size arbitrarily far out.

Geometrically, the real axis of convergence acts like a wall. Eventual one-sidedness turns the Mellin transform, after $x = e^y$, into the Laplace transform of a positive measure. Positivity forces convergence to advance whenever the transform analytically crosses that wall. An off-critical zero, however, creates a genuine pole just beyond it. The two statements are incompatible, so the assumed one-sided tail must fail repeatedly. Section 8 replaces this picture by the exact continuation, integrability, and pole calculations.

Nicolas's finite product $F(P)$ converts those analytic excursions into negative excursions of $\log F(P)$. To obtain actual integers, take

$$L_P = \operatorname{lcm}(1, 2, \ldots, P).$$

This least common multiple contains precisely the prime powers up to P and satisfies $\log L_P = \psi(P)$. An exact identity expresses its Robin margin as

$$\mathcal{M}(L_P) = -\log F(P) - D(P) - T(P),$$

where the height loss $D(P)$ and the finite-prime-power reserve $T(P)$ are nonnegative but smaller than the oscillation. The negative excursions of $\log F(P)$ therefore give positive values of $\mathcal{M}(L_P)$ for arbitrarily large P. These are violations of Robin's inequality above 5040. Hence universal validity of the inequality forces RH.

The technical proof that follows adds the exact thresholds, constants, convergence arguments, and finite data needed to make this outline rigorous; it does not change the four-step logical route.

3. Historical comparison and scope of novelty

The theorem, the use of colossally abundant extrema, the weighted prime-function estimates, and the Nicolas–Landau oscillation mechanism are classical. This paper claims no new unconditional theorem about the zeta function and no priority for those ingredients. Its contributions occur at three different levels: an English reconstruction of Robin's proof; several exact decompositions and interfaces chosen to make that proof formalizable; and a checked Lean artifact with independent statement and kernel replay. The distinction is summarized below.

stage	classical source	present reconstruction
criterion and cutoff	Robin's Theorem 1 [12]	The same theorem and cutoff, stated directly over Mathlib objects; no new mathematical claim.
CA reduction	Robin's Section 3, Proposition 1, compares successive CA integers using the convexity of $x \mapsto \varepsilon x - \log\log x$.	Proposition 5.2 expresses the same supporting-line geometry at an arbitrary target and separates the exact $5040 \leftrightarrow 55440$ startup transition. The reformulation and its quantifier-exact assembly are project work; the reduction is Robin's.
RH implication	Robin's Section 3, especially Lemmas 2–8 and Proposition 3, combines a weighted explicit formula with prime-product estimates, following Nicolas but requiring finer bounds.	The complete exponent budget packages every actual exponent into an error term and a nonnegative prime cost, then proves a single exact coefficient gap at $H \geq 74500$. This is a project-authored formal decomposition and explicit cutoff, not a claim of a new asymptotic estimate or an optimal constant.
finite range	Robin reports computing 2347 CA integers below the analytic range and then checking the remaining integers below 55440.	Exact startup factorizations, a tangent interval, and rational certificate rows replace the historical computer calculation by kernel-checkable data plus a general soundness theorem.
false-RH converse	Robin invokes Nicolas's oscillation theorem, itself based on Landau's principle, and transfers the conclusion directly to CA integers.	Sections 8 and 9 reconstruct the transform argument, isolate the positive-transform theorem, and use least common multiples to obtain literal integer counterexamples before applying the CA equivalence. This is an alternative formal proof architecture, not a new oscillation theorem attributed to this project.

Two formal interfaces merit a more precise explanation. First, a classical sum over zeta zeros is normally understood to repeat a zero according to its analytic multiplicity. That convention is not enough for the interchanges needed in Lean. The multiplicity-aware infrastructure supplies an explicit index type for the divisor of ξ, a value map from indices to nonzero zeros, inverse-norm summability, and the hypotheses licensing each exchange of an integral with an infinite sum. It proves no new fact about where the zeros lie; it turns the classical convention into a reusable formal interface.

Second, the complete prime cost in (7.9) is a bookkeeping device. Robin estimates several prime products separately. The present proof first sums the actual exponent choices of an arbitrary integer, then completes to all relevant primes. Each missing prime or prematurely terminated exponent pays the smaller of two explicit penalties, so the analytic error and arithmetic loss can be bounded in independent modules. This modularity is useful for verification, but no theorem here says that it is computationally minimal or analytically more efficient than every earlier argument.

Third, Lemma 8.1 is called a "finite Nicolas comparison" because it connects Nicolas's finite product to the analytic error used in his oscillation theorem. The product, Mertens series, and partial-summation idea are classical. The repository-owned contribution is the exact identity $\log F = R_M + R_\vartheta + K$, the explicit $4/x$ tail, and the separately proved signs $R_\vartheta \leq 0$ and

$K \le J$ that make the comparison usable without asymptotic shorthand. This is again a formal reconstruction and reusable interface, not a new oscillation phenomenon.

3.1. **Relation to prior formal work.** The adjacent formal literature concerns the analytic objects around the criterion rather than Robin's divisor-sum biconditional itself. Gomes and Kontorovich [4] formalized an equivalent zero-locus formulation of the Riemann hypothesis through the Dirichlet eta function and supplied a Mathlib implementation of its analytic prerequisites; their development does not state Robin's inequality. Loeffler and Stoll [7] describe the Mathlib development of zeta and L-functions, including the formal statement of the Riemann hypothesis used here. The PrimeNumberTheoremAnd project [11] supplies Hadamard-factorization, completed-zeta, Mertens, and explicit prime-function infrastructure on which the present reconstruction builds.

Those developments establish definitions and reusable analytic foundations. The present artifact connects them to a different endpoint: Robin's exact sum-of-divisors criterion, including the weighted endpoint, complete exponent budget, finite certificate cover, Nicolas–Landau continuation, and LCM transfer. The author is not aware of an earlier mechanically checked proof of this complete biconditional, but no exhaustive priority claim is made. The comparison asserted here is narrower and inspectable: none of the cited formal developments contains the two public declarations proved by this repository.

The primary intended audience is therefore formal mathematics and analytic number theory. For AI-assisted mathematics, the contribution is a concrete case study in maintaining fixed theorem targets, provenance, exact certificates, and independent acceptance gates; it is not a new general machine-learning model. Section 11 makes that case study and its evidence explicit.

4. Arithmetic preliminaries

For $n > e^e$, set

$$\mathcal{M}(n) = \log \mathcal{A}(n) - \gamma - \log\log\log n. \tag{4.1}$$

All logarithms in (4.1) are then defined, and Robin's inequality is equivalent to $\mathcal{M}(n) < 0$.

Write $n = \prod_p p^{a_p}$. For a prime p and an integer $j \ge 1$, define the marginal gain and threshold

$$g_{p,j} = \log\left(\frac{1-p^{-(j+1)}}{1-p^{-j}}\right), \qquad \tau_{p,j} = \frac{g_{p,j}}{\log p}. \tag{4.2}$$

The event (p, j) records the change from exponent $j-1$ to exponent j.

Lemma 4.1 (Prime-power event decomposition)**.** *For every positive integer n and every $\varepsilon > 0$,*

$$\log \mathcal{A}(n) = \sum_p \sum_{1 \le j \le a_p} g_{p,j}, \tag{4.3}$$

$$\log Q_\varepsilon(n) = \sum_p \sum_{1 \le j \le a_p} (g_{p,j} - \varepsilon \log p). \tag{4.4}$$

For fixed p, the sequences $g_{p,j}$ and $\tau_{p,j}$ are nonincreasing in j.

Proof. The geometric-series identity gives

$$\frac{\sigma(p^a)}{p^a} = \frac{1-p^{-(a+1)}}{1-p^{-1}} = \prod_{j=1}^{a} \frac{1-p^{-(j+1)}}{1-p^{-j}}.$$

Multiplicativity and logarithms give (4.3); subtracting $\varepsilon \log n$ gives (4.4). If $q = p^{-1} \in (0,1)$, direct cross-multiplication shows that $(1-q^{j+2})/(1-q^{j+1}) \le (1-q^{j+1})/(1-q^j)$. Taking logarithms and then dividing by $\log p > 0$ proves the monotonicity. □

The reduced weight of an event at parameter ε is

$$w_\varepsilon(p,j) = g_{p,j} - \varepsilon \log p = (\tau_{p,j} - \varepsilon)\log p. \tag{4.5}$$

Thus events above the threshold have nonnegative reduced weight and events below it have nonpositive reduced weight. This is the elementary complementary-slackness description of a colossally abundant factorization.

Lemma 4.2 (Existence of a maximizer)**.** *For every $\varepsilon > 0$ there exists an integer $N > 1$ maximizing Q_ε.*

Proof. The divisor bound

$$\mathcal{A}(n) = \sum_{d|n} \frac{1}{d} \le \sum_{1\le d\le n} \frac{1}{d} \le 1 + \log n$$

implies $Q_\varepsilon(n) \le (1+\log n)n^{-\varepsilon} \to 0$. Beyond a finite point the objective is below its value at 2. A maximum over the remaining finite set is therefore a global maximum. □

5. The colossally abundant tangent reduction

Put $H_n = \log n$ and

$$\lambda(n) = \frac{1}{H_n \log H_n}. \tag{5.1}$$

This is the derivative at H_n of the concave function $\phi(H) = \log\log H$.

Lemma 5.1 (Global tangent inequality)**.** *If $L, U > 1$, then*

$$\log\log U - \log\log L \le \frac{U-L}{L\log L}. \tag{5.2}$$

Proof. Apply $\log y \le y - 1$ first to $y = (\log U)/(\log L)$ and then to $y = U/L$. Division by the positive quantities L and $\log L$ gives (5.2). This proof applies whether U lies to the left or to the right of L. □

Proposition 5.2 (Tangent transfer)**.** *Let $n, q > 5040$. Suppose that*

$$\log Q_{\lambda(n)}(n) \le \log Q_{\lambda(n)}(q). \tag{5.3}$$

If Robin's inequality holds at q, then it holds at n.

Proof. Writing $L = H_n$ and $U = H_q$, inequality (5.3) becomes

$$\log \mathcal{A}(n) - \lambda(n)L \le \log \mathcal{A}(q) - \lambda(n)U.$$

By Lemma 5.1, $\phi(U) - \phi(L) \le \lambda(n)(U-L)$. Adding the two inequalities yields

$$\log \mathcal{A}(n) - \phi(L) \le \log \mathcal{A}(q) - \phi(U).$$

After subtracting γ, the left- and right-hand sides are precisely $\mathcal{M}(n)$ and $\mathcal{M}(q)$. Hence $\mathcal{M}(q) < 0$ implies $\mathcal{M}(n) < 0$. □

The remaining point is to ensure that a maximizer at the tangent parameter is itself above the exceptional cutoff. The exact startup transition is particularly useful. Let

$$\varepsilon_0 = \frac{\log(12/11)}{\log 11}. \tag{5.4}$$

The event $(11, 1)$ changes 5040 into $55440 = 11 \cdot 5040$ and has threshold ε_0.

Lemma 5.3 (The startup tie)**.** *Both* 5040 *and* 55440 *maximize* Q_{ε_0}*. Moreover*

$$\frac{1}{34} < \varepsilon_0 < \frac{1}{22}. \tag{5.5}$$

Proof. The factorization $5040 = 2^4 3^2 5 \cdot 7$ determines its included and first excluded events. Put

$$r_{p,j} = \frac{1 - p^{-(j+1)}}{1 - p^{-j}}, \qquad \tau_{p,j} = \frac{\log r_{p,j}}{\log p}.$$

The complete boundary calculation is small enough to display:

event	$r_{p,j}$	comparison with ε_0
$(2,4)$	$31/30$	$\tau_{2,4} > \varepsilon_0$
$(3,2)$	$13/12$	$\tau_{3,2} > \varepsilon_0$
$(5,1)$	$6/5$	$\tau_{5,1} > \varepsilon_0$
$(7,1)$	$8/7$	$\tau_{7,1} > \varepsilon_0$
$(2,5)$	$63/62$	$\tau_{2,5} < \varepsilon_0$
$(3,3)$	$40/39$	$\tau_{3,3} < \varepsilon_0$
$(5,2)$	$31/30$	$\tau_{5,2} < \varepsilon_0$
$(7,2)$	$57/56$	$\tau_{7,2} < \varepsilon_0$
$(11,1)$	$12/11$	$\tau_{11,1} = \varepsilon_0$

Here is the complete arithmetic behind the strict signs. For an event with base p and gain ratio r, positivity of logarithms gives

$$p < r^N \implies \frac{1}{N} < \frac{\log r}{\log p}, \qquad r^N < p \implies \frac{\log r}{\log p} < \frac{1}{N}.$$

The cutoff enclosure (5.5) follows from the two exact rational inequalities

$$11 < (12/11)^{34}, \qquad (12/11)^{22} < 11.$$

For the four included top events the kernel checks

$$2 < (31/30)^{22}, \quad 3 < (13/12)^{22}, \quad 5 < (6/5)^{22}, \quad 7 < (8/7)^{22};$$

hence their thresholds exceed $1/22 > \varepsilon_0$. For the four first-missing events it checks

$$(63/62)^{34} < 2, \quad (40/39)^{34} < 3, \quad (31/30)^{34} < 5, \quad (57/56)^{34} < 7;$$

hence their thresholds lie below $1/34 < \varepsilon_0$. These are integer comparisons after denominators are cleared; the tie $(11,1)$ is definitional.

For a fixed prime, monotonicity in j propagates the displayed top-event and first-missing-event signs through every layer. Every prime not dividing 5040 is at least 11, and the first-layer threshold is nonincreasing in the prime; hence its reduced weight is nonpositive, with equality only at $(11,1)$. Equation (4.5) now compares the event packet of 5040 term by term with every competitor. Thus 5040 maximizes Q_{ε_0}. The tied event has weight zero, and adjoining it multiplies 5040 by 11, proving the same maximality for 55440. □

Lemma 5.4. *If* $n \geq 720720$ *and* q *maximizes* $Q_{\lambda(n)}$*, then* $q > 5040$*.*

Proof. The function $\lambda(n)$ is nonincreasing above 5040. Exact logarithmic bounds give

$$\lambda(720720) < \frac{1}{34} < \varepsilon_0,$$

and hence $\lambda(n) < \varepsilon_0$. Write $\lambda = \lambda(n)$. If $q \leq 5040$, then maximality at ε_0 gives

$$\log Q_\lambda(q) - \log Q_\lambda(5040) = \log Q_{\varepsilon_0}(q) - \log Q_{\varepsilon_0}(5040) + (\varepsilon_0 - \lambda) \log \frac{q}{5040} \leq 0.$$

On the other hand, the exact tie of Lemma 5.3 gives

$$\log Q_\lambda(55440) - \log Q_\lambda(5040) = (\varepsilon_0 - \lambda)\log 11 > 0.$$

Thus no $q \le 5040$ can maximize Q_λ, because 55440 has strictly larger objective value. □

Proposition 5.5 (Reduction to colossally abundant integers)**.** *Assume Robin's inequality for* $5040 < n < 720720$*. Then the following are equivalent:*

(i) Robin's inequality holds for every $n > 5040$*;*
(ii) it holds for every colossally abundant $N > 5040$*.*

Proof. Only (ii)⇒(i) needs proof. Let $n \ge 720720$. By Lemma 4.2, choose a maximizer q for $Q_{\lambda(n)}$. Lemma 5.4 gives $q > 5040$, so (ii) gives Robin's inequality at q. Maximality gives (5.3), and Proposition 5.2 transfers the inequality to n. The assumed finite interval handles the remaining integers. □

6. The finite range

The analytic argument in Section 7 starts at $\log n = 74500$. The exact cutoff therefore requires a proof for all $5040 < n$ below that logarithmic height. We divide it into three pieces.

Proposition 6.1. *For* $5041 \le n < 7560$*,*

$$35\sigma(n) \le 127n. \tag{6.1}$$

Consequently Robin's inequality holds throughout this interval.

Proof. For each integer in the interval, an exact prime-power factorization is given. Pairwise coprimality and the identity $\sigma(p^a) = 1 + p + \cdots + p^a$ reduce (6.1) to natural-number arithmetic. To pass from (6.1) to Robin's inequality, use the exact bounds

$$e^\gamma > \frac{177}{100}, \qquad \log\log n > \frac{103}{50}, \qquad \frac{127}{35} < \frac{177}{100}\frac{103}{50}.$$

□

Proposition 6.2. *Robin's inequality holds for* $7560 \le n < 720720$*.*

Proof. The maximality of 5040 at ε_0 gives

$$\log \mathcal{A}(n) - \varepsilon_0 \log n \le \log \mathcal{A}(5040) - \varepsilon_0 \log 5040.$$

Set $y = \log n$. It is enough to prove positivity of the explicit gap

$$G(y) = \gamma + \log\log y - \log \mathcal{A}(5040) - \varepsilon_0(y - \log 5040),$$

because the preceding inequality then places $\log \mathcal{A}(n)$ strictly below $\gamma + \log\log y$. The curvature used at this step is explicit:

$$G''(y) = -\frac{\log y + 1}{y^2(\log y)^2} < 0 \qquad (y > 1).$$

Thus G is concave, and its value at every point of a closed interval is at least the smaller of its two endpoint values. The endpoint checks use the following strict rational bounds:

endpoint	certified inequalities
7560	$\gamma > 577/1000$, $\log\log\log 7560 > 3917/5000$, $\log \mathcal{A}(5040) < 269/200$, and $\varepsilon_0 \log(3/2) < 203/13500$
720720	$\gamma > 72/125$, $\log\log\log 720720 > 191/200$, $\log \mathcal{A}(5040) < 269/200$, and $\varepsilon_0 \log 143 < 5/27$

Here $\mathcal{A}(5040) = 403/105$, and the arguments of the last two logarithms arise from $7560/5040 = 3/2$ and $720720/5040 = 143$. Substitution makes each endpoint gap a strict rational inequality, so concavity proves positivity throughout the interval. □

For the much larger interval above 720720, a compact rational certificate replaces integer-by-integer factorization. A certificate row consists of a log-height interval $[L, U]$, a rational slope λ, prime-layer cutoffs, an upper bound for the corresponding product, a lower height bound, and lower bounds for $\log L$ and $\log U$. The layer inequalities certify which prime-power events can have positive reduced weight. The product inequality bounds the maximum possible value of $\log \mathcal{A}(n) - \lambda \log n$. Finally, concavity reduces the Robin comparison on $[L, U]$ to two endpoint inequalities.

More precisely, for a cutoff vector $c = (c_1, \dots, c_d)$ define the layer box

$$\mathcal{B}(c) = \{(p, j) : p \text{ prime},\ 1 \le j \le d,\ p \le c_j\}.$$

The reduced weight of an event is

$$w_{p,j}(\lambda) = g_{p,j} - \lambda \log p = \log p\,(\tau_{p,j} - \lambda). \tag{6.2}$$

To state the integer boundary checks without pretending that a cutoff is prime, extend this expression to every real base $x > 1$ by

$$\widetilde{w}_j(x; \lambda) = \log\left(1 + \frac{x - 1}{x(x^j - 1)}\right) - \lambda \log x.$$

For a prime p, this is exactly $w_{p,j}(\lambda)$. For every layer j, the row proves

$$\widetilde{w}_j(c_j; \lambda) \ge 0, \qquad \widetilde{w}_j(c_j + 1; \lambda) \le 0,$$

using the rational logarithm bounds below. The gain decreases with the base, so these two boundary signs classify every prime in that layer. The extra check $\widetilde{w}_{d+1}(2; \lambda) \le 0$, together with monotonicity in the layer and base, excludes every event above the stored vector. Thus events inside $\mathcal{B}(c)$ have nonnegative weight and all events outside have nonpositive weight. Any integer's actual exponent packet is consequently bounded above by the sum over this finite box; this is the precise sense in which the box is a valid worst case.

Appendix B prints the complete first row, the scaled logarithm checker, and the two endpoint predicates. Appendix C lists all 36 intervals and their main structural parameters. Keeping those arithmetic audits separate here leaves the main argument focused on the row schema and its soundness theorem.

The logarithm certificates used here are elementary and explicit. For a rational $q \ge 1$, put $z = (q - 1)/(q + 1)$. Then $0 \le z < 1$ and, for $N \ge 1$,

$$\begin{aligned} L_N(q) &= 2 \sum_{i=0}^{N-1} \frac{z^{2i+1}}{2i + 1}, \\ U_N(q) &= L_N(q) + \frac{2z^{2N+1}}{1 - z^2}, \\ L_N(q) &\le \log q \le U_N(q). \end{aligned} \tag{6.3}$$

The upper remainder is deliberately coarse: it follows by deleting the denominators $2i + 1$ from the positive tail of the atanh series. Consequently both bounds in (6.3) are rational and their soundness is proved once for arbitrary q and N. Every row in this proof uses $N = 20$.

Certificate discovery versus certificate checking. A practical way to find a row is to choose rational L, U, λ and then, for layer j, choose an integer cutoff c_j. The checker compares the threshold formula at c_j and $c_j + 1$: the former must have nonnegative reduced weight and the latter nonpositive reduced weight. Monotonicity in the base then proves the signs for every prime in that layer. A final check at base 2 proves that no still higher layer can contribute positively. From the resulting finite box one computes the exact integer products

$$A = \prod (p^{j+1} - 1), \qquad B = \prod p(p^j - 1), \qquad Q = \prod p.$$

All three products are over $\mathcal{B}(c)$. Their logarithms give exactly

$$\sum_{(p,j)\in\mathcal{B}(c)} w_{p,j}(\lambda) = \log A - \log B - \lambda \log Q.$$

For example, Row 00 has $c = [17, 5, 3, 2, 2]$. Its layers contain respectively the prime bases

$$\{2, 3, 5, 7, 11, 13, 17\}, \quad \{2, 3, 5\}, \quad \{2, 3\}, \quad \{2\}, \quad \{2\}.$$

Direct multiplication gives

$$A = 8\,490\,305\,136\,604\,741\,632\,000, \qquad B = 1\,651\,300\,226\,181\,365\,760\,000,$$
$$Q = 367\,567\,200.$$

The kernel then verifies, among the remaining row predicates,

$$A \le \frac{5141587823941}{10^{12}} B, \qquad 2^{28}\frac{11486475}{8388608} \le Q.$$

Rational bounds of the forms $A \le (\text{gain upper})B$ and $2^e m \le Q$, together with (6.3), produce the objective and height bounds. If either endpoint inequality fails, one changes the slope or shortens the interval; accepted endpoints are then chained exactly. This search is only certificate discovery. It need not be trusted or even reproduced: the stored row contains all of its outputs, the general soundness lemma proves what successful checks imply, and Lean's kernel evaluates every row predicate from those visible rational and integer values.

The discovery process was iterative, not an optimization computation. An earlier checked cover contained 39 rows. The cancellation-aware coefficient revision in `208b447` lowered the analytic handoff from 100000 to 74500 and removed the last three rows, leaving the present 36-row witness. Rejected candidate rows were not retained as a stable trace, so no numerical count of trial proposals is claimed. Nor is 36 claimed minimal: it is an upper bound supplied by one accepted cover, and there is no nontrivial lower bound or optimality theorem for this certificate format. Different slopes, interval widths, or logarithm approximants could change the row count without changing the soundness argument.

The historical discovery runs were not instrumented per candidate, so the paper does not report invented search times or a trial count. Appendix C instead reports a reproducible quantity: the elapsed time for each accepted row module in one identified clean CI build. Those times measure elaboration and kernel checking on that runner, not intrinsic certificate complexity; they vary with hardware, toolchain, and cache state.

The height data certify that the relevant prime-base product is at least $2^e m$; the checker turns this into a lower logarithmic height with a 20-term rational log series. The gain bound is likewise converted into an upper logarithmic objective. Thus this row checks the two literal rational inequalities displayed in Lemma 6.3, not floating-point approximations. Its complete definition and kernel proof are `Robin1984/Finite/Certificates/Row00.lean`. The other rows have the same schema; the full list and gap-free cover proof are `Robin1984/Finite/Certificates/AllRows.lean`, with all individual records in `Robin1984/Finite/Certificates`.

Lemma 6.3 (Soundness of a finite row)**.** *Suppose a row $[L, U]$ satisfies its layer-sign, product, height, and endpoint inequalities. If $n > 5040$ and $L \le \log n \le U$, then $\mathcal{M}(n) < 0$.*

Proof. The sign checks and (4.5) bound the actual factorization packet by the finite threshold box. The exact product bounds convert that box into an upper bound B for $\log \mathcal{A}(n) - \lambda \log n$. The two endpoint checks state

$$B < \gamma + \log\log L - \lambda L, \qquad B < \gamma + \log\log U - \lambda U.$$

Concavity of $y \mapsto \log\log y - \lambda y$ gives the same strict inequality at $y = \log n$. Rearrangement is $\mathcal{M}(n) < 0$. □

Proposition 6.4 (Certified finite cover)**.** *Robin's inequality holds whenever*

$$720720 \le n, \qquad \log n < 74500. \tag{6.4}$$

Proof. Thirty-six rational rows form a gap-free chain from $336/25$ to $37283397387/500000$, an endpoint strictly larger than 74500. Since $336/25 < \log 720720$, they cover (6.4). Every row satisfies the hypotheses of Lemma 6.3. In the final retained row the first-layer product includes every prime at most 70849; it is split into 105 contiguous blocks, whose exact products concatenate to the required prefix. The blocks cover the half-open base range $[0, 70850)$ without a gap, and each block proves its numerator, denominator, and prime-base product. Concatenation therefore reconstructs the same first-layer factors in A, B, Q that a monolithic product would contain; the split changes evaluation cost, not the mathematical packet. All comparisons are integer or rational inequalities. Thus the list of rows and their checks constitute a finite proof of the proposition. □

Combining Propositions 6.1, 6.2, and 6.4 gives the following.

Theorem 6.5 (Finite completion)**.** *If $5040 < n$ and $\log n < 74500$, then Robin's inequality holds.*

In particular, Proposition 5.5 is now unconditional.

7. The implication under the Riemann hypothesis

Let

$$\vartheta(x) = \sum_{p \le x} \log p, \qquad \psi(x) = \sum_{p^k \le x} \log p.$$

The analytic estimate is most transparent after setting $H = \log n$ and

$$S(H) = H^{-1/2}(\log H)^{-1}. \tag{7.1}$$

The section is easiest read backward from its last line. We seek $\mathcal{M}(n) < 0$. Proposition 7.2 bounds this margin by an analytic debit $\mathcal{E}(H)$ minus an arithmetic credit $\mathcal{C}(H)$. The weighted explicit formula controls the debit; a complete block of prime squares supplies the credit; and Proposition 7.3 proves that, once $H \ge 74500$, the credit is at least $(113/50)S(H)$ while the debit is strictly smaller. The formulas below establish these three statements but do not add another logical layer.

For a first reading, Lemma 7.1 may be used as the analytic input and its kernel-estimate proof skipped; the narrative resumes at the definition of $\mathcal{C}(H)$ and $\mathcal{E}(H)$. The intervening integration-by-parts calculation is retained because it explains the exact constant c_0, the asymmetric error bound, and the otherwise delicate endpoint $m = 1$.

We first record the weighted explicit-formula input. For $m \ge 1$, define

$$w_m(t) = t^{-m-1}\frac{m\log t + 1}{(\log t)^2}, \qquad I_m(x) = \int_x^\infty (\psi(t) - t) w_m(t)\, dt. \tag{7.2}$$

For a nontrivial zero ρ of ζ, counted with multiplicity, put

$$K_m(\rho, x) = \int_x^\infty t^{\rho-m-1} \frac{m \log t + 1}{(\log t)^2}\, dt.$$

Lemma 7.1 (Robin's weighted explicit formula)**.** *Assume* RH. *For $m \geq 1$ and $x \geq 2$,*

$$I_m(x) = -\sum_\rho \frac{K_m(\rho, x)}{\rho} - T_m(x), \tag{7.3}$$

where the pole, gamma-factor, and trivial-zero correction is the explicit nonnegative integral

$$T_m(x) = \int_x^\infty w_m(t) \left(\log(2\pi) + \frac{1}{2}\log(1 - t^{-2})\right) dt. \tag{7.4}$$

For $t \geq 2$ the parenthesis lies between 0 and $\log(2\pi)$, and hence

$$0 \leq T_m(x) \leq \frac{\log(2\pi) x^{-m}}{\log x}. \tag{7.5}$$

If

$$c_0 = \gamma + 2 - \log(4\pi),$$

then

$$-B_m(x) - \frac{\log(2\pi) x^{-m}}{\log x} \leq I_m(x) \leq B_m(x), \tag{7.6}$$

where

$$B_m(x) = c_0 x^{1/2-m} (\log x)^{-1} \left(m + (\log x)^{-1} + \frac{4}{(2m-1)(\log x)^2}\right). \tag{7.7}$$

Proof. Apply the explicit formula to the continuous cutoff whose Mellin transform is $K_m(s, x)/s$. The arithmetic side is the von Mangoldt sum associated with $I_m(x)$; moving the vertical line separates the pole at 1, the nontrivial zeros, and the negative even zeros. Every interchange is justified by an explicit summable norm majorant, so the zero sum retains analytic multiplicity rather than choosing an enumeration without proof.

The kernel calculation is not a black box. Since

$$-\frac{d}{dt}\frac{t^{-m}}{\log t} = t^{-m-1}\frac{m \log t + 1}{(\log t)^2},$$

one integration by parts gives

$$\frac{K_m(\rho, x)}{\rho} = \frac{x^{\rho-m}}{\rho \log x} \int_x^\infty \frac{t^{\rho-m-1}}{\log t}\, dt. \tag{7.8}$$

A second integration by parts, with $d(t^{\rho-m}) = (\rho - m) t^{\rho-m-1}\, dt$, separates the leading term

$$\frac{m}{\rho(m-\rho)} \frac{x^{\rho-m}}{\log x}$$

from a remainder containing $(\log t)^{-2}$. Bounding that remainder once more by integration against $(\log t)^{-3}$ gives, under RH and hence $\Re\rho = 1/2$, the pointwise majorant

$$\left|\frac{K_m(\rho, x)}{\rho}\right| \leq \frac{1}{|\rho|^2}\, x^{1/2-m} (\log x)^{-1} \left(m + (\log x)^{-1} + \frac{4}{(2m-1)(\log x)^2}\right).$$

Hadamard factorization of the completed zeta function supplies the exact multiplicity-counted identity

$$\sum_\rho |\rho|^{-2} = \gamma + 2 - \log(4\pi) = c_0.$$

Summing the majorant gives $B_m(x)$. The combined pole/trivial-zero factor is exactly (7.4); its integrand bound and $\int_x^\infty w_m(t)\,dt = x^{-m}/\log x$ prove (7.5), explaining the asymmetry in (7.6).

The contour identity initially covers $m \geq 2$. For the endpoint let

$$q(t) = \frac{t(\log t + 1)}{2\log t + 1}, \qquad q'(t) = \frac{\log t(2\log t + 3)}{(2\log t + 1)^2}.$$

Exact integration by parts gives

$$I_1(x) = q(x)I_2(x) + \int_x^\infty q'(t)I_2(t)\,dt.$$

The same identity is proved separately for the zero series and for T_m, with absolute-integrability bounds licensing each interchange. Substitution of the $m = 2$ formula therefore yields (7.3) at $m = 1$, rather than obtaining it by a limiting convention. This is Robin's Lemma 2 with no omitted endpoint. □

Costa Pereira's estimates for $\psi - \vartheta$, with the corrected constants from [3], control the prime-power part separated from (7.3). Abel summation supplies the complementary lower bound for a complete block of prime squares.

Define the complete prime cost

$$\mathcal{C}(H) = \sum_{p \leq H/2} \min\left\{ \frac{\log p}{H\log H}, \frac{1}{p^2} \right\}. \tag{7.9}$$

Also define

$$\begin{aligned} \mathcal{E}(H) = (2 + c_0)\frac{H^{-1/2}}{\log H} + (c_0 - 2)\frac{H^{-1/2}}{(\log H)^2} \\ + (8 + 4c_0)\frac{H^{-1/2}}{(\log H)^3} + 2\frac{H^{-2/3}}{\log H} + \log(2\pi)\frac{H^{-1}}{\log H}. \end{aligned} \tag{7.10}$$

Thus $\mathcal{E}(H)$ is a uniform upper bound assembled from the weighted prime-function errors, whereas $\mathcal{C}(H)$ is a nonnegative deduction forced by completing the exponent sum to all small primes. Their names record their roles in (7.11); neither is a new number-theoretic function outside this proof.

Proposition 7.2 (Complete exponent budget)**.** *Assume* RH. *If* $n \geq 1$ *and* $H = \log n \geq 20000$, *then*

$$\mathcal{M}(n) \leq \mathcal{E}(H) - \mathcal{C}(H). \tag{7.11}$$

Proof. For a prime p occurring in n with exponent a, the telescoping identity from Lemma 4.1 and $\log y \leq y - 1$ give

$$\log \frac{1 - p^{-(a+1)}}{1 - p^{-1}} \leq \frac{1}{p} - p^{-(a+1)} - m_p, \qquad m_p = \log(1 - p^{-1}) + p^{-1} \leq 0. \tag{7.12}$$

Subtract $a\log p/(H\log H)$ from each prime contribution. Completing the sum to all primes $p \leq H$ does not lose the omitted primes: if $a = 0$ and $p \leq H/2$, the benchmark contribution pays at least the minimum in (7.9). Summation therefore yields a prime-reciprocal expression minus $\mathcal{C}(H)$.

Partial summation writes the prime-reciprocal expression as $\log\log H$ plus the Meissel–Mertens constant and a weighted ϑ-error. Replacing ϑ by ψ separates the nonnegative prime-power tail. Lemma 7.1 and the Costa Pereira bounds show that the remaining weighted error is at most (7.10). The Mertens constants cancel with the sum of the m_p, giving exactly (7.11). □

The numerical heart of the large-height proof is the following strict gap.

Proposition 7.3 (Explicit coefficient gap)**.** *Assume* RH. *For* $H \geq 74500$,

$$\mathcal{E}(H) < \frac{113}{50} S(H), \qquad \mathcal{C}(H) \geq \frac{113}{50} S(H). \tag{7.13}$$

Proof. For the first inequality, divide (7.10) by $S(H)$. The exact identity becomes

$$2 + c_0 + (c_0 - 2)(\log H)^{-1} + (8 + 4c_0)(\log H)^{-2} + 2H^{-1/6} + \log(2\pi) H^{-1/2}.$$

Uniform rational bounds valid for every $H \geq 74500$ control $(\log H)^{-1}$, $H^{-1/6}$, $H^{-1/2}$, c_0, and $\log(2\pi)$. The values actually used are

$$\begin{aligned} &\log H \geq \frac{56}{5}, \quad (\log H)^{-1} \leq \frac{5}{56}, \quad (\log H) H^{-1/6} \leq \frac{87}{50}, \\ &H^{-1/2} \leq \frac{1}{272}, \qquad c_0 \leq \frac{1}{20}, \qquad \log(2\pi) \leq \frac{46}{25}. \end{aligned} \tag{7.14}$$

There is no numerical conflict in the bound for c_0:

$$c_0 = \gamma + 2 - \log(4\pi) = 0.0461914179\ldots < 0.05 = \frac{1}{20}.$$

The proof does not depend on that decimal evaluation. The harmonic-number and logarithm bounds in Appendix A prove $c_0 < 154/3125 < 1/20$ rationally. Substitution into (7.14), while retaining the favorable negative coefficient $c_0 - 2$, gives the exact margin

$$\frac{\mathcal{E}(H)}{S(H)} \leq \frac{3010451}{1332800} = \frac{113}{50} - \frac{1677}{1332800} < \frac{113}{50}.$$

For the second inequality, take the complete prime-square block between $\sqrt{2H}$ and $H/2$. Abel summation gives a dominant normalized main term tending to $2\sqrt{2}$, while the prime-function error, endpoint, and prime-power terms are lower order. The calculation retains the RH error at the far endpoint instead of replacing the whole ϑ tail by a coarse global bound; only after that cancellation is the remaining nonnegative prime-power contribution discarded. Appendix A gives the complete normalized expression and every rational substitution. The resulting exact lower margin is

$$\frac{\mathcal{C}(H)}{S(H)} \geq \frac{15537277889}{6868000000} = \frac{113}{50} + \frac{15597889}{6868000000} > \frac{113}{50}.$$

Thus the two sides of the gap are proved independently and meet with positive rational margins. □

Theorem 7.4 (The RH implication)**.** *If* RH *holds, then Robin's inequality holds for every* $n > 5040$.

Proof. If $\log n < 74500$, use Theorem 6.5. Otherwise combine Propositions 7.2 and 7.3:

$$\mathcal{M}(n) < \left(\frac{113}{50} - \frac{113}{50} \right) S(\log n) = 0.$$

This is equivalent to Robin's inequality. □

8. The Nicolas–Landau oscillation

Define Nicolas's finite product and normalized function by

$$\mathfrak{P}(x) = \prod_{p \leq x} \left(1 - \frac{1}{p}\right), \qquad F(x) = e^{\gamma} \log \vartheta(x)\, \mathfrak{P}(x). \tag{8.1}$$

The signed quantity used below is $\log F(x)$.

Let

$$k(t) = \frac{(\log t)^{-1} + (\log t)^{-2}}{t^2},$$

and put

$$K(x) = \int_x^\infty (\vartheta(t) - t)k(t)\,dt, \qquad J(x) = \int_x^\infty (\psi(t) - t)k(t)\,dt. \tag{8.2}$$

For functions $g, h : \mathbb{R} \to \mathbb{R}$, write $g = \Omega_-(h)$ at infinity if there is $c > 0$ such that, for every X, some $x \geq X$ satisfies $g(x) \leq -ch(x)$.

Only two outputs of this section are needed later. First, $\log F(x)$ is bounded above by the weighted prime error $J(x)$ plus a term of order x^{-1}. Second, if RH fails, Landau's principle forces $J(x)$ below a negative multiple of x^{-b} arbitrarily far out for some $b < 1/2$. Since $x^{-1} = o(x^{-b})$, the first output transfers the second one to $\log F$. Section 9 then absorbs the remaining LCM losses, which are only of order $P^{-1/2} \log P$.

Lemma 8.1 (Finite Nicolas comparison)**.** *For all sufficiently large x,*

$$\log F(x) \leq J(x) + \frac{4}{x}. \tag{8.3}$$

Proof. The exact decomposition, with no asymptotic notation, is

$$\begin{aligned} \log F(x) &= R_M(x) + R_\vartheta(x) + K(x), \\ R_M(x) &= \sum_{0<n\leq\lfloor x\rfloor} \Big(\log(1 - n^{-1}) + n^{-1}\Big) - (M - \gamma), \\ R_\vartheta(x) &= \log\log\vartheta(x) - \log\log x - \frac{\vartheta(x) - x}{x\log x}, \end{aligned} \tag{8.4}$$

where the summand is understood to be zero for nonprimes and M is the Meissel–Mertens constant in the convergent prime summand series. The explicit tail estimate gives $|R_M(x)| \leq 4/x$. The inequality $\log y \leq y - 1$, applied to the ratio of the two logarithmic heights, is exactly $R_\vartheta(x) \leq 0$ after simplification.

Finally

$$J(x) - K(x) = \int_x^\infty (\psi(t) - \vartheta(t))k(t)\,dt \geq 0,$$

because both factors are nonnegative. Applying these three bounds to (8.4) gives (8.3). □

Proposition 8.2 (Nicolas–Landau negative oscillation)**.** *If* RH *is false, then there is b with*

$$0 < b < \frac{1}{2} \tag{8.5}$$

such that

$$J(x) = \Omega_-(x^{-b}). \tag{8.6}$$

Proof. Failure of RH, together with the functional equation and conjugation symmetry, gives a nontrivial zero ρ of ζ with $1/2 < \Re\rho < 1$. On its horizontal line choose a zero ρ_* with maximal real part. Choose

$$1 - \Re\rho_* < b < \frac{1}{2}.$$

The contradiction is an instance of a simple positivity obstruction. If the corrected tail $J(x) + x^{-b}$ were eventually positive, its Mellin transform would be the Laplace transform, under $x = e^y$, of a positive measure. Positivity makes convergence propagate to every real point that the analytic continuation can cross. The zeta logarithmic derivative, however, gives that continuation

a pole at $s = 1 - \rho_*$, strictly inside the propagated half-plane. The remainder of the proof makes each part of this obstruction precise.

Suppose (8.6) fails. Then, beyond some X,

$$J(x) + x^{-b} > 0. \tag{8.7}$$

Define the positive measure

$$d\mu_b(x) = \mathbf{1}_{(X,\infty)}(x)x^{-1}\big(J(x) + x^{-b}\big)\,dx$$

and use $Y(x) = \log x$ as the random variable. Its complex moment-generating function is the Mellin transform

$$M_b(s) = \int_X^\infty x^{s-1}\big(J(x) + x^{-b}\big)\,dx. \tag{8.8}$$

For $\Re s < 0$ the integral is computed directly. Integration by parts in the J term expresses it through the Mellin transform of $\psi(x) - x$ and hence through $-\zeta'(1-s)/\zeta(1-s)$ after its pole at 1 is removed. The added power has the exact transform

$$\int_X^\infty x^{s-b-1}\,dx = \frac{X^{s-b}}{b-s} \qquad (\Re s < b).$$

These formulas define an analytic continuation $H_b(s)$ wherever the zeta factor is regular.

Here is the transform calculation used in that sentence. The Frullani identity

$$k(t) = \int_0^\infty (u+1)t^{-(u+2)}\,du$$

and Fubini's theorem give, for $\Re s < 0$,

$$\int_3^\infty x^{s-1}J(x)\,dx = \int_0^\infty (u+1)\frac{\Phi_3(u+1-s) - 3^s\Phi_3(u+1)}{s}\,du, \tag{8.9}$$

where

$$\Phi_3(z) = \left(\frac{-\zeta'(z)}{z\zeta(z)} - \frac{1}{z-1}\right) - \int_1^3 (\psi(t) - t)t^{-z-1}\,dt. \tag{8.10}$$

The apparent value at $s = 0$ is filled by analytic continuation. Absolute integrability of the three-variable carrier $3 < x < t$, $u > 0$ licenses the two exchanges of integration in (8.9); compact and large-u majorants then prove analyticity of the filled expression. Formula (8.10) makes the singularity source explicit: the finite startup integral is entire, while the logarithmic derivative records the nontrivial zeros of ζ.

We now separate the two roles of Landau's argument. First, positivity of μ_b makes the real convergence set of (8.8) convex. At an interior point a, the derivatives of M_b are its nonnegative moments

$$\int (\log x)^k e^{a\log x}\,d\mu_b(x).$$

If an analytic continuation has a power-series radius reaching $a + d$, its Taylor series bounds the sum of these moments at $d \geq 0$; monotone convergence then proves actual integrability at $a + d$. Iterating this strict crossing shows that every real $\sigma < b$ lies in the interior of the true integrability set. Consequently M_b itself is analytic throughout the half-plane $\Re s < b$ and agrees there with H_b by the identity theorem.

Second, put $s_0 = 1 - \rho_*$. Since $\Re s_0 = 1 - \Re\rho_* < b$, the preceding conclusion makes M_b analytic, and therefore locally bounded, at s_0. Approach that point along

$$s_\delta = 1 - \rho_* - \delta, \qquad \delta > 0.$$

Maximality of $\Re\rho_*$ on its horizontal line says that $\zeta(\rho_*+\delta)\neq 0$, so $H_b(s_\delta)$ is regular for every $\delta>0$. At the endpoint, however, $1-s_0=\rho_*$ and the logarithmic derivative has the zero pole; the norm of $H_b(s_\delta)$ tends to infinity as $\delta\downarrow 0$. Since $M_b(s_\delta)=H_b(s_\delta)$, analytic boundedness and divergence contradict one another. Thus the assumed eventual positivity is impossible, proving (8.6). □

The proof just given is the strict form of the oscillation argument used by Nicolas and Robin. Its quantifiers are important: the excursions occur arbitrarily far out, rather than at one selected numerical frontier. The positive-transform step is not imported as an unnamed principle. Its reusable formal statement proves that if a nonnegative exponential transform converges at every real exponent below a frontier and has an analytic continuation across the frontier, then it actually converges at some strictly larger exponent. The proof expands the continuation into its power series, identifies its coefficients with nonnegative moments, and uses monotone convergence to recover exponential integrability. This generic theorem is linked explicitly in Section 11; the subsequent Nicolas modules supply (8.9), the rightmost-zero ray, and the blow-up contradiction.

Corollary 8.3. *If* RH *is false, there exists* $0<b<1/2$ *such that*

$$\log F(x)=\Omega_-(x^{-b}). \tag{8.11}$$

Proof. Use Lemma 8.1 and $x^{-1}=o(x^{-b})$, which holds because $b<1$. □

9. Least common multiples and the converse

Let

$$L_P=\operatorname{lcm}(1,2,\dots,P).$$

The classical identity $\log L_P=\psi(P)$ connects this finite integer with the analytic frontier. Define the logarithmic height loss

$$D(P)=\log\log\psi(P)-\log\log\vartheta(P). \tag{9.1}$$

The saturated Euler product differs from the actual abundancy of L_P by the nonnegative prime-power reserve

$$T(P)=\sum_{p\le P}-\log\Bigl(1-p^{-a_p-1}\Bigr),\qquad a_p=\max\{a:p^a\le P\}. \tag{9.2}$$

Lemma 9.1 (Exact LCM identity)**.** *For all sufficiently large* P*,*

$$\mathcal{M}(L_P)=-\log F(P)-D(P)-T(P). \tag{9.3}$$

Proof. The finite Euler product gives

$$\sum_{p\le P}-\log(1-p^{-1})=-\log\mathfrak{P}(P),$$

where $\mathfrak{P}(P)$ denotes the product in (8.1). The prime-power factorization of L_P gives

$$-\log\mathfrak{P}(P)=\log\mathcal{A}(L_P)+T(P).$$

Now expand $\log F(P)$ and use $\log L_P=\psi(P)$. Rearrangement is exactly (9.3). □

Lemma 9.2 (The transfer losses)**.** *For every* $b<1/2$*,*

$$D(P)+T(P)=o(P^{-b}). \tag{9.4}$$

Proof. The elementary bound

$$0 \le \psi(P) - \vartheta(P) \le 2\sqrt{P}\log P$$

and a linear lower bound for $\vartheta(P)$ give

$$0 \le D(P) \ll \frac{\log P}{\sqrt{P}}.$$

Splitting the reserve at $\sqrt{P}$ and summing the geometric tails gives $T(P) \ll P^{-1/2}$. Since $P^{-1/2}\log P = o(P^{-b})$ whenever $b < 1/2$, the claim follows. □

Theorem 9.3 (The converse implication)**.** *If Robin's inequality holds for every $n > 5040$, then* RH *holds.*

Proof. Suppose RH is false. Choose b as in Corollary 8.3. Negative excursions of $\log F(P)$ give positive excursions of $-\log F(P)$ of order P^{-b}. By Lemma 9.2, subtracting $D(P) + T(P)$ preserves such excursions. Lemma 9.1 therefore gives arbitrarily large P with $\mathcal{M}(L_P) > 0$. Since $L_P \ge P$, these least common multiples eventually exceed 5040, and $\mathcal{M}(L_P) > 0$ is a literal failure of Robin's inequality. This contradicts the hypothesis. □

10. Proofs of the main theorems

Proof of Theorem 1.1. The implication RH $\Rightarrow$(1.1) is Theorem 7.4. The reverse implication is Theorem 9.3. □

Proof of Theorem 1.2. If RH holds, Theorem 1.1 gives Robin's inequality for every integer above 5040, hence in particular for every colossally abundant one. Conversely, assume it for every colossally abundant integer above 5040. The finite startup theorem and Proposition 5.5 extend it to every integer above 5040. Theorem 1.1 then gives RH. □

11. Formal verification and reproducibility

11.1. **Public statements and the checked proof graph.** The preceding proof has been formalized in Lean 4.33.1 using Mathlib [6, 8]. The two public declarations have the mathematical types

$$(\forall n \in \mathbb{N},\ 5040 < n \Rightarrow \sigma(n) < e^{\gamma} n \log\log n) \iff \mathrm{RH},$$

$$\mathrm{RH} \iff (\forall n, \varepsilon,\ 5040 < n \Rightarrow n \text{ is colossally abundant at } \varepsilon \Rightarrow \sigma(n) < e^{\gamma} n \log\log n).$$

They are named

`Robin1984.robin_inequality_iff_riemannHypothesis`
`Robin1984.riemannHypothesis_iff_colossallyAbundant_robin`.

The small Palomar challenge file fixes these statements using Mathlib only; the solution file supplies the proof terms, so statement comparison is independent of the internal proof graph.

Here the symbol RH is not an informal placeholder. Expanding Mathlib's definition, it is the proposition

$$\forall s \in \mathbb{C},\quad \zeta(s) = 0 \implies (\nexists k \in \mathbb{N},\ s = -2(k+1)) \implies s \neq 1 \implies \Re s = \frac{1}{2}.$$

Thus the public biconditional uses the usual statement that every nontrivial zero of the analytically continued zeta function lies on the critical line. The converse proof uses the functional equation and conjugation symmetry to turn a violation of this predicate into a zero with $1/2 < \Re s < 1$. No alternate project definition of RH is introduced. The exact reduction from the negation of Mathlib's predicate to such a right-half zero is proved in `NicolasLandau`.

The shortest route through the checked declarations is the following table. An arrow here means use as a theorem, not merely an import.

mathematical branch	checked Lean route
public boundary	`Challenge` ↔ `Solution` by exact statement comparison
equivalence assembly	`Equivalence` combines the forward and converse declarations
forward implication	`LargeHeightRobin` above 74500, and `FiniteComplete` below it
converse implication	`NicolasLandauRobinBridge` combines the negative oscillation with the LCM transfer
CA restriction	`CAReduction` combines startup completion with tangent coverage

Thus the full biconditional does not emerge from one opaque analytic theorem: the finite and large-height forward branches meet only in the final case split, and the converse is a separate contradiction branch.

The two advertised biconditionals also have different quantifier surfaces but no gap between them. From RH, Theorem 1.1 immediately specializes to every colossally abundant integer. Conversely, the CA hypothesis and Proposition 5.5 extend the inequality from those extrema to every integer above 5040; Theorem 1.1 then returns RH. Thus the second result is assembled from the all-integer criterion and the independently proved extremal reduction; Sections 6 and 7 are not being mistaken for proofs only about CA integers.

For the technically dense statements, the correspondence is declaration level rather than merely thematic:

paper statement	checked Lean declaration and relation
Lemma 7.1	`robinPsiWeightedErrorIntegral_eq_zero_sum_correction_all` is (7.3) with (7.4); `robinPsiWeightedErrorIntegral_bounds_all` is (7.6), both in `RobinLemmaTwo`
Proposition 7.2	`robin_log_abundancy_le_complete_prime_cost` in `AbundancyExponentBound` is the same inequality after unfolding $\mathcal{M}$, $\mathcal{E}$, and $\mathcal{C}$
Proposition 7.3	`robinMertensWeightedScalar_lt_113_div_50` and `robin_minimum_prime_cost_ge_large_scalar` in the linked `Mertens coefficient` and `prime-cost coefficient` modules are the two inequalities in (7.13)
Lemma 8.1	`nicolasLogMertensOscillation_le_J_add_four_div` in `NicolasOscillation` is (8.3)
Landau crossing	`exists_integrableExpSet_gt_of_analyticContinuationAt` in `MGFAnalyticContinuation` is the generic positive-transform theorem
Proposition 8.2	`exists_nicolasJ_omegaMinus_of_not_riemannHypothesis` in `NicolasLandauRightmostRay` has exactly the quantifiers in (8.5)–(8.6)
Theorem 9.3	`not_nativeRobinInequalityAll_of_not_riemannHypothesis` in `NicolasLandauRobinBridge` is its contrapositive

11.2. **How the checked artifact was produced.** What the final proof says and how it was produced are different questions. The former is answered by the declarations above and by the mathematical proof in Sections 5–9. The production process was a source-guided reconstruction with the following stages.

(1) The public statements were fixed first in `Challenge`, directly over Mathlib definitions and without importing any project module. This prevented later proof engineering from silently changing the target.

(2) Robin's argument and the cited inputs from Nicolas, Landau, Alaoglu– Erdős, Rosser–Schoenfeld, and Costa Pereira were decomposed into named intermediate statements. Source-specific results remained in their mathematical families, while reusable order, integration, logarithm, Mellin-transform, and prime-power lemmas were separated from the final theorem assembly.

(3) Analytic identities were proved symbolically before numerical constants were introduced. Infinite sums and integrals were accompanied by explicit summability or integrability proofs; the weighted zero sums retained analytic multiplicity.

(4) The finite-range checker and its soundness theorem were proved before certificate data were accepted. Candidate rows and prime blocks could be found by ordinary computation, but only their exact integer and rational outputs entered Lean, where kernel reduction checked them.

(5) The resulting source graph was reduced to the declarations needed by the advertised results and the public convenience facades. Direct-import, provenance, placeholder, metadata, statement-boundary, and independent replay checks were then applied to the completed graph.

Human direction selected the theorem statements, mathematical route, source correspondences, and final assertions. AI-assisted proof engineering helped with implementation, refactoring, certificate organization, and audits. The production history is not itself a proof: correctness rests on the resulting proof terms, exact certificates, and the validation pipeline described below.

11.3. **Automated research and formalization context system.** The AI assistance was not a single prompt followed by unchecked code generation. The project used a persistent, structured context system that connected four kinds of state:

(1) *Research state* associated claims with exact source passages, bibliographic identities, chronology, and provenance classifications. Later reconstructions and formalization choices were kept distinct from results stated by Robin, Nicolas, Landau, and the other cited authors.

(2) *Mathematical state* recorded exact definitions, quantifier signatures, sign conventions, normalizations, numerical domains, and the current hard lemma connecting an open branch to the advertised theorem. Numerical experiments were attached only as evidence for a matching domain, never promoted to hypotheses.

(3) *Formal state* tracked the exact Lean declaration under attack, its direct and transitive dependencies, import providers, proof-state obligations, prior failed approaches, focused-build evidence, and the declarations on the final dependency path.

(4) *Assurance state* tracked placeholder and axiom audits, provenance coverage, direct-import checks, certificate coverage, statement-boundary comparison, kernel replay, and release status. Substantive mistakes were converted into explicit preflight rules before the same class of action could be repeated.

Reusable task-specific *skills* formed the control layer over that state. A skill is a versioned workflow package activated for a class of work, not a learned mathematical oracle. The packages used here covered:

- source and provenance review, including chronology and the distinction between direct source formalization and project-original glue;
- hard-theorem discipline, which fixes the live declaration, forbids replacing a universal parameter by an unjustified finite experiment, and separates proved results from certificate evidence and open obligations;

- numerical and certificate work, requiring the proposed inequality and its domain to be checked independently before Lean formalization and requiring every retained certificate to pass a general soundness theorem;
- repository engineering, including exact import providers, dependency closure, placeholder scans, focused builds, and staged-surface checks; and
- paper and artifact production, requiring source-to-declaration correspondence, exact-link checks, repeated PDF builds, extracted-text checks, and page-by-page visual inspection.

The context router activated the smallest applicable set and supplied each skill with the same live theorem identity and evidence ledger. Skills could compose—for example, a numerical bound proposed during proof work could not enter the paper until the numerical-domain rule, Lean proof rule, provenance rule, and document-verification rule all agreed on its status. Their durable preflight and hard-stop conditions also carried lessons from a failed build or incorrect claim into later sessions, instead of relying on conversational memory alone.

The skill library also evolved during the project. When a build failure, mathematical contradiction, inaccurate claim, or user correction exposed a substantive failure class, work on the affected path stopped. The system recorded four fields: the triggering condition, a check that must run before the action is attempted again, the action now prohibited, and the corrected action. A repeated failure strengthened the entry with a mechanically tested hard stop. Rules that applied only to one declaration remained in the project context; rules that generalized across formalizations were promoted into, or used to revise, a reusable skill. Subsequent sessions loaded the revised package before acting, so failure-derived knowledge changed future behavior even when the earlier conversation was no longer present.

Worked mathematical case: closing the large-height gap. One retained task packet fixed the upper-error and lower-cost Lean targets in `Mertens coefficient` and `prime-cost coefficient`, the universal domain $H \geq 74500$, the normalization $S(H)$, and the prohibition on adding an assumption or changing the public theorem. The prior checked route began at $H = 100000$. The automated proposal that survived review had three parts: retain the favorable sign of $(c_0 - 2)/\log H$ instead of taking an absolute value; complete a block of prime squares between $\sqrt{2H}$ and $H/2$; and carry the far-end RH error through the Abel-summation cancellation before discarding nonnegative terms. It then proposed separate rational subgoals for every monotone factor.

Human acceptance was conditional on two independent checks: the normalized upper and lower expressions had to be numerically true on the stated domain, and Lean had to prove their symbolic monotonicity and exact rational margins. The accepted output is not the proposal text but the declarations in `MertensExplicitCoefficient` and `PrimeCostExplicitCoefficient`, together with the fractions printed in Appendix A. Commit `208b447` records the resulting 100000-to-74500 handoff and the removal of three now-unneeded certificate rows. The numerical check was evidence; only the Lean proofs and subsequent kernel checks were accepted into the artifact.

Worked failure-to-rule case: theorem and import boundaries. During submission hardening, a Comparator failure reported that the Challenge and Solution theorem statements did not match even though their displayed formulas appeared identical. The failure class was dependence on implicit or project-local declaration identity at the public boundary. The corrected rule requires `autoImplicit false`, permits only Mathlib imports in `Challenge.lean`, defines the statement vocabulary directly there, and requires exported declaration equality before release. The public changes are visible in commits `9612f5e` and `160cd08`.

A second failure occurred after import narrowing: modules using `Nat.lcmUpto` no longer imported its defining provider and failed with an unknown-constant error. The corresponding

rule now resolves the exact provider, requires a direct import in every consumer, and builds the first affected consumer before any full CI run. Commits `1d60111` and `0544629` apply the correction. The mechanical hard stop is encoded in `check-imports.ps1`; the statement hard stop is Comparator itself. A later complete successful run exercised the Lean build, documentation, Comparator, Lean replay, NanoDa replay, and release gate.

These examples expose the task inputs, the accepted decomposition, the human acceptance criterion, the retained source output, and the recurrence barrier. They are the two documented workflow cases evaluated here, not a complete telemetry record of every exploratory interaction. No baseline run without the system was performed, so the paper makes no causal claim that automation reduced elapsed time, human effort, or the total number of errors. The public before-and-after commits, exact proof terms, check scripts, and CI logs make the narrower claims about retained outputs and recurrence barriers testable.

This evolution is controlled rather than autonomous self-modification of the proof environment. Skill revisions are explicit, reviewable instruction artifacts; they are checked on representative tasks and remain subordinate to human direction and repository policy. They may alter search order, preflight checks, context selection, or stopping conditions, but they cannot alter a theorem statement, convert evidence into a proof, whitelist a new axiom, or bypass Lean's kernel and the independent replay gates. The public artifact therefore depends on the skills for engineering discipline, not for logical soundness.

At each stage, the system selected only the source excerpts, definitions, theorem types, dependency fragments, and earlier failure evidence relevant to the live obligation. This constrained context mattered more than raw context volume: it reduced namespace confusion, accidental quantifier changes, unjustified numerical specialization, duplicate lemmas, and work on declarations outside the final proof graph. Candidate formal steps were tested first against the smallest relevant module; expensive full builds were run only for coherent repository states.

Human control remained at the mathematical boundary. The maintainer chose the target theorems, sources, proof route, attribution, and public claims. Automated agents did propose intermediate definitions and lemma decompositions, search library interfaces, construct candidate proof terms, generate exact certificate candidates, refactor dependencies, and run audits. Human review selected which definitions, statements, and proofs were retained. Every accepted output is visible as Lean source, rational data, provenance metadata, or a validation log; no model response, research summary, numerical search, confidence score, or context record is trusted by either advertised theorem.

The methodological claim is therefore limited to the architecture and the two cases above. Reproducing the checked result requires the public source and certificates, not a replay of the exploratory interaction history: the build, Comparator, and two kernel implementations check the retained artifact.

11.4. **Reusable formalization outputs.** Beyond the two final biconditionals, the development leaves several units intended for reuse:

- a generic Landau positive-transform theorem in the `ProbabilityTheory` namespace, proving strict extension of exponential integrability from analytic continuation and nonnegative moments;
- multiplicity-aware zero-divisor and summability interfaces for completed zeta, together with explicit Mellin/zero-series interchange lemmas;
- exact rational logarithm bounds and a theorem-parametric finite-row checker that separates certificate discovery from proof;

- reusable bounds for prime-power layers, least-common-multiple towers, and normalized explicit coefficients; and
- provenance, dependency, Challenge/Solution, replay, and release gates for a large formal-mathematics artifact.

The first and other project-independent units are isolated under `Robin1984/Mathlib`; their proposed upstream destinations are recorded in `MATHLIB_PORTING.md`. The result is therefore both a checked reconstruction of Robin's theorem and a case study in turning source mathematics, exploratory computation, and AI-assisted proof engineering into a narrowly trusted formal artifact.

For an AI or machine-learning audience, the relevant contribution is not a claim that a model resolved an open conjecture. It is a controlled human–agent workflow whose endpoint is falsifiable at declaration level: fixed targets, explicit dependencies, exact certificates, recorded failure classes, and independent kernel acceptance. No new general-purpose Lean tactic is claimed. The reusable outputs are the mathematical interfaces, certificate architecture, context-and-skill discipline, and validation pipeline described above.

11.5. **Project-local mathematics and provenance.** The formalization status of the principal arguments is as follows.

- The Robin biconditional and CA restriction are project-local compositions in `Solution` and `Equivalence.Theorem`.
- The weighted explicit formula, including $m = 1$, is a project-local source formalization in `weighted formula` and $m = 1$ `endpoint`.
- The Nicolas–Landau negative oscillation is a project-local source formalization in `oscillation proof`.
- The LCM transfer and converse are project-local source formalizations in `LCM bridge`.
- Finite-row soundness and all row data are project-original certificate formalizations in `row soundness` and `all rows`.

The dependency boundary can be summarized without inspecting import syntax:

source	mathematical role
Mathlib	definitions of ζ, ξ, and RH; complex and real analysis, measure and integration theory, infinite sums, arithmetic, and elementary asymptotics
pinned PNTAnd fork	Hadamard-factorization and divisor infrastructure, completed-zeta facts, Mertens and prime-function infrastructure, and the explicit Rosser–Schoenfeld and Costa Pereira estimates
Robin1984	the weighted formula in the form used here, its zero-series interchanges and $m = 1$ endpoint, the complete exponent budget and coefficient gap, the Nicolas–Landau continuation and oscillation, the LCM transfer, the CA reduction, and every finite certificate

It would therefore be incorrect to describe every dependency as proved from scratch in repository-owned modules. Mathlib and the pinned PNTAnd revision supply a large body of foundational and explicit analytic number theory. The project-local claim is narrower and checkable: the source-specific chain from those imported results to Robin's two advertised biconditionals, including the weighted endpoint, Nicolas–Landau continuation, LCM transfer, CA reduction, and finite certificates, is proved here without adding a new axiom.

In particular, Lemma 7.1 and Proposition 8.2 are not assumptions imported as opaque theorems: their full proofs occur in the linked repository-owned modules. The lower-level imported analysis is substantial, as it should be for a theorem about zeta zeros, but no imported declaration states

Robin's biconditional or the Nicolas–Landau conclusion needed here. The exact distinction is recorded declaration by declaration in the provenance metadata.

The formal development separates arithmetic identities, colossally abundant extrema, the RH estimates, the Nicolas–Landau argument, finite certificates, and final assembly. Reusable order, logarithm, Mellin-transform, and prime-power lemmas are exposed in their mathematical families. Each owned Lean file carries a content-reviewed provenance label distinguishing direct source formalization, another published source, standard mathematics without a claimed individual originator, and primarily project-authored formalization work. The complete reviewed ledger is distributed with the repository as the `provenance guide` and its `machine-readable ledger`.

For more detailed reading, the assembled `startup transition`, the `Mertens coefficient`, the `prime-cost coefficient`, and the `finite certificate data` correspond directly to Sections 5–7. These source links follow the main branch for browsing; each release archives the exact commit from which its paper and Lean artifacts were built.

11.6. **Exact finite certificates.** The finite proof contains no floating-point oracle. The startup factorizations and the 36 interval rows reduce to exact natural-number and rational inequalities. The large final retained prime product is split into 105 blocks only to keep individual kernel computations manageable. The certificate checkers are proved sound once; each data row is then accepted by kernel reduction.

Appendix B quotes the checked interface and expands Row 00 into its stored record, scaled logarithm bounds, product tests, and endpoint predicates. The conceptual separation is simple: one ordinary theorem proves that any successful row implies Robin's inequality throughout its interval; another declaration asks Lean's kernel to evaluate the visible rational row. The aggregate theorem checks that the accepted endpoints form a gap-free cover. The full definitions are in `FiniteRowCertificate` and `Row00`.

11.7. **What Comparator and the kernel replays establish.** The validation stages have different jobs. The ordinary Lean build checks that every project module elaborates and that Lean's kernel accepts the final proof terms. Comparator then protects the public theorem boundary against a different risk: a purported solution might prove an easier proposition with the same informal description, or redefine a symbol occurring in the target.

The trusted `Challenge` module imports only four narrow Mathlib modules. It defines the divisor sum, Robin inequality, and CA predicate directly, and states the two advertised theorems with deliberate proof holes. The `Solution` module imports the project and supplies the actual proofs. The checked `Comparator configuration` names exactly those two theorems, permits only `propext`, `Quot.sound`, and `Classical.choice`, requires NanoDa, and declares no definition holes. In particular, the solution is not allowed to choose a convenient meaning for σ, Robin's inequality, CA, or RH.

The complete semantic policy is short enough to display:

```
"theorem_names": [
  "Robin1984.robin_inequality_iff_riemannHypothesis",
  "Robin1984.riemannHypothesis_iff_colossallyAbundant_robin"
],
"definition_names": [],
"permitted_axioms": ["propext", "Quot.sound", "Classical.choice"],
"enable_nanoda": true
```

The omitted JSON fields only name the modules `Challenge` and `Solution`; there is no additional whitelist of project definitions.

Using pinned tool revisions, the Comparator job performs the following checks.

(1) It builds and exports the Challenge and Solution environments separately under the Landrun sandbox.

(2) For each advertised theorem it compares the exported statement and every declaration on which that statement depends. This is a comparison of Lean declarations, not of displayed text or theorem names alone.

(3) It follows the Solution theorem bodies transitively and rejects any axiom outside the three permitted standard principles.

(4) It replays the exported Solution environment through Lean's default kernel.

(5) Because `enable_nanoda` is true, it also submits the export to NanoDa, an independent kernel implementation.

Passing Comparator therefore establishes statement identity, the configured axiom bound, and kernel acceptance through two implementations. It does not establish that the trusted Challenge statement is the historically intended reading of Robin's paper, nor that the exposition is persuasive. The explicit Challenge definitions, the expansion of Mathlib's RH predicate above, and ordinary mathematical review address that separate question.

11.8. **Axioms and dependency boundary.** The repository contains no project-local axiom declarations. Outside the two deliberate statement holes in the Palomar challenge, project proof sources contain no `sorry`, `admit`, or `native_decide`. Diagnostic commands such as `#print axioms` occur neither in repository-owned Lean sources nor in the timed external import closure. A completed kernel audit of both advertised declarations returned exactly the standard principles

`propext`, `Classical.choice`, `Quot.sound`.

This audit follows the proof terms of the two advertised declarations through their actual transitive closure. It therefore distinguishes library files that are merely present in a dependency checkout from declarations genuinely used by this proof: an unused theorem elsewhere cannot enter the result's axiom set.

The ordinary `lake build` command is the full formalization build. Its Lake-native default target reads a checked-in topological order of 257 modules, schedules one module job at a time, bounds each Lean process to one internal task thread, and ends with `Solution`. Thus the submitted project itself controls the memory-intensive finite-certificate order; no external runner or local patching step is part of the proof build.

All dependencies are pinned. The required Lean 4.33.1 changes to `PrimeNumberTheoremAnd` are published in the public fork `https://github.com/kimihiro64/PrimeNumberTheoremAnd` at commit `f8f58c7`, based on upstream commit `47fa486`. The full hashes are recorded in the Lake manifest. Lake reconstructs this dependency directly from public Git history; no local patching step is required.

The fork is intended as a temporary compatibility and extension branch, not as a permanent divergence. At the time of this revision, the upstream `main` toolchain file still declares Lean 4.32.2, and upstream `main` does not yet contain the xi-divisor source used here. Replacing the fork therefore requires two things: an upstream revision compatible with the project's Lean toolchain, and upstream versions or equivalent replacements for the additional xi-divisor and Mertens interfaces.

The fork has two distinct roles. First, it updates proof scripts and package pins for Lean 4.33.1 and exposes named consequences of the existing Mertens development needed by this project. These changes repair elaboration against the new toolchain; they do not replace an analytic theorem by an assumption or alter Mathlib's definition of RH. Second, it adds the multiplicity-aware

interface used by the weighted zero sums. The principal declarations in `RiemannXiDivisorZeros` are:

- `RiemannXiDivisorZeroIndex` and `riemannXiDivisorZeroValue` give an index type and value map for the nonzero divisor of ξ, retaining analytic multiplicity;
- `riemannXiDivisorZeroValue_ne_zero` and `riemannXiDivisorZeroValue_eq_zero` prove that every indexed value is nonzero and is an actual zero of ξ;
- `summable_riemannXiDivisorZero_norm_inv_sq` proves summability of the inverse-square norms of those values;
- `completedRiemannZeta_eq_cpow_mul_Gamma_mul_riemannZeta` expands completed zeta under its stated nonvanishing hypotheses; and
- `neg_two_mul_logDeriv_riemannXi_zero_eq` proves the exact identity $-2(\xi'/\xi)(0) = \gamma + 2 - \log(4\pi)$.

The companion `RobinXiZeroSummability` proves, as `summable_riemannXiDivisorZero_norm_inv_rpow`, the fractional inverse-norm summability used to interchange the zero sum and Mellin integral. These results are derived from the dependency's existing divisor support, Hadamard growth and summability, completed-zeta, and differentiability theorems. The fork adds no axiom, proof placeholder, or numerical oracle. The xi-divisor files retain Matteo Cipollina's authorship and Apache-2.0 notice; the upstream project is due to Alex Kontorovich, Terence Tao, and its contributors. The complete fork diff is public.

Project-independent helper results are isolated under `Robin1984/Mathlib` and mapped to proposed upstream Mathlib paths in `MATHLIB_PORTING.md`. The inventory records which files are ready but not yet submitted. The intended maintenance path is to upstream those generic units separately and to propose the xi-divisor interface and compatibility repairs to `PrimeNumberTheoremAnd`; until acceptance, the pinned public fork is the reproducible source of the checked dependency. When both replacement conditions are met, the project can repin directly to upstream, but only after the full Lean build and Comparator/NanoDa pipeline pass against that revision.

11.9. **Reproducing and independently checking the result.** The audit record is public and readable without repository membership or GitHub authentication:

Source and history.:
`https://github.com/kimihiro64/Robin1984`

Complete CI runs and job logs.:
`https://github.com/kimihiro64/Robin1984/actions/workflows/ci.yml`

Tagged source, paper, API site, compiled Lean tree, and checksums.:
`https://github.com/kimihiro64/Robin1984/releases`

Machine-readable theorem and provenance metadata.:
`formalization.yaml`, `comparator.json`, and `provenance/ledger.json`

Trusted statement boundary and implementation.:
`Challenge.lean` and `Solution.lean`

The string `kimihiro64` is simply the public GitHub owner namespace, not an indication of a private repository. Each release is tied to an exact Git commit and is therefore a citable immutable snapshot even though the `main` branch continues to develop. A DOI-bearing archival mirror would add another persistence layer, but is not required to inspect or reproduce the present artifact.

There are three useful levels of reproduction. First, a reader can inspect a completed run without compiling Lean. The public CI history exposes the separate job results and logs. After

every required job succeeds, the workflow publishes a GitHub release containing the paper, an offline API site, and the complete Linux `.lake/build` tree. The release identifies the exact source commit and records checksums for its assets. The compiled objects are a convenience for inspection and compatible incremental use; the source and pinned revisions remain authoritative.

Second, a source-level kernel rebuild can be made from a clean release tag:

```
git clone https://github.com/kimihiro64/Robin1984.git
cd Robin1984
git checkout TAG
lake update
lake exe cache get
lake build
```

Here `TAG` is the tag named on the chosen GitHub release. `lake update` reconstructs the exact dependency revisions from the manifest, and `lake exe cache get` downloads the matching Mathlib cache; omitting the cache step gives a larger cold build. The final `lake build` follows the repository's checked-in sequential module order and ends by building `Solution`. On Windows, `bootstrap.ps1` installs and checks the same project prerequisites before these commands.

The repository-level structural checks can then be run as

```
pwsh -File scripts/check-imports.ps1
pwsh -File scripts/check-provenance.ps1
ruby scripts/validate-formalization.rb
```

They check narrow imports, the candidate-module boundary, per-file provenance and documentation, and the Palomar metadata. They are consistency checks, not substitutes for kernel checking.

Third, on Linux with Git, Lake, Go, Rust/Cargo, Python 3, and Landrun, the independent statement and replay pipeline is

```
./scripts/verify-comparator.sh
```

The checked-in `verification script` checks out exact reviewed revisions of Comparator, lean4export, Landrun, and NanoDa; requires the project's Lean 4.33.1 toolchain; performs the sequential project build; and runs the five Comparator stages listed above. The checked-in `CI workflow` is the complete executable recipe: it additionally validates licensing and release metadata, builds the paper and API documentation, transfers the completed Lean artifacts to the downstream jobs, and publishes a release only if every dependency job, including Comparator, succeeds.

Observed cost and hardware. The hosted workflow uses GitHub's standard public `ubuntu-latest` runner, which at the time of this revision provides four x64 CPU cores, 16 GB RAM, and 14 GB SSD storage. This is a tested sufficient environment, not a proven minimum; peak memory below 16 GB and a smaller storage floor have not been benchmarked. The project build is deliberately serialized to control memory, so a clean `lake build` should be budgeted at roughly two hours; one measured cold build took about 1 hour 42 minutes.

The independent pipeline is longer. In the representative source-changing successful run of commit `4d0b59e`, the Lean job took 1 hour 16 minutes, API documentation 16 minutes 58 seconds, Comparator and its two replays 2 hours 39 minutes, and total wall time was 3 hours 56 minutes 59 seconds because downstream jobs overlap. In a cache-reusing successful run, the Lean job fell to 15 minutes 36 seconds, while Comparator still took 2 hours 17 minutes. Accordingly, an independent team should budget about four hours for a full uncached assurance run, not merely the duration of `lake build`. Source changes invalidate the project-object cache; paper-only changes can reuse it. The release's compiled tree avoids the first rebuild for a matching environment but does not replace source-level verification.

Portability and maintenance. The authoritative artifact is the source at an immutable release tag. The downloadable `.olean` tree is a Linux x64 convenience built for Lean 4.33.1; compiled objects are not expected to work across Lean versions, platforms, or incompatible dependency revisions. A source rebuild uses `lean-toolchain` and `lake-manifest.json`, which currently pin Mathlib commit `0df444a360eaa60ab8c11dca51a86af692955474` and the PNTAnd fork commit `f8f58c749d6cde8a641348fcd5e4702993651cd6`. The Windows bootstrap covers the main Lean build; the full Comparator/Landrun/NanoDa recipe is presently Linux-oriented and additionally requires the tools listed above.

Maintenance is deliberate rather than automatic. A toolchain or dependency repin is tested first on the affected modules and then must pass the complete Lean, provenance, Comparator, Lean-kernel, and NanoDa gates before release. Generic project-owned modules are tracked for upstreaming in `MATHLIB_PORTING.md`. As verified for this revision, PNTAnd upstream `main` still declares Lean 4.32.2 and lacks the xi-divisor interface, so the public fork remains necessary. The fork will be removed only when both compatibility and interface replacements exist and the full pipeline passes after repinning. Future maintenance may be required as Lean evolves; immutable tagged releases preserve the exact historical inputs rather than promising that old compiled objects will load in new toolchains.

The formalization was developed with human mathematical direction and AI-assisted proof engineering. Lean's kernel checks that the linked proof terms inhabit the exact public types above, using only the reported logical principles. It does not by itself certify that those types are the intended reading of Robin's article, that the reconstruction is historically faithful, that the result is novel, or that the exposition is persuasive. Those are mathematical and scholarly review obligations. The explicit expansion of Mathlib's RH predicate, the source-to-module table, and the provenance ledger make that comparison inspectable rather than implicit. The present status is self-assessed; no independent expert review is claimed.

Appendix A. Exact rational coefficient audit

This appendix collects the arithmetic intentionally omitted from the main narrative of Proposition 7.3. First, the formal bound for the zero constant uses

$$\gamma < H_{256} - \log 256, \qquad H_{256} \le \frac{15311}{2500}, \qquad \log 2 \ge \frac{69314}{100000}, \qquad \log(4\pi) \ge \frac{253}{100}.$$

Consequently

$$c_0 < \frac{15311}{2500} - 8\frac{69314}{100000} + 2 - \frac{253}{100} = \frac{154}{3125} < \frac{1}{20}.$$

This is the rational proof behind the decimal orientation $c_0 = 0.0461914179\ldots$.

For the error bound, put $q = (\log H)^{-1}$. The $H^{-1/6}$ estimate in (7.14) turns $2H^{-1/6}$ into at most $(87/25)q$. Combining this with the favorable term $-2q$ and applying the remaining bounds gives

$$\begin{aligned}\frac{\mathcal{E}(H)}{S(H)} &\le 2 + \frac{37}{25}\frac{5}{56} + 8\left(\frac{5}{56}\right)^2 \\ &\quad + \frac{1}{20}\left(1 + \frac{5}{56} + 4\left(\frac{5}{56}\right)^2\right) + \frac{46}{25}\frac{1}{272} \\ &= \frac{3010451}{1332800} = \frac{113}{50} - \frac{1677}{1332800} < \frac{113}{50}.\end{aligned}$$

For the complete prime cost, Abel summation over the prime-square block $\sqrt{2H} < p \le H/2$ gives the following normalized lower expression:

$$\begin{aligned}
&2\sqrt{2}\frac{\log H}{\log H+\log 2}\left(1-\frac{1}{\log H+\log 2}\right)\\
&\quad-\frac{9}{4}\,2^{1/4}H^{-1/4}\frac{\log H}{\log H+\log 2}-\log(2\pi)H^{-1/2}\frac{\log H}{\log H+\log 2}\\
&\quad-4H^{-1/2}\frac{\log H}{\log H-\log 2}\\
&\quad-2\sqrt{2}\,c_0H^{-1}\frac{\log H}{\log H-\log 2}\left(2+\frac{1}{\log H-\log 2}+\frac{4}{3(\log H-\log 2)^2}\right).
\end{aligned}$$

The last line is the far-end RH term retained through the cancellation. The rational bounds used for $H \ge 74500$ are

$$\begin{aligned}
\sqrt{2} &\ge \frac{7071}{5000}, & 2^{1/4} &\le \frac{119}{100},\\
\frac{\log H}{\log H+\log 2} &\ge \frac{16}{17}, & \frac{1}{\log H+\log 2} &\le \frac{17}{202},\\
H^{-1/4}\frac{\log H}{\log H+\log 2} &\le \frac{573}{10000}, & H^{-1/2}\frac{\log H}{\log H+\log 2} &\le \frac{7}{2000},\\
H^{-1/2}\frac{\log H}{\log H-\log 2} &\le \frac{197}{50000}, & \frac{\log H}{\log H-\log 2} &\le \frac{107}{100}.
\end{aligned} \tag{A.1}$$

The normalized far-end term is at most $1/10000$, so the final exact substitution is

$$\begin{aligned}
&2\frac{7071}{5000}\frac{16}{17}\left(1-\frac{17}{202}\right)-\frac{9}{4}\frac{119}{100}\frac{573}{10000}-\frac{46}{25}\frac{7}{2000}\\
&\qquad-4\frac{197}{50000}-\frac{1}{10000}\\
&\qquad=\frac{15537277889}{6868000000}=\frac{113}{50}+\frac{15597889}{6868000000}>\frac{113}{50}.
\end{aligned}$$

Appendix B. A finite certificate in full

The first of the 36 accepted rows stores the following exact record:

field	stored rational value
$[L,U]$	$[336/25,\ 22117363/1000000]$
λ	9771914406623/500000000000000
layer cutoffs	$[17,5,3,2,2]$
gain upper bound	5141587823941/1000000000000
height data (e,m)	$(28,\ 11486475/8388608)$
lower bounds for $\log L, \log U$	$1299117667/500000000,\ 619272591/200000000$

The essential checked interface, with one proof body elided, is

```
theorem RobinFiniteRowChecks.sound {r : RobinFiniteRow}
    (h : RobinFiniteRowChecks r) {n : Nat} (hn : 5040 < n)
    (hLo : (r.lo : Real) <= Real.log (n : Real))
    (hHi : Real.log (n : Real) <= (r.hi : Real)) :
    Robin1984.Core.NativeRobinInequality n := by ...

theorem robinFiniteRow00_checks :
```

```
  RobinFiniteRowChecks robinFiniteRow00 := by
decide +kernel
```

Let $g = 5141587823941/10^{12}$ be the stored gain bound and $k(q) = \lfloor \log_2 q \rfloor$. In terms of the functions in (6.3), the scaled upper checker is

$$U_{20}^{\mathrm{sc}}(q) = k(q)\frac{693147180559945310}{10^{18}} + U_{20}\left(\frac{q}{2^{k(q)}}\right),$$

and the lower checker uses $693147180559945309/10^{18}$ and L_{20}. With the stored λ, exponent $e = 28$, and mantissa $m = 11486475/8388608$, the objective supplied to both endpoints is

$$B_0 = U_{20}^{\mathrm{sc}}(g) - \lambda\left(28\frac{693147180559945309}{10^{18}} + L_{20}(m)\right).$$

The kernel decision verifies the integer product inequalities $A \le gB$ and $2^{28}m \le Q$, every adjacent layer-sign comparison, and

$$B_0 < \frac{57721}{100000} + L_{20}^{\mathrm{sc}}(\ell_L) - \lambda L, \qquad B_0 < \frac{57721}{100000} + L_{20}^{\mathrm{sc}}(\ell_U) - \lambda U.$$

The general logarithm and Euler-constant theorems then convert these rational facts into the real endpoint inequalities used by Lemma 6.3. The record and its kernel proof are `Row00`; the reusable soundness theorem is `FiniteRowCertificate`.

Appendix C. Complete finite-cover inventory

Table 1 lists every accepted row. The columns L, U and λ are the exact stored rationals. The entry c_1/d gives the first-layer cutoff and the number of stored layers; "blocks" is the number of separately checked first-layer prime-product blocks, with zero denoting the direct monolithic calculation used by Row 00. The final column is the elapsed module time reported by the clean CI build of commit `4d0b59e`. These are observed elaboration and kernel-checking times on that runner, not discovery times or hardware-independent complexity bounds.

Table 1: The complete 36-row certificate cover.

row	L	U	λ	c_1/d	blocks	seconds
00	$\frac{336}{25}$	$\frac{22117363}{1000000}$	$\frac{9771914406623}{500000000000000}$	17/5	0	2.3
01	$\frac{22117363}{1000000}$	$\frac{35788699}{1000000}$	$\frac{2565510579241}{250000000000000}$	28/6	2	2.3
02	$\frac{35788699}{1000000}$	$\frac{29149999}{500000}$	$\frac{5519615043151}{1000000000000000}$	46/7	3	2.4
03	$\frac{29149999}{500000}$	$\frac{47544747}{500000}$	$\frac{3004424868557}{1000000000000000}$	76/7	4	2.4
04	$\frac{47544747}{500000}$	$\frac{146801491}{1000000}$	$\frac{1724212981141}{1000000000000000}$	120/8	5	2.5
05	$\frac{146801491}{1000000}$	$\frac{111686563}{500000}$	$\frac{8278925801}{8000000000000}$	184/9	6	2.5
06	$\frac{111686563}{500000}$	$\frac{163594351}{500000}$	$\frac{80829194299}{125000000000000}$	274/10	7	2.6
07	$\frac{163594351}{500000}$	$\frac{464815043}{1000000}$	$\frac{422180858753}{1000000000000000}$	395/10	8	2.6
08	$\frac{464815043}{1000000}$	$\frac{130165559}{200000}$	$\frac{70868018799}{250000000000000}$	557/11	9	2.7
09	$\frac{130165559}{200000}$	$\frac{893036263}{1000000}$	$\frac{38967393937}{200000000000000}$	771/11	10	2.7
10	$\frac{893036263}{1000000}$	$\frac{240770673}{200000}$	$\frac{137136588803}{1000000000000000}$	1048/12	12	2.8
11	$\frac{240770673}{200000}$	$\frac{63176719}{40000}$	$\frac{99275427189}{1000000000000000}$	1391/12	13	2.8
12	$\frac{63176719}{40000}$	$\frac{2035337871}{1000000}$	$\frac{73775311097}{1000000000000000}$	1806/13	14	2.9
13	$\frac{2035337871}{1000000}$	$\frac{518900301}{200000}$	$\frac{11152011487}{200000000000000}$	2314/13	16	2.9
14	$\frac{518900301}{200000}$	$\frac{815385537}{250000}$	$\frac{42786768497}{1000000000000000}$	2927/14	17	3.1
15	$\frac{815385537}{250000}$	$\frac{406392661}{100000}$	$\frac{6654185407}{200000000000000}$	3662/14	19	3.1
16	$\frac{406392661}{100000}$	$\frac{1248690297}{250000}$	$\frac{2622639723}{100000000000000}$	4528/14	21	3.1
17	$\frac{1248690297}{250000}$	$\frac{3045165617}{500000}$	$\frac{20930175389}{1000000000000000}$	5542/15	23	3.3
18	$\frac{3045165617}{500000}$	$\frac{458996953}{62500}$	$\frac{16893531041}{1000000000000000}$	6716/15	25	3.4
19	$\frac{458996953}{62500}$	$\frac{4387795691}{500000}$	$\frac{13794103247}{1000000000000000}$	8059/15	27	3.5
20	$\frac{4387795691}{500000}$	$\frac{10433152137}{1000000}$	$\frac{11354367573}{1000000000000000}$	9603/15	30	3.6
21	$\frac{10433152137}{1000000}$	$\frac{12279993269}{1000000}$	$\frac{2357543377}{250000000000000}$	11356/16	33	3.8
22	$\frac{12279993269}{1000000}$	$\frac{179906601}{12500}$	$\frac{1973616211}{250000000000000}$	13335/16	36	4.0
23	$\frac{179906601}{12500}$	$\frac{16784073643}{1000000}$	$\frac{664793849}{1000000000000000}$	15587/16	39	4.3
24	$\frac{16784073643}{1000000}$	$\frac{9718846411}{500000}$	$\frac{1126354703}{200000000000000}$	18110/16	42	4.6
25	$\frac{9718846411}{500000}$	$\frac{22433188173}{1000000}$	$\frac{1200244709}{250000000000000}$	20934/17	46	4.9
26	$\frac{22433188173}{1000000}$	$\frac{1287242339}{50000}$	$\frac{2057220229}{500000000000000}$	24088/17	50	5.4
27	$\frac{1287242339}{50000}$	$\frac{7349340853}{250000}$	$\frac{3547339421}{1000000000000000}$	27570/17	54	5.8
28	$\frac{7349340853}{250000}$	$\frac{1338267981}{40000}$	$\frac{3072760897}{1000000000000000}$	31426/17	59	6.4
29	$\frac{1338267981}{40000}$	$\frac{9481137173}{250000}$	$\frac{2672852847}{1000000000000000}$	35690/18	64	7.1
30	$\frac{9481137173}{250000}$	$\frac{1070418171}{25000}$	$\frac{2335546837}{1000000000000000}$	40370/18	70	8.2
31	$\frac{1070418171}{25000}$	$\frac{48110318623}{1000000}$	$\frac{2050940953}{1000000000000000}$	45463/18	76	9.6
32	$\frac{48110318623}{1000000}$	$\frac{10789326747}{200000}$	$\frac{1807809003}{1000000000000000}$	51028/18	82	10
33	$\frac{10789326747}{200000}$	$\frac{30165067587}{500000}$	$\frac{1597827091}{1000000000000000}$	57137/18	89	12
34	$\frac{30165067587}{500000}$	$\frac{33566442549}{500000}$	$\frac{1418388113}{1000000000000000}$	63731/18	97	15
35	$\frac{33566442549}{500000}$	$\frac{37283397387}{500000}$	$\frac{157973183}{125000000000000}$	70849/19	105	19

The endpoint list satisfies $U_i = L_{i+1}$ for every adjacent pair. Two declarations in the linked `AllRows module` complete the aggregate check:

- `robinFiniteRows_cover` checks the gap-free equalities and outer endpoints by kernel reduction;
- `robinFiniteRows_checks` dispatches every row to its separately checked predicate.

The exact remaining cutoffs, gain bound, height mantissa and exponent, and endpoint logarithm bounds are stored in the linked `certificate directory`.

Appendix D. Notation map

symbol	meaning	first definition
$\sigma(n)$	Sum of the positive divisors of n.	Section 2
$\mathcal{A}(n)$	Abundancy ratio $\sigma(n)/n$.	Section 2
$Q_\varepsilon(n)$	Penalized abundancy $\mathcal{A}(n)n^{-\varepsilon}$; its global maximizers are colossally abundant at parameter ε.	(1.2)
$\mathcal{M}(n)$	Logarithmic Robin margin; Robin's inequality is $\mathcal{M}(n) < 0$.	(4.1)
$g_{p,j}, \tau_{p,j}$	Marginal log gain for raising p^{j-1} to p^j, and its gain per unit log height.	(4.2)
$\lambda(n)$	Tangent slope $1/(\log n, \log\log n)$.	(5.1)
ε_0	Exact parameter $\log(12/11)/\log 11$ tying 5040 and 55440.	(5.4)
$\mathcal{B}(c)$	Finite prime-layer box determined by the cutoff vector c.	Section 6
L_N, U_N	Rational lower and upper bounds for the logarithm.	(6.3)
H	Logarithmic height $\log n$ in the analytic argument.	Section 7
$S(H)$	Natural RH comparison scale $H^{-1/2}/\log H$.	(7.1)
I_m, K_m, T_m, B_m	Weighted prime error, one-zero kernel, explicit correction, and its RH majorant.	(7.2)–(7.7)
c_0	Multiplicity-counted inverse-square zero sum $\gamma + 2 - \log(4\pi)$.	(7.7)
$\mathcal{E}(H)$	Uniform analytic debit in the complete exponent budget.	(7.10)
$\mathcal{C}(H)$	Nonnegative arithmetic cost forced by omitted primes or finite exponents.	(7.9)
ϑ, ψ	Chebyshev prime and prime-power counting functions.	Section 7
$\mathfrak{P}(x), F(x)$	Nicolas's finite Euler product and its normalized oscillation function.	(8.1)
k, K, J	Positive kernel and the weighted ϑ- and ψ-error tails used in the Landau argument.	(8.2)
L_P	Least common multiple of $1, \dots, P$.	Section 9
$D(P), T(P)$	Log-height loss and finite prime-power reserve in the LCM transfer.	(9.1)–(9.2)

Acknowledgements and attribution

The central mathematical result is Guy Robin's. The oscillation input and finite product F are due to Jean-Louis Nicolas, and the positivity principle used in that argument is due to Edmund Landau. The colossally abundant framework is due to Leonidas Alaoglu and Paul Erdős. Explicit prime-function estimates used in the reconstruction are credited to J. Barkley Rosser, Lowell Schoenfeld, and N. Costa Pereira, including Pereira's corrigendum. The formal reconstruction, certificate design, and exposition are the work of the Robin1984 project; credits for external Lean infrastructure are recorded above and in the repository metadata.

References

[1] L. Alaoglu and P. Erdős, *On highly composite and similar numbers*, Trans. Amer. Math. Soc. 56 (1944), 448–469.

[2] N. Costa Pereira, *Estimates for the Chebyshev function $\psi(x) - \vartheta(x)$*, Math. Comp. 44 (1985), no. 169, 211–221, doi:10.2307/2007805.

[3] N. Costa Pereira, *Corrigendum to "Estimates for the Chebyshev function $\psi(x) - \vartheta(x)$"*, Math. Comp. 48 (1987), no. 177, 447.

[4] B. Gomes and A. Kontorovich, *Riemann hypothesis in Lean (with or without Mathlib)*, formal development, 2020, https://github.com/AlexKontorovich/Lean-RH.

[5] E. Landau, *Über einen Satz von Tschebyschef*, Math. Ann. 61 (1905), 527–550, https://eudml.org/doc/158244.

[6] L. de Moura and S. Ullrich, *The Lean 4 theorem prover and programming language*, in: Automated Deduction–CADE 28, Lecture Notes in Computer Science 12699, Springer, 2021, 625–635.

[7] D. Loeffler and M. Stoll, *Formalizing zeta and L-functions in Lean*, Ann. Formalized Math. 1 (2025), doi:10.46298/afm.15328.

[8] The Mathlib Community, *The Lean mathematical library*, CPP 2020, 367–381.

[9] J.-L. Nicolas, *Petites valeurs de la fonction d'Euler*, J. Number Theory 17 (1983), no. 3, 375–388, doi:10.1016/0022-314X(83)90055-0.

[10] J.-L. Nicolas, *Small values of the Euler function and the Riemann hypothesis*, Acta Arith. 155 (2012), no. 3, 311–321, doi:10.4064/aa155-3-7.

[11] The PrimeNumberTheoremAnd contributors, *PrimeNumberTheoremAnd: a Lean formalization of the prime number theorem and related results*, source repository, https://github.com/AlexKontorovich/PrimeNumberTheoremAnd.

[12] G. Robin, *Grandes valeurs de la fonction somme des diviseurs et hypothèse de Riemann*, J. Math. Pures Appl. (9) 63 (1984), no. 2, 187–213, MR0774171 (in French).

[13] J. B. Rosser and L. Schoenfeld, *Approximate formulas for some functions of prime numbers*, Illinois J. Math. 6 (1962), 64–94, doi:10.1215/ijm/1255631807.