The Three-Prime Running Least Common Multiple

A two-prime transcendence theorem, exact three-prime structure, and a conditional residue criterion

Will Cook

July 2026

Abstract

The main contribution is a complete answer when the prime set has size two. De-duplication means retaining one reciprocal for each distinct value of the running least common multiple, rather than repeating that reciprocal at every smooth integer. For every pair of distinct primes, both the de-duplicated sum and Erdős's original sum with repeated running-LCM values are transcendental. Ordering the two pure-power channels by a Beatty word expresses the de-duplicated sum as a nonconstant algebraic affine function of a quantity A that is itself an affine function of a Hecke–Mahler value; the repeated sum factorises as $A(1+B)$ and is a nonconstant quadratic in the same A. The transcendence engine is Loxton and van der Poorten's 1977 theorem, of which the Bugeaud–Laurent statement quoted here is the modern form. These are the unconditional transcendence conclusions of the note. Both are paper proofs using that external theorem, not Lean-formalised results, and neither applies to any three-prime instance. They are also not the first public proof: Steve Fan posted the same reduction and conclusion on the erdosproblems.com discussion page for this problem on 26 June 2026, before this note was finalised. The argument here was found independently, and this note claims no priority for it.

For three pairwise distinct primes p, q, r, the contribution is exact finite structure, not an irrationality theorem. Writing $\mathrm{L}(x)$ for the least common multiple of the $\{p, q, r\}$-smooth numbers not exceeding x, we Lean-check

$$\mathrm{L}(x) = p^{\lfloor \log_p x \rfloor} q^{\lfloor \log_q x \rfloor} r^{\lfloor \log_r x \rfloor} \qquad (x \geq 1).$$

This yields the logarithmic cells and jump ratios, a finite height-fibre normal form, and the shell bound $9\,\#\delta \leq (j+3)^2$ under the stated ordered box hypothesis. Two statements in the smallest case $\{2, 3, 5\}$ are exact. First, group the jumps of L, that is, the points at which its value increases, into the blocks cut out by consecutive powers of two. Each open block $(2^a, 2^{a+1})$ contains at most one pure 3-power and at most one pure 5-power, so the block radix β_a, the product of the terminal factor 2 with a factor 3 when the block contains an internal 3-power and a factor 5 when it contains an internal 5-power, belongs to the exact four-element alphabet $\{2, 6, 10, 30\}$, rather than merely lying in a coarse bounded interval. All four letters already occur among the first five blocks: $\beta_0 = 2$, $\beta_1 = 6$, $\beta_2 = 10$ and

Authorship and AI use. The prose, formal proofs, and software in this version were drafted and revised by large-language-model agents under Will Cook's direction. Cook set the objectives and acceptance criteria, reviewed and authorised the manuscript, and is responsible for its contents. The agents are production tools, not authors. See *Statements and declarations*.

$\beta_4 = 30$. Second, the reciprocal of the running least common multiple at a smooth point is not a product of functions of the three exponents separately: on the smallest nontrivial two-by-two exponent rectangle its determinant is exactly $-1/15$. Thus that matrix has rank two, so the kernel is not a single product of three one-variable functions on that rectangle and no argument may assume such a form for it.

Problem #269 remains open from three primes onward: we prove no irrationality or transcendence statement in any three-prime case. What we prove towards it is conditional: after cancellation of the $\{2,3,5\}$-smooth part of a hypothetical denominator, the remaining denominator B is coprime to 30, and a denominator-dependent cofinal local-window residue escape rules out every positive integer solution d_n of the cleared recurrence $d_{n+1} = b_n d_n - B m_n$ that obeys the bound clearing the denominator delivers, where (b_n) is a radix word and (m_n) a forcing word. Lean checks the algebraic core of that reduction, namely that a fixed $\{2,3,5\}$-smooth factor is absorbed into the running height and that a common factor cancels from the carry recurrence, together with the resulting implication. An integer-only checker constructs the corresponding block digits, reproduces three small certificates, and finds an escaping window of length at most 18 for each of 106,666 tested denominator/start pairs. This is finite experimental evidence, not a cofinal theorem. Two problem-specific statements are not proved here: the rationality-to-carry identification, including the divisibility needed to perform that cancellation, and the cofinal escape property itself.

1 Introduction

Let P be a finite set of primes with $|P| \geq 2$, and let $a_1 < a_2 < \cdots$ enumerate the positive integers all of whose prime factors lie in P. Erdős Problem #269 asks whether

$$\sum_{n \geq 1} \frac{1}{[a_1, \ldots, a_n]}$$

is irrational, where $[a_1, \ldots, a_n]$ is the least common multiple [1, p. 65][2, p. 106]. Numbering and current reported status follow Bloom's catalogue, whose current record reproduces the displayed problem and explicitly warns that its OPEN label is the website owner's present assessment and may omit relevant literature [4]. The original publications, not the catalogue, carry the mathematical claims. The universal problem remains open, but this note settles every two-prime instance; unresolved finite cases begin with $|P| = 3$. In the primary 1988 source, Erdős states the infinite-P assertion as a simple exercise and presents persistence for a finite number of primes greater than one as a probable extension, not as a theorem; the current open status is supplied by the catalogue rather than inferred from that conjectural wording.

Three things are known and fix the shape of the question. The restriction $|P| \geq 2$ is necessary: for $P = \{p\}$ the enumeration is $a_n = p^{n-1}$, so $[a_1, \ldots, a_n] = p^{n-1}$ and the sum is $p/(p-1)$, which is rational. For that reason we use the modern restriction $|P| \geq 2$; the 1974 letter itself writes "given primes $p_1, \ldots, p_r$" without explicitly restricting r. For infinite P the sum is always irrational, which Erdős calls a simple exercise [2, p. 106]. And in a letter of 1 January 1973 he recorded that he could prove irrationality once duplicate running-LCM values are removed [3, p. 335]; the one-page letter states this result but does not include its proof.

That last remark is the one this note is closest to. The running least common multiple is not injective in n: it is constant along stretches of the enumeration and changes only at certain points. Removing duplicate summands means summing over the distinct values rather than over n. Section 4 proves that both this de-duplicated sum and the original repeated sum are transcendental when $|P| = 2$, by reducing them to a Hecke–Mahler value whose transcendence goes back to Loxton and van der Poorten [10, Theorem 8], quoted here in the modern form of Bugeaud and Laurent's Theorem 1.1 [9, p. 61, Theorem 1.1]. This is an independent route, not a recovery of the unprinted argument in the letter, and it is not the first public proof: Steve Fan posted the same factorisation, the same Hecke–Mahler reduction, and the same conclusion in the discussion thread of the problem's page on 26 June 2026 [11], with follow-up remarks there extending the argument to arbitrary coprime pairs. Sections 3 and 5 make the finite ingredients of the three-prime reindexing exact: they identify where the value is constant, by exactly what factor it changes when it changes, and give a finite rectangular-box fibre identity. They do not establish the ordered three-prime infinite passage and imply no irrationality result for three primes.

Throughout, p, q, r are pairwise distinct primes, and a *pure power* is a power b^e of a single base; a *pure* 3*-power* is a power of 3, and similarly for the other bases. Call n *smooth* when $n = p^i q^j r^k$ for some $i, j, k \geq 0$; this is the smooth lattice value. For $x \geq 1$ write

$$\mathrm{L}(x) = \operatorname{lcm}\{\, n \leq x : \ n \text{ smooth} \,\}, \qquad \mathrm{H}(x) = p^{\lfloor \log_p x \rfloor} q^{\lfloor \log_q x \rfloor} r^{\lfloor \log_r x \rfloor},$$

the running least common multiple and the pure-power height, where $\lfloor \log_b x \rfloor$ is the integer logarithm, the largest e with $b^e \leq x$. Since $a_1, \ldots, a_n$ are exactly the smooth numbers up to a_n, we have $[a_1, \ldots, a_n] = \mathrm{L}(a_n)$, and the summands of the problem are the reciprocals of L along the enumeration. The reciprocal of the height at a smooth point is the lattice kernel

$$\mathrm{K}(i, j, k) = \frac{1}{\mathrm{H}(p^i q^j r^k)}.$$

For $b \in \{p, q, r\}$ we call the set of positive powers of b the b*-channel*.

We treat $|P| = 3$ throughout, writing $P = \{p, q, r\}$, and the smallest instance is $\{2, 3, 5\}$. Two of the statements proved here are unconditional and exact rather than approximate. Theorem 8.1 determines exactly the alphabet of the dyadic compression of the *jump word*, that is, of the sequence of prime multipliers of L read in increasing order of the points where L increases (Section 3): the radix β_a of the block between 2^a and 2^{a+1}, meaning the product of the multipliers that occur in that block, takes one of the four values $2, 6, 10, 30$ and no others. Theorem 7.1 shows that the smallest two-by-two restriction of the kernel at $\{2, 3, 5\}$ has determinant $-1/15$ and hence rank two. In particular, it is not a product $f(i)g(j)h(k)$, so no argument may assume that form for the kernel on that rectangle. A rank-two matrix is itself a sum of two rank-one matrices, so decompositions into several separable terms are not excluded; what is excluded is the single product form. The mechanism is a small computation: $\mathrm{H}(6) = 4 \cdot 3 \cdot 5 = 60$ and not 6, because the running least common multiple at a smooth cutoff already sees powers of the other primes that the cutoff itself does not contain. The two-prime proof of Section 4 does not require such a separation.

The identification $\mathrm{L}(x) = \mathrm{H}(x)$ of Theorem 2.1 is what makes every statement about L below computable from three integer logarithms. In the unrestricted case that identification follows by iterating the prime-exponent maximum rule for least common multiples: $\operatorname{lcm}(1,\dots,N) = \prod_{t\le N} t^{\lfloor \log_t N\rfloor}$, the product being over the primes $t \le N$ [6, Ex. 1.21(a), p. 22]. Chebyshev's function is the corresponding prime-power sum, so equivalently $\log \operatorname{lcm}(1,\dots,N) = \psi(N)$ [6, §4.2, p. 75]; Montgomery–Vaughan state this exact identity directly in Exercise 6.2.7 [7, p. 183]. We record the P-restricted form because every later statement in the formal development is derived from it, and we claim nothing new for it.

The line of argument developed below has four stages. First, identify L exactly as a product of three pure powers, so that every later statement is a statement about three integer logarithms. Second, compress its jump word into dyadic blocks, whose radix takes only the four values of Theorem 8.1. Third, assume the sum rational with reduced denominator D, factor $D = D_{\mathrm{sm}}B$ where every prime divisor of D_{sm} lies in $\{2,3,5\}$ and $\gcd(B,30) = 1$, prove that the integral carry states, the integers the argument tracks from step to step (Section 8), share the factor D_{sm}, cancel it, and obtain a positive reduced carry d_n bounded by a denominator-dependent $K(B,n)$ and satisfying $d_{n+1} = b_n d_n - Bm_n$. Fourth, find a stretch of consecutive steps over which the accumulated base and forcing put the residue of that carry above the bound, which is impossible. The first two stages and the finite core of the fourth are proved and formalised here. The algebraic absorption-and-cancellation core of the third stage is also checked, but the problem-specific claim that the actual carry shares D_{sm} is not. That rationality-to-carry instantiation and the cofinal existence of windows in the fourth stage are the two obligations of Section 9. The third and fourth stages follow a pattern standard in irrationality proofs; what is specific here is that the radix word driving the carry is the block alphabet of Theorem 8.1.

The statement of Problem #269 has been formalised before, as a conjecture with an unfilled proof, in the *Formal Conjectures* collection [5]. That is a formal statement of the question up to a harmless rational normalisation: its Nat-indexed series includes the empty-prefix least-common-multiple term. Its rational, irrational, and infinite-prime assertions all end in `sorry`. The declarations described below are propositions about the objects the question is posed over. No claim of priority is made for any of these declarations. Kovač and Tao [8] treat several irrationality problems of Erdős for series of unit fractions by elementary means; nothing from that work is used here. We offer no numerical evidence about the value of the sum itself.

Status. The problem treated here is open, and this note does not close it. Every statement below marked as checked is a proposition that the pinned Lean kernel accepts from the sources this note links to, with no `sorry`, no added axiom, and no unchecked evaluation. That is a claim about the formal statement, not about its mathematical interest, its novelty, or the original problem. The unresolved obligations are named exactly, in their own section, and none of the finite computations, reductions, or no-go results here removes one of them.

The table records two independent facts. The *authority* column says whether the statement is supported by an argument in this paper together with a linked Lean decla-

ration, by a direct computation, or by no result in this paper. The last column states its logical reach and the nearest boundary. Thus a conditional implication may be checked in Lean while its hypothesis remains open, and the dyadic-window scan is a direct computation over 106,666 instances rather than an unbounded theorem.

Table 1: Authority and logical reach of the statements discussed in this note.

Statement	Authority	Logical reach and nearest boundary
A. Open target and prior literature		
Irrationality in any three-prime case	**No result here**	Open.
Irrationality after deleting duplicate running-LCM values	**Literature** (Erdős 1974 letter)	Asserted on p. 335, but no proof is supplied there and the letter's argument is not recovered here.
Transcendence after deleting duplicate values when $\lvert P\rvert = 2$	**Paper + external theorem** (Bugeaud–Laurent Theorem 1.1)	Unconditional for every pair of distinct primes; Theorem 4.1.
Transcendence of the original repeated sum when $\lvert P\rvert = 2$	**Paper + external theorem** (Bugeaud–Laurent Theorem 1.1)	Unconditional for every pair of distinct primes; Theorem 4.2. Neither two-prime theorem applies to three primes.
B. Unconditional finite structure		
$\mathrm{L}(x) = \mathrm{H}(x)$	**Paper + Lean**	Unconditional for distinct primes; Theorem 2.1. The unrestricted identity is classical.
$\mathrm{H}(x) \le x^3$	**Paper + Lean**	Unconditional; Proposition 2.2.
Constancy on logarithmic cells	**Paper + Lean**	Unconditional; Theorem 3.1.
Single-coordinate jump ratios	**Paper + Lean**	Unconditional; Theorem 3.2.
Exactly $3n + 1$ jump points	**Paper + Lean**	Unconditional; Theorem 3.3.
Height-fibre normal form	**Paper + Lean**	Exact finite identity; Theorem 5.1.
Infinite limit of that expansion	**No result here**	Not proved here; Section 5.
Shell multiplicity $9\,\#\delta \le (j+3)^2$	**Paper + Lean**	Unconditional; Theorem 6.3.
2×2 kernel restriction has rank two at $\{2,3,5\}$	**Paper + Lean**	Unconditional local obstruction; determinant $-1/15$ in Theorem 7.1.
Dyadic block radix lies in $\{2,6,10,30\}$	**Paper + Lean**	Unconditional; Theorem 8.1.
Canonical least-positive-residue arithmetic	**Paper + Lean**	Unconditional finite lemma; Theorem 8.3.
C. Conditional bridge		
Carry contradiction assuming (E)	**Paper + Lean**	Conditional implication only; Theorem 8.4.
Smooth-factor absorption and carry cancellation	**Paper + Lean**	Conditional core; Theorem 8.2. Divisibility of the actual carry remains unproved.
D. Bounded computation and remaining frontier		
Dyadic-window scan for $B \le 1000$	**Direct computation**	106,666 pairs with no failures; bounded evidence, not a theorem.
Bridge from the actual summands to (b_n, m_n, K)	**No result here**	Not proved here; Section 9.

Continued on the next page.

Statement	Authority	Logical reach and nearest boundary
Local-window escape for the actual $\{2,3,5\}$ word	No result here	Open; Problem 9.3.

Structure

Section 2 identifies the running value. Section 3 develops its jump structure, and Section 4 uses that structure to prove transcendence of both two-prime sums. Section 5 gives the finite three-prime normal form, Section 6 bounds fibre multiplicity, and Section 7 records the obstruction to separating the kernel. Section 8 determines the block alphabet, proves the conditional contradiction, and states the remaining unproved hypothesis. Section 9 collects the questions that remain. Linked phrases open the corresponding Lean declaration at the pinned source revision 76b5b0a7ed5d; the two transcendence theorems are analytic arguments using the cited external theorem and are not Lean-formalised here.

Keywords. irrationality; least common multiple; smooth numbers; lattice sums; Lean 4.
MSC 2020. 11J72 (primary); 11A05, 11N25, 68V20 (secondary).

2 The running least common multiple as a product of pure powers

The smooth numbers up to x are indexed by the exponent triples (i,j,k) with $i \le \lfloor \log_p x \rfloor$, $j \le \lfloor \log_q x \rfloor$, $k \le \lfloor \log_r x \rfloor$ and $p^i q^j r^k \le x$: the coordinate bounds make the index set finite, and the last condition is the actual constraint. This is the smooth prefix index set. Both conditions matter: the coordinate box is strictly larger than the prefix, since a product of three large pure powers can exceed x while each factor does not.

The identification below is routine, and its unrestricted analogue is classical, as recalled in the introduction; the short proof is given because every later statement is derived from it in the formal development.

Theorem 2.1 (the running least common multiple)**.** *Let p,q,r be pairwise distinct primes and $x \ge 1$. Then $\mathrm{L}(x) = \mathrm{H}(x)$.*

Proof. For divisibility in one direction, every smooth $n \le x$ has exponents bounded by the corresponding integer logarithms, so $n \mid \mathrm{H}(x)$, and hence $\mathrm{L}(x) \mid \mathrm{H}(x)$. For the other, the three pure powers $p^{\lfloor \log_p x \rfloor}$, $q^{\lfloor \log_q x \rfloor}$ and $r^{\lfloor \log_r x \rfloor}$ are themselves smooth numbers not exceeding x, so each divides $\mathrm{L}(x)$. Distinct primes have coprime powers, so their product divides $\mathrm{L}(x)$ as well. The two divisibilities give equality. □

Formalised as the running-lcm identity, from the divisibility into the height and the three membership statements for the pure powers, namely the first, second, and third pure-power memberships.

The hypothesis that the primes are pairwise distinct is used exactly once, in the coprimality step, and it is not decorative: without it the three pure powers need not have coprime orders and their product need not divide the least common multiple. The identity is what makes every later statement about L computable from three integer logarithms.

At $(p,q,r) = (2,3,5)$ it reads $\mathrm{L}(x) = 2^{\lfloor \log_2 x \rfloor} 3^{\lfloor \log_3 x \rfloor} 5^{\lfloor \log_5 x \rfloor}$, whose first ten values are

x	1	2	3	4	5	6	7	8	9	10
$\mathrm{L}(x)$	1	2	6	12	60	60	60	120	360	360

So $\mathrm{L}(10) = 8 \cdot 9 \cdot 5 = 360$, which is indeed the least common multiple of the smooth numbers $1,2,3,4,5,6,8,9,10$, and $\mathrm{L}(6) = 4 \cdot 3 \cdot 5 = 60$ rather than 6: the running value at a smooth cutoff already contains powers of the other two primes that the cutoff itself does not. These ten values also illustrate two of the statements of Section 3, since the value is constant on $\{5,6,7\}$ and on $\{9,10\}$, and each change multiplies by a single prime, by 2 at $x = 2,4,8$, by 3 at $x = 3,9$, and by 5 at $x = 5$.

Proposition 2.2. *For every $x \geq 1$, $\mathrm{H}(x) \leq x^3$.*

Proof. Each of the three factors is a power of its base not exceeding x. □

Formalised as the cubic majorant, which needs no primality. The exponent 3 is the number of generating primes, and the bound is the reason the reciprocal kernel is comparable to x^{-3} rather than to x^{-1}.

3 Constancy on logarithmic cells and the jump points

By Theorem 2.1 the running value depends on x only through the three integer logarithms. It is therefore constant wherever none of them changes and moves only where one of them does, and the next definition names the sets on which they are all constant, so that the next two theorems can say where the value stands still and by what factor it moves.

Say that x and y lie in the same *logarithmic cell* when $\lfloor \log_b x \rfloor = \lfloor \log_b y \rfloor$ for each of $b = p,q,r$: the cell relation.

Theorem 3.1 (constancy on cells)**.** *If $x,y \geq 1$ lie in the same logarithmic cell, then $\mathrm{L}(x) = \mathrm{L}(y)$. The same holds for the kernel at two smooth points whose values lie in one cell.*

Proof. Immediate from Theorem 2.1, since H depends on x only through the three integer logarithms. □

Formalised as the cell constancy of the running value, the cell constancy of the height, and the cell constancy of the kernel. The running value therefore changes only when one of the three logarithms changes, and the next theorem says by exactly how much.

Theorem 3.2 (single-coordinate jump ratios)**.** *Let $x,y \geq 1$. If $\lfloor \log_p y \rfloor = \lfloor \log_p x \rfloor + 1$ while the other two logarithms agree, then $\mathrm{L}(y) = p\,\mathrm{L}(x)$; and similarly with q or r in place of p.*

Proof. By Theorem 2.1 both sides are heights, and one factor of the height gains one in its exponent while the others are unchanged. □

Formalised as the first, second, and third coordinate steps, over the corresponding statements for the height alone, which need no primality: the first height step, the second, and the third.

The jump points are therefore the pure prime powers, and they do not collide across channels.

Theorem 3.3 (jump count)**.** *Let $n \geq 0$. The set of the first n positive powers of p, of q, and of r has exactly $3n$ elements, and adjoining the common origin 1 gives exactly $3n + 1$.*

Proof. The count is routine. Within one channel the powers $b, b^2, \dots, b^n$ are distinct because $b \geq 2$. Across two channels a common value would be a positive power of two distinct primes, which unique factorisation forbids. Finally 1 is not a positive power of any prime. □

Formalised as the positive jump count and the jump count with the origin, over the channel cardinality, the channel disjointness, and the exclusion of the origin; the channels themselves are the positive power sets. The exponent 0 is omitted from each channel because it is the shared initial value, which is why the origin is counted once rather than three times.

Two channels never meet, so at each positive pure power exactly one of the three logarithms advances, and by Theorem 3.2 the running value is multiplied there by the prime of that channel. Reading those multipliers in increasing order of the pure powers gives the *jump word* of L: one letter from $\{p, q, r\}$ for each positive pure power, recording which prime the value is multiplied by at that point. At $\{2, 3, 5\}$ the pure powers in increasing order are $2, 3, 4, 5, 8, 9, 16, 25, 27, 32, \dots$, so the jump word begins

$$2,\ 3,2,\ 5,2,\ 3,2,\ 5,3,2,\ \dots$$

where the grouping is the one used in Section 8: each group ends at a power of two.

4 The two-prime sums are transcendental

The jump description gives a complete result in rank two. Temporarily let $P = \{p, q\}$ with $p < q$, and write

$$L_{p,q}(t) = p^{\lfloor \log_p t \rfloor} q^{\lfloor \log_q t \rfloor}.$$

The argument of Theorem 2.1, with one coordinate omitted, identifies this as the running least common multiple of the $\{p, q\}$-smooth numbers up to t. Let

$$\mathcal{D}_{p,q} = 1 + \sum_{t \in \{p, p^2, \dots\} \cup \{q, q^2, \dots\}} \frac{1}{L_{p,q}(t)}.$$

Thus $\mathcal{D}_{p,q}$ retains the initial value and exactly one reciprocal for every later distinct running-LCM value.

Theorem 4.1 (two-prime de-duplicated transcendence)**.** *For every pair of distinct primes p, q, the number $\mathcal{D}_{p,q}$ is transcendental.*

Proof. Interchanging the primes if necessary, assume $p < q$, and set

$$\theta = \frac{\log p}{\log q}, \qquad x = \frac{1}{p}, \qquad y = \frac{1}{q}, \qquad m_n = \lfloor n\theta \rfloor, \qquad \delta_n = m_{n+1} - m_n.$$

Here $0 < \theta < 1$, and θ is irrational: a rational relation would give $p^b = q^a$ for positive integers a, b. Consequently $\delta_n \in \{0, 1\}$.

The initial value together with the p-channel contributes

$$A = \sum_{n\geq 0} x^n y^{m_n}.$$

There is a q-power strictly between p^n and p^{n+1} exactly when $\delta_n = 1$; it is then q^{m_n+1}, and its post-jump reciprocal is $x^n y^{m_n+1}$. Hence the q-channel contributes

$$B = \sum_{n\geq 0} \delta_n x^n y^{m_n+1}, \qquad \mathcal{D}_{p,q} = A + B.$$

All series here converge absolutely. Since $y^{m_{n+1}} - y^{m_n} = \delta_n y^{m_n}(y-1)$, shifting the sum for A gives

$$A - 1 - xA = x(y-1)\sum_{n\geq 0} \delta_n x^n y^{m_n}.$$

It follows that

$$\mathcal{D}_{p,q} = \frac{y-x}{x(y-1)}A - \frac{y}{x(y-1)}. \tag{1}$$

In the notation

$$F_\theta(x,y) = \sum_{n\geq 1}\sum_{k=1}^{\lfloor n\theta\rfloor} x^n y^k$$

for the Hecke–Mahler series, a finite geometric sum gives

$$A = \frac{1}{1-x} - \frac{1-y}{y}F_\theta(x,y). \tag{2}$$

Bugeaud and Laurent's Theorem 1.1 states, in particular, that $F_\theta(\beta,\alpha)$ is transcendental when $\theta \in (0,1)$ is irrational, α, β are nonzero algebraic numbers, $|\beta\alpha^\theta| < 1$, and $|\beta| < 1$ [9, p. 61, Theorem 1.1]; the $\rho = 0$ case used here goes back to Loxton and van der Poorten [10, Theorem 8]. We may take $(\beta,\alpha) = (x,y)$, because

$$|xy^\theta| = \frac{1}{p}\left(\frac{1}{q}\right)^{\log p/\log q} = \frac{1}{p^2} < 1.$$

Thus $F_\theta(x,y)$ is transcendental, and (2) makes A transcendental. Finally the coefficient of A in (1) is a nonzero algebraic number, since $x \neq y$; therefore $\mathcal{D}_{p,q}$ is transcendental. □

The same value also controls the series before repeated running-LCM values are removed. Because each $\{p,q\}$-smooth number is uniquely $p^i q^j$, that original series is

$$\mathcal{R}_{p,q} = \sum_{i,j\geq 0} \frac{1}{L_{p,q}(p^i q^j)}.$$

Theorem 4.2 (two-prime repeated-sum transcendence)**.** *For every pair of distinct primes p, q, the number $\mathcal{R}_{p,q}$ is transcendental.*

Proof. Again assume $p < q$ and use $\theta, x, y, m_n, \delta_n, A, B$ from the preceding proof. At the smooth point $p^i q^j$,

$$L_{p,q}(p^i q^j) = p^{i+\lfloor j/\theta \rfloor} q^{j+m_i},$$

because $\log_p(p^i q^j) = i + j/\theta$ and $\log_q(p^i q^j) = j + i\theta$. Absolute convergence therefore permits the factorisation

$$\mathcal{R}_{p,q} = AC, \qquad C = \sum_{j \geq 0} y^j x^{\lfloor j/\theta \rfloor}.$$

For each $j \geq 1$, put $n = \lfloor j/\theta \rfloor$. Irrationality of θ and $0 < \theta < 1$ give $m_n = j - 1$, $m_{n+1} = j$, and $\delta_n = 1$; conversely, every n with $\delta_n = 1$ arises from the unique $j = m_{n+1}$. Hence

$$C = 1 + B.$$

Equation (1), together with $\mathcal{D}_{p,q} = A + B$, gives

$$B = \frac{y(1-x)}{x(y-1)} A - \frac{y}{x(y-1)}.$$

Consequently

$$\mathcal{R}_{p,q} = \frac{y(1-x)}{x(y-1)} A^2 + \left(1 - \frac{y}{x(y-1)}\right) A.$$

The quadratic coefficient is a nonzero algebraic number. If $\mathcal{R}_{p,q}$ were algebraic, this display would make the transcendental number A a root of a nonzero polynomial over the algebraic numbers, a contradiction. □

Thus both versions are settled, at the stronger level of transcendence, for $|P| = 2$. A third prime replaces the single Beatty boundary by a genuinely two-dimensional ordering problem, so neither theorem supplies a three-prime irrationality result.

5 The height-fibre normal form

Grouping a lattice sum by the value of the running least common multiple turns it into a sum over heights whose coefficients are multiplicities. We prove this exactly, on a finite rectangular box.

Write $\mathcal{B}(h_p, h_q, h_r)$ for the box of exponent triples with $i \leq h_p$, $j \leq h_q$, $k \leq h_r$, the exponent box, and, for a height H, write $F(H)$ for the set of points of the box whose running height $\mathrm{H}(p^i q^j r^k)$ equals H, the height fibre. Distinct lattice points can carry the same height, and the coefficients of the normal form are exactly the sizes of these fibres.

Theorem 5.1 (finite normal form)**.** *For every box $\mathcal{B} = \mathcal{B}(h_p, h_q, h_r)$,*

$$\sum_{(i,j,k) \in \mathcal{B}} \mathrm{K}(i,j,k) = \sum_{H} \frac{\#F(H)}{H},$$

the outer sum ranging over the heights actually attained on $\mathcal{B}$.

Proof. The identity is a regrouping. Partition $\mathcal{B}$ into the fibres of the height map. On $F(H)$ every summand is $1/H$ by definition of the kernel, so the fibre contributes $\#F(H)/H$. □

Formalised as the height-fibre normal form, over the fibre sum, with the height of a lattice point given by the point height.

For a worked instance take $(p,q,r) = (2,3,5)$ and the box $\mathcal{B}(1,1,1)$, whose eight points carry the smooth values $1,2,3,5,6,10,15,30$ and the heights

$p^iq^jr^k$	1	2	3	5	6	10	15	30
H	1	2	6	60	60	360	360	10800

Six heights occur, two of them twice: the points 5 and 6 share the height 60, and 10 and 15 share the height 360. Theorem 5.1 here reads

$$1+\tfrac{1}{2}+\tfrac{1}{6}+\tfrac{1}{60}+\tfrac{1}{60}+\tfrac{1}{360}+\tfrac{1}{360}+\tfrac{1}{10800} = 1+\tfrac{1}{2}+\tfrac{1}{6}+\tfrac{2}{60}+\tfrac{2}{360}+\tfrac{1}{10800} = \tfrac{18421}{10800}.$$

The two coefficients 2 carry the whole content of the regrouping on this box; where the heights are pairwise distinct the identity is a relabelling and nothing more.

Together with Theorem 3.1 this is the finite core of the ordered prime-power jump expansion: the value is constant on cells, the cells are indexed by heights, and the coefficient of a height is the number of lattice points it collects. The passage to the infinite sum, and the explicit ordering of the pure powers that would make the expansion a series in the jumps, are not proved here. Theorem 5.1 is an identity between two finite sums, and it is stated over the full box rather than the smooth prefix, so it is not a statement about L at a cutoff.

The same module records a one-step map $\tau(b,d,s) = b(s-d)$, the variable-base tail step, with the rewriting that names its expanded form. No orbit of this map is analysed.

6 A quadratic bound for smooth exponent shells

A tail estimate needs to know how many lattice points can share a short multiplicative interval. Fix a box $\mathcal{B}(h_p,h_q,h_r)$ and an interval $[\lambda,\eta)$, and write δ for the set of exponent triples of that box whose smooth value lies in that interval, the smooth exponent shell. Every bound below assumes the interval short in the multiplicative sense, meaning that η is at most a stated multiple of λ.

Informally, the next lemma says that an interval whose right endpoint is at most b times its left endpoint contains at most one of the numbers b^aw.

Lemma 6.1 (uniqueness in a short interval)**.** *Let $b \geq 1$ and suppose $\eta \leq b\lambda$. If b^aw and $b^{a'}w$ both lie in $[\lambda,\eta)$, then $a = a'$.*

Proof. If $a < a'$ then $\eta \leq b\lambda \leq b \cdot b^aw = b^{a+1}w \leq b^{a'}w < \eta$, which is impossible; the case $a > a'$ is symmetric. □

Formalised as the short-interval uniqueness. The hypothesis is that the interval has multiplicative width at most b.

Proposition 6.2. *Let $\mathcal{S}$ be the shell of the box $\mathcal{B}(h_p, h_q, h_r)$ in $[\lambda, \eta)$. If $\eta \le r\lambda$ then $\#\mathcal{S} \le (h_p+1)(h_q+1)$; if $\eta \le p\lambda$ then $\#\mathcal{S} \le (h_q+1)(h_r+1)$.*

Proof. Suppose $\eta \le r\lambda$. If two triples of $\mathcal{S}$ agree in their first two coordinates, Lemma 6.1 applied with $b = r$ and $w = p^i q^j$ forces their third coordinates to agree as well. So the projection forgetting the third coordinate is injective on $\mathcal{S}$, and its image lies in a rectangle with $(h_p+1)(h_q+1)$ points. The case $\eta \le p\lambda$ is the same with the first coordinate projected away. □

Formalised as the third-coordinate projection and the first-coordinate projection.

Theorem 6.3 (quadratic multiplicity bound)**.** *Let $\mathcal{S}$ be the shell of the box $\mathcal{B}(h_p, h_q, h_r)$ in $[\lambda, \eta)$. Suppose $\eta \le r\lambda$ and $h_p \le h_q \le h_r$ with $h_p + h_q + h_r = j$. Then*

$$9\,\#\mathcal{S} \le (j+3)^2.$$

Proof. By Proposition 6.2, $\#\mathcal{S} \le (h_p+1)(h_q+1)$; under the sorting hypothesis the two surviving coordinates are the two smallest. It therefore suffices to prove that $a \le b \le c$ with $a+b+c = j$ gives $9(a+1)(b+1) \le (j+3)^2$. From $a \le b \le c$ we get $a + 2b \le j$, so it is enough that $9(a+1)(b+1) \le (a+2b+3)^2$; writing $b = a + d$ with $d \ge 0$, the difference of the two sides is $d(3a + 4d + 3) \ge 0$. □

Formalised as the quadratic shell bound, over the sorted quadratic estimate. The constant 9 is the square of the number of generating primes and appears because the bound is the arithmetic–geometric comparison for a sum of three sorted coordinates; sorting is a hypothesis, not a normalisation, since the shell itself is not symmetric in the three bases. The last inequality of the proof is an equality when $h_p = h_q = h_r$, both sides then being $9(h_p+1)^2$, so no constant larger than 9 survives that step.

The bound is uniform in λ and η subject to the width condition, and it is stated for the actual filtered shell rather than for a lattice model of it. It is an input to a tail estimate and is not itself one: no series is bounded here. The estimate is elementary and uses no analytic input on the distribution of smooth numbers, only the projection of Proposition 6.2.

7 Non-separability of the three-prime kernel

One might hope to write the three-prime kernel as $f(i)g(j)h(k)$ and so reduce the problem to one-dimensional criteria. Such a factorisation would force the value at $(1,1,0)$ to be determined by the values at $(0,0,0)$, $(1,0,0)$ and $(0,1,0)$, since the four values would then satisfy $\mathrm{K}(0,0,0)\ \mathrm{K}(1,1,0) = \mathrm{K}(1,0,0)\ \mathrm{K}(0,1,0)$. Theorem 7.1 computes those four values at $\{2,3,5\}$ and finds that they do not: the failure is exact, and it occurs on the smallest rectangle on which it could occur.

Theorem 7.1 (non-separability at $\{2,3,5\}$)**.** *With $(p,q,r) = (2,3,5)$,*

$$\mathrm{K}(0,0,0)\ \mathrm{K}(1,1,0) - \mathrm{K}(1,0,0)\ \mathrm{K}(0,1,0) = -\tfrac{1}{15} \neq 0.$$

Consequently this two-by-two kernel restriction has rank two over $\mathbb{Q}$.

Proof. The four values are $K(0,0,0) = 1$, $K(1,0,0) = 1/2$, $K(0,1,0) = 1/6$ and $K(1,1,0) = 1/60$, computed from $H(1) = 1$, $H(2) = 2$, $H(3) = 2 \cdot 3 = 6$ and $H(6) = 4 \cdot 3 \cdot 5 = 60$. Hence the determinant is $1/60 - 1/12 = -1/15$, and a two-by-two matrix with nonzero determinant has rank two. □

Formalised as the non-separation witness, over the four exact values, the origin, value at two, value at three, and value at six. The exact determinant calculation is the rank-two certificate. The height at 6 is 60 rather than 6 because the maximal pure powers below 6 are 4, 3 and 5: the running least common multiple at a smooth cutoff sees powers of the other primes that the cutoff itself does not contain. That is the mechanism behind the non-separation.

A nonzero two-by-two minor rules out writing the kernel as $f(i)g(j)h(k)$ on this box, so no argument may assume that single product form here. It rules out nothing further: a rank-two matrix is itself a sum of two rank-one matrices, so sums of separable terms, higher-rank decompositions, separable majorants, changes of variable, and one-dimensional estimates applied after a decomposition all remain available. It gives no rank lower bound beyond two, and it is not an independence or irrationality statement.

8 Dyadic blocks and a conditional carry contradiction

The section makes three moves and leaves one hypothesis standing. We first compress the jump word of Section 3 into the blocks cut out by consecutive powers of two, and show that the multiplier of a block takes only four values (Theorem 8.1). We then record what can be cancelled from a hypothetical denominator, and exactly where that cancellation is still conditional (Theorem 8.2). Finally we leave the smooth numbers behind and argue with integer sequences alone: for a multiplier coprime to 30, no positive sequence obeying the cleared recurrence can stay inside its bound once a certain residue condition holds arbitrarily far out (Theorem 8.4). That residue condition, condition (E) below, is the hypothesis; it is not proved here, and the closing subsection reports a finite computation, which is evidence for it and not a proof of it.

8.1 The four-element block alphabet

Take $(p,q,r) = (2,3,5)$ and compress the jump word of Section 3 between consecutive powers of two: a block starts just after 2^a, includes every pure 3- or 5-power strictly between 2^a and 2^{a+1}, and ends with the jump at 2^{a+1}. A channel cannot occur twice inside one block. Indeed, if

$$2^a < b^e, b^f < 2^{a+1} \qquad (b \geq 2),$$

then the ratio between the interval endpoints is $2 \leq b$, so strict monotonicity of the powers forces $e = f$. This is the checked internal-power uniqueness lemma. Each block therefore contributes at most one letter from each of the 3- and 5-channels, and exactly one letter 2 at its right end.

Let β_a, the *block radix*, be the product of the terminal dyadic factor 2, a factor 3 when the block contains an internal 3-power, and a factor 5 when it contains an internal 5-power; equivalently, β_a is the product of the letters of the jump word lying in block a.

Theorem 8.1 (the dyadic block alphabet)**.** *With* $(p,q,r) = (2,3,5)$ *and* β_a *as above, for every* a*,*

$$\beta_a \in \{2,6,10,30\}; \qquad \textit{in particular} \qquad 2 \le \beta_a \le 30.$$

Proof. Internal-power uniqueness leaves two independent yes/no choices, one for the 3-channel and one for the 5-channel. Multiplying the terminal factor 2 by the selected channel factors gives exactly the displayed four cases. □

All four letters already occur among the first five of the six blocks tabulated below:

a	$(2^a, 2^{a+1})$	internal 3-power	internal 5-power	β_a
0	$(1,2)$	none	none	2
1	$(2,4)$	3	none	6
2	$(4,8)$	none	5	10
3	$(8,16)$	9	none	6
4	$(16,32)$	27	25	30
5	$(32,64)$	none	none	2

Block 4 is the only one of these six carrying an internal power in both channels, and block 5 carries neither, so its radix falls back to the terminal factor alone. Multiplying the radices along a run of blocks gives the product of the jump-word letters over that run: for instance $\beta_1\beta_2 = 60$ is the product of the four multipliers at $3,4,5,8$.

The definition is the dyadic block base; Lean checks both the exact four-case alphabet and the bounded-radix consequence. The radix word is therefore constrained to four values, and no growth hypothesis on it is needed. Section 9 gives a literal finite formula for the corresponding block digit m_a^{235}, matching the integer-only checker. What is not yet checked in Lean is the theorem that identifies that digit, its tail and its sharp carry bound with the original repeated series under a rationality hypothesis.

8.2 The denominator reduction and its boundary

Let D be the reduced denominator of a hypothetical rational value, and write

$$D = D_{\mathrm{sm}}B, \qquad D_{\mathrm{sm}} = 2^u 3^v 5^w, \qquad \gcd(B,30) = 1.$$

The coprimality condition below is therefore intended as the endpoint of a reduction from an arbitrary D, not as a restriction on which rational values are being considered.

The argument runs on a sequence of integers c_n, one for each step, called the *carry states*; they satisfy a recurrence $c_{n+1} = b_n c_n - D m_n$ of the shape displayed below, driven by a radix word b_n and a forcing word m_n. The name is meant to suggest the integer left after clearing D from the n-th tail of the series. Supplying that reading for the actual series, and with it the divisibility used in the next paragraph, is the unproved identification of Section 9; nothing in this subsection or the next depends on the reading, only on the recurrence.

Every fixed $\{2,3,5\}$-smooth factor divides the running height once the cutoff reaches that factor (checked absorption). If the denominator-cleared carry states c_n share the

absorbed factor, so that $c_n = D_{\mathrm{sm}} d_n$, Lean cancels it from

$$c_{n+1} = b_n c_n - D_{\mathrm{sm}} B\, m_n$$

and obtains

$$d_{n+1} = b_n d_n - B m_n$$

(checked cancellation). Positivity and the sharp denominator-dependent upper bound descend through the same positive factor (checked bound transfer), and the reduced carry inherits the exact window identity (checked window transfer).

Informally: provided the smooth part D_{sm} divides every carry state, it can be divided out of the whole system, leaving the same four statements with a multiplier coprime to 30 in place of D. That proviso is the hypothesis of the theorem, and it is not proved here.

Theorem 8.2 (conditional denominator reduction)**.** *If the actual denominator-cleared carry for $D = D_{\mathrm{sm}} B$ has the common-factor form $c_n = D_{\mathrm{sm}} d_n$ with $D_{\mathrm{sm}} > 0$, then its recurrence, positivity, bound, and window identity reduce to the same statements with multiplier B, where* $\gcd(B, 30) = 1$.

The theorem is conditional at exactly one point: height absorption does not by itself prove that the *carry state* is divisible by D_{sm}. That divisibility must come from the still-unproved rationality-to-carry identification for the actual series. The phrase "B coprime to 30" below is valid only downstream of this bridge.

8.3 The window recurrence and the residue contradiction

The statements of this subsection are about integer sequences: they do not refer to L or to the smooth numbers, and the identification of those sequences with the block data of the preceding subsection is the obligation just recorded.

Call a pair (ℓ, h) with $h > 0$ a *window*: the stretch of h consecutive steps beginning at index ℓ. To locate a carry at the end of a window without performing any division, one needs the multiplier and the forcing accumulated across the window, and those are the two sequences defined next. For an integer radix word b_n and forcing word m_n, define

$$\begin{aligned} W_{\ell,0} &= 1, \quad W_{\ell,h+1} = b_{\ell+h} W_{\ell,h}, \\ F_{\ell,0} &= 0, \quad F_{\ell,h+1} = b_{\ell+h} F_{\ell,h} + m_{\ell+h}. \end{aligned}$$

These are the window base and window forcing. If an integral carry satisfies

$$d_{n+1} = b_n d_n - B m_n,$$

then induction gives the exact division-free identity

$$d_{\ell+h} = W_{\ell,h} d_\ell - B F_{\ell,h};$$

this is the checked window identity.

For $C > 0$, let $\mathrm{lpr}_C(x) \in \{1, \dots, C\}$ be the least positive representative of x mod C, with a zero residue represented by C. This is the canonical representative. Lean checks

both its positive range and its congruence to the source integer. This convention matters: replacing $\mathrm{lpr}_C(x)$ by $|x|$ would not be a modular statement.

Let $K(B,n)$ be a bound on the reduced carry at denominator B and step n. It enters as a parameter: the statements below hold for whichever function K is supplied, and deriving the correct one for the actual series is part of the identification of Section 9. Define *cofinal local-window escape* to mean that for every $B > 0$ coprime to 30 and every ℓ_0, there are $\ell \geq \ell_0$ and $h > 0$ such that

$$C_{\ell,h} := |W_{\ell,h}| > 0 \quad \text{and} \quad K(B, \ell + h) < \mathrm{lpr}_{C_{\ell,h}}(-BF_{\ell,h}). \tag{E}$$

Both the quantifier over B and the dependence of K on B are part of the statement. Informally, and separately for each fixed $B > 0$ coprime to 30: however far out one starts, some window has a nonzero accumulated base, and the least positive residue of its accumulated forcing, weighted by $-B$ and taken modulo that base, exceeds the carry bound at the endpoint of the window. The exact unproved proposition is the cofinal local-window escape condition.

The next statement is the finite core of the argument. Informally, it says that a positive integer of size at most K cannot be congruent modulo C to a number whose canonical positive residue exceeds K. The key point is the convention just fixed: because lpr_C represents a vanishing residue by C and not by 0, the proof has to treat that case separately, and the two branches conclude for different reasons.

Theorem 8.3 (the finite residue contradiction)**.** *Let $C > 0$, $c > 0$, and $|c| \leq K$. If $c \equiv x \pmod{C}$ and $K < \mathrm{lpr}_C(x)$, then the hypotheses are contradictory.*

Proof. The canonical representative lies in $\{1, \dots, C\}$. The inequalities put c strictly between 0 and C. If $x \equiv 0 \pmod{C}$, its positive representative is C while $c \bmod C = c \neq 0$. Otherwise both c and $\mathrm{lpr}_C(x)$ are their own residues and congruence makes them equal, contradicting $|c| \leq K < \mathrm{lpr}_C(x)$. □

The natural-state version is the finite contradiction; the integer carry version is the integer form of the contradiction.

The same finite arithmetic gives an exact classifier, not merely a contradiction. If $C > 0$, $c > 0$, $|c| \leq C$, and $c \equiv x \pmod{C}$, then

$$\mathrm{lpr}_C(x) = |c|.$$

This includes the endpoint correctly: the zero congruence class is represented by C, so $c = C$ gives $\mathrm{lpr}_C(x) = C$, not zero. Lean checks this as exact least-positive-residue classifier. This is only a finite theorem. Applying it to the series still requires two independent results: rationality must produce a positive bounded integral carry with the required smooth-divisibility congruence, and one must construct arbitrarily late escaping windows.

Theorem 8.4 (a conditional contradiction for bounded carries)**.** *Let (b_n), (m_n) and K be the radix word, the forcing word and the bound of this subsection, and assume the cofinal escape property (E). Fix $B > 0$ coprime to 30. There is no integral sequence d_n satisfying simultaneously*

$$d_{n+1} = b_n d_n - B m_n, \qquad d_n > 0, \qquad |d_n| \leq K(B,n) \quad (n \geq 0).$$

Proof. Choose one window supplied by (E). The checked window identity gives

$$d_{\ell+h} \equiv -BF_{\ell,h} \pmod{|W_{\ell,h}|}.$$

The endpoint state is positive and at most $K(B, \ell + h)$, whereas the canonical positive residue of the right-hand side is larger than this bound. Theorem 8.3 is the contradiction. □

This is formalised as the reduced-carry extinction theorem. Coprimality with 30 is used by (E) to select a window; once a window has been chosen, the finite contradiction does not use it. The formalisation carries the edge cases: $|W_{\ell,h}| = 0$ is excluded, a zero residue is represented by the full modulus, and positivity prevents the endpoint carry from being zero.

8.4 A finite check of the escape condition

A *certificate* here is a finite tuple of integers recording one instance of the inequality in (E) for the data the checker constructs. It exists so that a reader can recheck the instance in a line, without running the checker and without touching the infinite part of the argument. The integer-only dyadic-window checker constructs the ordered pure-power jumps, the block bases, and the block digits from exact multiplicity counts. It reproduces the following certificates; the columns are denominator B, dyadic start a, window length h, endpoint jump index n, window base W, forcing F, least positive residue $R = \mathrm{lpr}_W(-BF)$, and short bound K.

B	a	h	n	W	F	R	K
1	1	2	4	60	47	13	9
7	1	3	6	360	289	137	95
16	1	4	9	10800	8735	640	352

The first row reads as follows. The window starts at $a = 1$ and has length 2, so its base is the product of the two block radices, $W = \beta_1\beta_2 = 6 \cdot 10 = 60$; the forcing accumulated over the window is $F = 47$; and $\mathrm{lpr}_{60}(-47) = 13$, since $-47+60 = 13$, which exceeds the bound $K = 9$. The other two rows are read the same way, with $W = \beta_1\beta_2\beta_3 = 360$ and $W = \beta_1\beta_2\beta_3\beta_4 = 10800$. The third row lies outside the domain of the escape condition, since $\gcd(16, 30) = 2$ while both (E) and Theorem 8.4 quantify only over $B > 0$ coprime to 30. It is displayed to illustrate the window arithmetic at greater depth, not as an instance of the escape condition.

A fresh scan over every $B \le 1000$ coprime to 30 and every $100 \le a \le 500$ tested 106,666 pairs. In every case a window of length at most 18 made both $W > K$ and $\mathrm{lpr}_W(-BF) > K$; the largest first successful length was 14. The computation uses integers only and is reproducible from the pinned checker. Neither the scan nor the three displayed certificates proves escape for unbounded B or for cofinally many starts.

9 Complements and further questions

Theorems 4.1 and 4.2 close both two-prime questions, at the stronger level of transcendence. The three-prime de-duplicated and repeated series are outside the one-

dimensional Hecke–Mahler reduction used in those theorems and remain open. At three primes the exact unresolved statement is best separated from the bridge and from the possible methods for proving it.

9.1 The actual block data and the missing bridge

The checker already uses literal data, which we record so that the open problems have no unspecified forcing word. For $p \in \{2,3,5\}$ let q, r be the other two primes and put

$$A_p(e) = \#\{(i,j) \in \mathbb{N}^2 : q^i r^j < p^e\}, \qquad C_p(e) = \sum_{u=1}^{e} A_p(u).$$

Let I_a be the increasing list of internal jumps (p,e) with $p \in \{3,5\}$ and $2^a < p^e < 2^{a+1}$. For $(p,e) \in I_a$ let

$$\sigma_a(p,e) = \prod_{(q,f)\in I_a,\ p^e<q^f} q.$$

Then, for $a \geq 1$, the exact checker definitions are

$$\beta_a = 2 \prod_{(p,e)\in I_a} p, \qquad m_a^{235} = A_2(a+1) + \sum_{(p,e)\in I_a} (p-1)\sigma_a(p,e)(C_p(e) - C_2(a)). \tag{9.1}$$

The first formula is the four-letter radix of Theorem 8.1; the second is the integer implemented by the pinned checker. If

$$\nu_a = \#\{p^e : p \in \{2,3,5\},\ e \geq 1,\ p^e < 2^{a+1}\},$$

its exact tested short bound is

$$K^{235}(B,a) = \left\lfloor \frac{B(\nu_a^2 + 10\nu_a + 27)}{9} \right\rfloor. \tag{9.2}$$

Finally define the actual scaled block tail

$$T_a = \sum_{j\geq a} \frac{m_j^{235}}{\beta_a \beta_{a+1} \cdots \beta_j}; \qquad T_{a+1} = \beta_a T_a - m_a^{235}. \tag{9.3}$$

Problem 9.1 (actual rationality-to-carry identification)**.** Starting from the original repeated $\{2,3,5\}$ series and a reduced denominator $D = D_{\mathrm{sm}}B$, prove with a stated onset $a \geq a_D$ that the literal objects in (9.1)–(9.3) give positive integral carries

$$c_a = DT_a, \qquad c_{a+1} = \beta_a c_a - Dm_a^{235}, \qquad 0 < c_a \leq K^{235}(D,a), \qquad D_{\mathrm{sm}} \mid c_a.$$

Writing $c_a = D_{\mathrm{sm}} d_a$, prove the reduced recurrence and its sharp bound

$$d_{a+1} = \beta_a d_a - Bm_a^{235}, \qquad 0 < d_a \leq K^{235}(B,a).$$

This is the missing connection, not a notational convenience: the generic carry theorems of Section 8 do not identify themselves with the original series. The problem fixes the digit, tail, onset and bound against which a proposed proof can be tested.

9.2 The intrinsic tail question

Problem 9.2 (exact nonintegrality of every reduced tail)**.** For every $B \geq 1$ with $\gcd(B, 30) = 1$ and every $a \geq 1$, prove

$$BT_a \notin \mathbb{Z}. \tag{9.4}$$

This pointwise form is stronger-looking but cleaner than "cofinally nonintegral": if BT_a is integral at one index, the recurrence

$$BT_{a+1} = \beta_a BT_a - Bm_a^{235}$$

makes it integral at every later index. Thus a direct solution of Problem 9.2, joined to Problem 9.1, bypasses all residue-window machinery.

The bounded-radix theorem gives a useful exact reduction. Since $2 \leq \beta_a \leq 30$, any real affine tail orbit either hits an integer or is, cofinally often, at distance at least $1/31$ from every integer (checked). It does not exclude the integral branch; Problem 9.2 is exactly what must do so for the actual orbit.

9.3 A denominator-adaptive sufficient criterion

For the literal pair (m^{235}, K^{235}), retain the window definitions of Section 8:

$$W_{\ell,h} = \prod_{j=0}^{h-1} \beta_{\ell+j}, \qquad F_{\ell,0} = 0, \qquad F_{\ell,h+1} = \beta_{\ell+h} F_{\ell,h} + m_{\ell+h}^{235}.$$

Problem 9.3 (actual cofinal local-window escape)**.** Prove the displayed quantifier order

$$\forall B \geq 1\ (\gcd(B, 30) = 1),\ \forall a_0 \geq 1,\ \exists \ell \geq a_0\ \exists h \geq 1 : \quad \operatorname{lpr}_{W_{\ell,h}}(-BF_{\ell,h}) > K^{235}(B, \ell + h). \tag{9.5}$$

Every β_a is positive, so $W_{\ell,h} > 0$ is automatic. The cofinal quantifier is present for a substantive reason: Problem 9.1 may supply the reduced recurrence only after the denominator-dependent onset a_D, and (9.5) then supplies a window beyond that onset. Once such a window is chosen, the positive endpoint carry must equal the canonical residue exactly (checked), yet it lies in the possible carry set $\{1, \dots, K^{235}(B, \ell + h)\}$; the strict inequality excludes that set. This is a one-sided least-positive-residue statement, not two symmetric arcs around zero.

Two exact countermodels rule out tempting shortcuts. For $(W, F, B) = (6, 4, 1)$,

$$\operatorname{lpr}_6(-4) = 2,$$

so the canonical residue need not be coprime to W. For $(W, F, B) = (60, 47, 37)$,

$$\operatorname{lpr}_{60}(-37 \cdot 47) = 1,$$

so a fixed window has no denominator-independent positive residue lower bound. Therefore the window may genuinely depend on B; neither a fixed finite computation nor a universal bounded-length guess addresses the unbounded denominator and cofinal-start quantifiers. The finite checker can reject a proposed sufficient condition, but supplies no evidence for either shortcut.

9.4 A literal two-dimensional analytic route

For the de-duplicated series define

$$\mathcal{D}_{2,3,5} = 1 + \sum_{t \in \{2^n, 3^n, 5^n : n \geq 1\}} \frac{1}{2^{\lfloor \log_2 t \rfloor} 3^{\lfloor \log_3 t \rfloor} 5^{\lfloor \log_5 t \rfloor}}.$$

The dyadic coding is the joint rotation word

$$\delta_{3,a} = \lfloor (a+1)\theta_3 \rfloor - \lfloor a\theta_3 \rfloor, \quad \theta_3 = \frac{\log 2}{\log 3}, \qquad \delta_{5,a} = \lfloor (a+1)\theta_5 \rfloor - \lfloor a\theta_5 \rfloor, \quad \theta_5 = \frac{\log 2}{\log 5},$$

with $\beta_a = 2\,3^{\delta_{3,a}} 5^{\delta_{5,a}}$.

Problem 9.4 (function-faithful two-dimensional representation)**.** Express $\mathcal{D}_{2,3,5}$ as a nonconstant algebraic combination of values of a specified two-dimensional Hecke–Mahler, cone-generating or multivariate Mahler function and verify every hypothesis of a published value theorem; or give a conditional theorem under an explicit logarithmic nondegeneracy hypothesis; or prove that the literal series has no representation in the specified finite-dimensional class.

Pairwise irrationality of θ_3 and θ_5 is not silently promoted to the orbit-closure or equidistribution hypothesis a two-dimensional theorem may need. The finite-observer formalisation isolates the precise faithfulness requirement: equality in a finite observer must imply equality after symbolic realisation, and a genuine finite-dimensional factorisation forces the realised symbolic span to be finite-dimensional (residue–coboundary form, system, checked). No theorem here proves that the literal realised span is infinite, so a general finite separable decomposition is neither assumed nor declared excluded.

9.5 The structural frontier

The radix alphabet itself extends without difficulty. For ordered primes $p_1 < \cdots < p_s$, an interval (p_1^a, p_1^{a+1}) contains at most one power from each other channel, because consecutive p_i-powers have ratio $p_i > p_1$. Hence its block radix belongs to the 2^{s-1}-letter alphabet

$$\left\{ p_1 \prod_{i=2}^{s} p_i^{\varepsilon_i} : \varepsilon_i \in \{0,1\} \right\}.$$

The nontrivial question is quantitative: prove effective recurrence or discrepancy for the actual four-letter $\{2, 6, 10, 30\}$ word, an asymptotic formula with an error term for the restricted two-dimensional shell counts that generate m_a^{235}, or the exact finite-separation rank of the literal three-prime kernel under a specified family of shifts.

The three-channel rigidity and carry-lift extinction theorems already exclude one false route in four exact steps. Under channel surjectivity, ordinary block-nullity is equivalent to the perturbation being a coboundary of a channel potential (potential classifier). Zero perturbations on genuine $2 \to 3$ and $2 \to 5$ transitions then identify all three potential values and force the perturbation to vanish at every index. For an integral lift, that vanishing makes a nonzero initial lift error grow by the exact product

of the successive bases; bases at least two make its absolute value at least 2^N, contradicting even a single index-N bound strictly below 2^N (one-index extinction). A single index therefore already suffices, and consequently no uniform bound on the lift error can hold either: under the same hypotheses a nonzero initial error is incompatible with any bound valid at every index (uniform extinction). A separate four-state calculation reaches the obstruction earlier: four real states in $(0, 1)$ with unit-accuracy integral lifts and the two anchor equalities force the first complete 2-block sum to be 1, hence that block cannot be null (first-block sum, four-state obstruction).

These conclusions remain conditional. No theorem constructs the actual #269 orbit, its ordered-power word, an integral carry lift, the two anchors, or block-nullity. The affine alternative above may produce an arbitrary integral state, not necessarily zero; its cofinal 1/31 separation is neither eventual nor positive-density and gives no unbounded distance. The four-state calculation supplies no cofinal windows, and `rational_of_scaledTail_integer` classifies no denominators. The actual carry supplies a weighted block defect instead, so any successful argument must use that weighted identity or construct a different faithful lift. None of these results is an irrationality theorem for #269.

Until the bridge and either the direct tail problem or an adequate substitute are proved, Problem #269 remains open.

Statements and declarations

Artefact and data availability. The pinned formal-source revision contains the Lean sources, the fixed toolchain, the library manifest, and the exact dyadic-window checker used in the finite experiment. This manuscript provides navigation rather than proof authority.

Declaration of generative AI use. Large-language-model agents were used throughout development to draft and revise prose, formal proofs, and software. The author set the objectives and acceptance criteria, selected and reviewed the claims, and approved the published version. The author assumes responsibility for the accuracy, interpretation, and presentation of the work. Generative systems are production tools; they are not authors and supply no independent authority. Lean checks each proof term against the fixed library version, and the sources linked here contain no proof placeholders and no project-defined axioms; Lean does not authorise the exposition, the citation choices, or the interpretation, for which the author remains responsible.

Funding and competing interests. This work received no external funding. The author declares no competing interests.

Acknowledgements. The problem numbering and status follow the Erdős Problems catalogue maintained by Thomas Bloom [4].

A Guide to the formal sources

Each linked phrase opens its Lean declaration at the pinned source revision 76b5b0a7ed5d. The running-LCM structure, residue arithmetic, and local-window

bridge occupy three separate modules. Four distinctions are worth carrying into the source. The height statements hold for arbitrary bases, while the statements about L need the three primes to be distinct. The normal form of Section 5 is a finite identity over a rectangular box, not a convergence theorem. The smooth part of a hypothetical denominator can be cancelled only after divisibility of the actual carry states by that factor is proved. And the formal cofinal-escape predicate is an unproved hypothesis of Theorem 8.4; its application to the actual $\{2, 3, 5\}$ word is not asserted.

References

[1] P. Erdős and R. L. Graham, *Old and New Problems and Results in Combinatorial Number Theory*, Monogr. Enseign. Math. 28, Geneva, 1980, p. 65. For a possibly infinite prime set Q, the page states the infinite-Q irrationality and asks what happens for finite Q with more than one element.

[2] P. Erdős, *On the irrationality of certain series: problems and results*, in A. Baker (ed.), *New Advances in Transcendence Theory*, Cambridge UP, 1988, pp. 102–109, doi:10.1017/CBO9780511897184.009.

[3] P. Erdős, *Letter to the Editor* (written 1 January 1973), Fibonacci Quart. **12** (1974), no. 4, p. 335. The letter poses the full series as a conjecture and asserts irrationality after retaining only the distinct running-LCM values.

[4] T. F. Bloom, *Erdős Problem #269*, `erdosproblems.com/269`, accessed 28 July 2026 (page displays "last edited 28 December 2025"). The current record labels the finite-support problem open, cites `[ErGr80, p. 65]` and `[Er88c, p. 106]`, routes to the 1974 letter on p. 335, records the infinite-prime and de-duplicated variants, and explicitly describes its status as the website owner's present assessment rather than a literature-completeness guarantee.

[5] The Formal Conjectures Authors, *FormalConjectures.ErdosProblems.269*, Lean source at commit `f776d2f`, 2025, accessed 28 July 2026.

[6] T. M. Apostol, *Introduction to Analytic Number Theory*, Springer, New York, 1976.

[7] H. L. Montgomery and R. C. Vaughan, *The Prime Number Theorem*, in *Multiplicative Number Theory I: Classical Theory*, Cambridge Studies in Advanced Mathematics 97, Cambridge University Press, 2007, pp. 168–198, doi:10.1017/CBO9780511618314.008.

[8] V. Kovač and T. Tao, *On several irrationality problems for Ahmes series*, Acta Math. Hungar. **175** (2025), 572–608, doi:10.1007/s10474-025-01528-0; arXiv:2406.17593, 2024.

[9] Y. Bugeaud and M. Laurent, *Transcendence and continued fraction expansion of values of Hecke–Mahler series*, Acta Arith. **209** (2023), 59–90, doi:10.4064/aa220323-18-1; authors' publisher-layout PDF, arXiv:2203.12901v1. Theorem 1.1 is on journal p. 61 (publisher-layout PDF p. 3); the authors note there that its $\rho = 0$ case was already obtained by Loxton and van der Poorten.

[10] J. H. Loxton and A. J. van der Poorten, *Arithmetic properties of certain functions in several variables III*, Bull. Austral. Math. Soc. **16** (1977), 15–47.

[11] S. Fan, comment on Erdős Problem #269, erdosproblems.com forum, thread 269, 26 June 2026. The comment gives the two-channel factorisation, the Hecke–Mahler reduction, and the transcendence conclusion for $|P| = 2$; follow-up comments there note the extension to arbitrary coprime pairs.