Reciprocal Mersenne Subseries

A claim-bounded reasoning surface for Erdős Problem #257

Will Cook

July 2026

Authorship and AI use. The prose in this version was generated 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*.

Abstract

For $n \in \mathbb{N}_{>0}$ let $w_n = (2^n - 1)^{-1}$, and for $J \subseteq \mathbb{N}_{>0}$ let $\mathcal{A}_J = \{\sum_{n \in J} \varepsilon_n w_n : \varepsilon_n \in \{0,1\}\}$. Erdős Problem 257 asks whether every infinite-support sum is irrational. This paper is a Lean-backed structural audit, not a solution. Its unconditional checked content includes exact finite-denominator noncollapse; unique coding, compactness, and nowhere density for every restricted achievement set $\mathcal{A}_J$; perfectness when J is infinite; and the measure dichotomy $2^{-|F|}$ when $J^c = F$ is finite, versus zero when J^c is infinite. For squarefree support the divisor incidence is $2^{\omega(n)} - 1$: at every even base this blocks two divisibility-first certificate schemas, while adjoining 1 preserves the irrationality question and restores arbitrarily long base-2 opening blocks. Adjacent suffix-cylinder words also admit an exact carry-pivot normal form. The targets $1/2$ and $1/21$ have no finite-support representation in the stated normalisation and reduce to explicit infinite-orbit alternatives; neither membership question is decided. Published work independently settles prime support at base 2 and squarefree support at every power-of-two base. We claim no priority for those cited results. Universal #257 and both target membership questions remain open.

How to read this document

This is a reasoning surface, not an exhaustive corpus export or a literature survey. It exposes the principal premises, method boundaries, and open interfaces; the machine-readable corpus remains the current inventory when later theorem waves outrun the exposition. Three conventions carry the paper's claim boundary.

Evidence bands. Every statement is tagged. [Lean] means a proof term was checked by the pinned Lean kernel. [Cert] means an exact finite computation, with no floating-point decision anywhere in it. [Math] means proved in ordinary mathematics in the sources but not formalised. [Cited] means established in the published literature. [Open] means not proved. These are never blurred, and a statement carrying one band is never described in language belonging to another.

Scale. Every parameter-indexed statement is tagged scale:fixed (proved at specific listed values), scale:bounded (proved below an explicit bound), scale:cofinal (proved for

arbitrarily large parameters), or scale:uniform (proved for all parameters). The corpus contains theorems at all four scales; the unresolved endpoints require particular infinite-orbit or cofinal producers that are not proved. Sorting by scale keeps that boundary visible.

Coordinates. Every statement records the representation it is expressed in. An obstruction is a fact about a coordinate, not about the object, and a wall measured in one representation may simply not exist in another; the atlas section gives the transport maps, so any obstruction recorded here can be re-measured elsewhere.

Erdős Problem 257 is open. No section of this document claims otherwise, and a reduction of the problem to another statement is recorded as a reduction, never as progress toward a solution.

1 The problem, and what is actually known

1.1 The object and the two exact statements

For an integer $n \geq 2$ set $x_n := 1/(2^n - 1)$, and for $A \subseteq \{2, 3, 4, \dots\}$ write $x_A := \sum_{n \in A} x_n$; the sum always converges, since $x_n = O(2^{-n})$. Erdős's Problem #257 asks a single universally quantified question about this family.

Definition 1.1 (Universal #257, written (U))**.** Is $x_A = \sum_{n \in A} 1/(2^n - 1)$ irrational for *every* infinite $A \subseteq \mathbb{N}$? **OPEN.** [Open]

Definition 1.2 (The half-value question, written (H))**.** Let the *Mersenne achievement set* be $\mathcal{A} := \{\, y \in \mathbb{R} : \exists A \subseteq \mathbb{N},\ 0 \notin A,\ y = \sum_{n \in A} x_n \,\}$ (Lean source). Is $1/2 \in \mathcal{A}$? **OPEN.** [Open]

The two are joined by one elementary and completely one-sided implication. Every $2^n - 1$ is odd, so every *finite* subset sum of $\{x_n\}$ has odd reduced denominator and cannot equal $1/2$ (Lean source). Hence a witness to $1/2 \in \mathcal{A}$ is necessarily an infinite A with $x_A = 1/2$ rational, which refutes (U) outright. Conversely, (U) implies $1/2 \notin \mathcal{A}$ as one instance among uncountably many, and a proof of $1/2 \notin \mathcal{A}$ leaves (U) entirely open. So (H) is the maximally symmetric candidate counterexample to (U), not a restatement of it; the asymmetry is load-bearing and recurs below as Barrier 2.5. [Math]

Nothing in this document decides (U) or (H). What it does contain is (i) the unconditional record, stated with hypotheses and evidence bands, and (ii) a characterisation of *why* the recorded attacks fail — assembled so that the failures function as measurements of one obstruction rather than as a list of disappointments.

1.2 The unconditional record, universal direction

Every landed unconditional irrationality theorem for (U) instantiates one engine: a weighted-coefficient block certificate over the divisor incidence $\mathrm{sc}_A(n) := \#\{d \mid n : d \in A\}$, the generic form being `irrational_coeff_series_of_weighted_coeff_block_certificates`. The complete inventory of what that engine has been made to prove, uniformly in the base $b \geq 2$:

Support class	**Exact hypotheses**	**Site** [Lean]
Full support $A = \mathbb{N}$	$b \geq 2$	Lean source
Multiples $A = d\mathbb{N}$	$b \geq 2$, $d \geq 1$	Lean source
Purely periodic	m-periodic, contains a positive element	Lean source
Eventually periodic	m-periodic from N_0, A infinite	Lean source
Residue class	any $m \geq 1$, any c	Lean source
Odd support (density 1/2)	—	Lean source
Lacunary (lcm gap)	a_k outgrows $\operatorname{lcm}\{a_0, \dots, a_{k-1}\}$	Lean source
Factorials; powers of two	instances of the lcm-gap theorem	Lean source, Lean source
Pairwise coprime (Erdős 1968)	A infinite, pairwise coprime, *and* $\sum_{a \in A} 1/a < \infty$	Lean source
Signed periodic weights	periodic integer weights, nonnegative, frequently nonzero	Lean source

Two facts about this table are worth a specialist's attention more than any single row. First, the base restriction of the classical results is gone: the $b = 2$ case of the first row recovers Erdős's 1948 irrationality theorem for the Erdős–Borwein constant $E = \sum_{n \geq 1} 1/(2^n - 1)$ as a one-line corollary, and every row holds for all $b \geq 2$. Second, and more important, *every* row lies inside a two-axis box — divisor independence plus a global reciprocal-tail budget — and the box is not an accident of effort. That is Barrier 2.7.

The one conditional extension worth naming is Lean source, which strictly generalises pairwise coprimality (cores may be arbitrary divisors of one fixed Q) but still demands globally pairwise-coprime petals and summable reciprocal petal mass, and whose selector hypothesis `SunflowerForcedSlotTailSelection` is itself an unproved cofinal supply. Despite the file name there is no dichotomy theorem in that file. [Lean][Open]

Two one-sided rigidity results constrain a hypothetical rational-valued support, both only from the sparse side. For any A with a positive element and x_A rational, divisor-free windows are at most logarithmic: for every $\varepsilon > 0$ there is B with every zero window of sc_A beyond the scale threshold of length $\leq \varepsilon \log_2 N + B$, uniformly in A, N and the window position (Lean source). [Lean] This is vacuous whenever $2 \in A$, since then every zero window already has length ≤ 1. The companion Lean source is essentially tight and is best read as a restatement rather than a constraint.

Remark 1.3 (What is and is not covered in the universal direction). The pinned Lean corpus contains no theorem for $A =$ the primes, but this is a formalisation boundary, not

a mathematical open case: Tao and Teräväinen prove the base-2 prime-support value irrational by first identifying it with $\sum_{n\geq 1}\omega(n)2^{-n}$. Duverney and Tachiya likewise settle squarefree support at every power-of-two base. Neither result is formalised here. The corpus still has no general theorem for an arbitrary positive-density nonperiodic or divisor-dense support, and the universal statement remains open.

1.3 The unconditional record, half-value direction

Theorem 1.4 (Achievement-set geometry). *$\mathcal{A}$ is compact, closed, perfect, totally disconnected and nowhere dense, and* volume$(\mathcal{A}) = 1$ *— inside an ambient span of $E \approx 1.6067$, a relative density of about 62%. The positive-index digit coding onto $\mathcal{A}$ is injective: each achievable real has* exactly one *support.* *coord:achievement-set scale:uniform [Lean]*
Sites: *Lean source*, *Lean source*, *Lean source*, *Lean source*, *Lean source*, *Lean source*.

Theorem 1.5 (Support-restricted refinement). *For the achievement set built from a coordinate subset $J \subseteq \mathbb{N}$: either the omitted coordinate set is finite, and the volume is exactly $2^{-|J^c|}$, or infinitely many coordinates are omitted and the volume is 0. Injectivity survives every restriction; perfectness is proved when J is infinite. No perfectness claim is made for finite J, whose coding range is finite.* *coord:achievement-set scale:uniform [Lean]*
Sites: *Lean source*, *Lean source*, *Lean source*, *Lean source*.

Theorem 1.6 (Membership equals greedy survival; the fatal-gap dichotomy). *$x \in \mathcal{A}$ iff $0 \leq x$ and the canonical greedy (take-if-possible) residual after every rank n satisfies $\rho(x,n) \leq T_{n+1} := \sum_{j>n} x_j$. Unconditionally, either $1/2 \in \mathcal{A}$ or a finite fatal gap exists, and the two are exactly complementary.* *coord:greedy-orbit scale:uniform [Lean]*
Sites: *Lean source*; *Lean source*; *Lean source*; *Lean source*; *witness type Lean source*.

Two boundary degeneracies are excluded at every rank, unconditionally: $\rho_k \neq x_k$ by the odd-denominator argument above, and $\rho_k \neq T_{k+1}$ because the tail differs from E by a rational while E is irrational. Both greedy decisions are therefore strict at every rank. [Math]

Finally, the certified computation. All of it is exact integer arithmetic, no floating point. Truncation rungs $J = 3, \ldots, 22$ all *proved* to survive (not merely computed). Seam orbit certified to row 200,000: branch counts R:M:U = 100197:49899:49898, the target inequality failing at exactly one reset row, $r = 7$; maximal pure-R run 19, attained at row 158,096. Independently, rows $6 \leq n \leq 2500$: 1209 resets, zero classification anomalies, minimum in-scope margin +1.1119 bits at row 14 growing to $\approx$ 1246 bits by row 2500 (mean 627), never re-approaching the threshold. Sharp tail margin, ranks 2 to 3000: no fatal skip, take fraction 0.4992, tightest take-margin +0.0340 bits at rank 7 (the near-tie $1/126 \geq 1/127$ — the single tightest real data point in the corpus, and the first case any proposed proof should be tested against). Greedy set certified to $m = 200{,}000$ with $0 \leq 1/2 - \sum_{d\in G} x_d < 2^{-199999}$. A half-trapping macro receipt reaches depth 13,548,057 (labelled in its own source as a private exact finite computation receipt, not publication authority). Kernel-checked corroboration of the low rows: Lean source verifies rows 13–30 by `decide`. [Cert]

Remark 1.7 (Literature, exactly). E is irrational (Erdős 1948). Zudilin's irrationality measure is $\mu(E) \leq 2.42343562\ldots$ (*Math. Notes* 72 (2002) 858–862; the erratum, *Acta Arith.*

111 (2004) 153–164, gives 2.46497868 …). Transcendence of E is unknown (Duverney–Tachiya, *Forum Math.* 31 (2019) 1557–1566), and the Erdős–Graham conjecture that every such subset sum is irrational — i.e. (U) — is itself open. For binary digit runs of E the only rigorous result is a *lower* bound: Erdős constructs a 0-run of length $\geq c \log^{1/10} N$ among the first N bits; Crandall (*INTEGERS* 12 (2012) #A23, §7) records the longest observed 0-run as 47 in the first 2^{43} bits and states that he does not know how to show that 11 appears infinitely often. Campbell (arXiv:2605.24160, 2026) settles that occurrence question affirmatively. [Cited]

That last pair of citations is the sharpest single orientation fact available. Every rigorous result in the literature about the digits of E points the wrong way: they produce long runs, or guarantee that patterns occur. What (H) needs is that a pattern *never* occurs for long. No result of that logical type exists.

2 The wall

2.1 The thesis

This programme has produced many reformulations of #257 and no proof of it. That is the honest summary, and expert readers have said so. The claim of this section is that the reformulations are not a hundred separate failures. They are repeated measurements of *one* obstruction, taken in different coordinates, and once they are assembled the obstruction has a describable shape. A systematic elimination that says exactly which classes of argument cannot work, and why, is a result about the problem; a reformulation is not. The elimination is what follows.

The generating fact is a single structural one, and it is not a convenience of the formalisation.

Observation 2.1 (The problem admits no ensemble). The Mersenne weights are strictly superincreasing. Consequently the digit coding is injective (Theorem 1.4), the greedy orbit is the unique candidate support for any target, and distinct subset sums of a row of seam weights are separated by at least 2^{s+1}. Every one of the nine coordinates the corpus has built — truncation rung J, seam row s, integer margin, sharp tail margin, half-carry K_n, reverse-carry word, Boolean–Möbius exact row, fatal-cell interval packing, Dedekind cut-locator — is therefore a coordinate on the binary digits of *one* specific real orbit. There is no family to range over. coord:meta scale:uniform [Lean][Math]
Sites: Lean source, Lean source, Lean source.

All six barriers below are consequences of Observation 2.1 in different directions: there is nothing to choose (2.2), nothing weaker to prove (2.3), nothing to average over (2.6), no information at the available scales (2.4), no certificate in the provable direction (2.5), and, on the universal side, no way to make a general support satisfy the certificate engine's first hypothesis (2.7).

2.2 Barrier 1: canonicalisation — the witness space is a singleton

Proposition 2.2 (Canonicalisation; PROVED)**.** *For the half-target in strictly superincreasing Mersenne weights, every finite object satisfying an approximate-hit condition at depth d* is *the*

canonical greedy prefix at depth d. Precisely: (i) every straddling word at depth d agrees bit-for-bit with the greedy support of $1/2$ on $(0, d]$; (ii) every below-half core D whose crossing deficit is smaller than the next weight is forced to equal `halfGreedyPrefixSupport`$(c-1)$*, so the universally quantified $\forall D$ in the lane's headline socket collapses onto one orbit; (iii) a remainder below the gap is equivalent to being the integer greedy bit-word.* *coord:straddle-prefix / integer-greedy scale:uniform [Lean]*

Sites: *Lean source*; *Lean source*; *Lean source*.

Rules out. Every argument whose leverage comes from *choosing* a favourable witness: a support D, a Boolean word, a straddle prefix, an adversarial late word, or a lattice/covering configuration. At each depth the admissible set has one element, so there is no choice to exploit and no adversary to defeat.

Explains. Three recorded attack families die for exactly this reason, and their deaths look superficially unrelated. (1) *Generic lattice covering* proves the wrong sign: the free prefix lattice breaches the $2^{t/2}$ safety radius in 74% of $t \in [8, 26]$ — witness $t = 15$, where fans of $\{2, 3, 5, 6, 10\}$ sum to $2^{15} + 1$, at circular distance 1; expected near-hits grow like $2^{t/6}$. Any generic bound on prefix cancellation therefore proves the opposite of what is wanted. [Cert] (2) *Digit-cascade impossibility* is false statically: explicit words exist at $t = 20$ (top 18 of 20 bits all-1) and $t = 22$ (top 20 of 22 all-0), and an unconditional construction supplies about $t/12$ independent single-bit knobs, verified with zero exceptions to $t = 601$. [Cert] (3) *Integer-window static adversary search* finds nothing: zero hits, exhaustively for $t \le 22$ and exactly at $t = 40, 42, 45$ over all cut positions. [Cert] The reconciliation is the content: residues mod 2^t are adversary-tunable, but the integer *magnitude* is not, because the row weights are gap-dominated. So the static tier and the dynamic tier coincide identically, and all three families were, without knowing it, measuring the same single orbit.

Remark 2.3 (A recorded hazard). The natural sufficient condition via fractional split-mass is *not* necessary on real crossing cores: the fixture $D = \{2, 3\}$, $c = 5$ is an explicit counterexample (Lean source). Any attack on the capacity socket that routes through fractional mass is already refuted. [Lean]

2.3 Barrier 2: route collapse — every "weaker producer" is the goal

This is the barrier that most consistently wastes effort, because the sockets it kills look strictly weaker on the page: they are Π^0_2 statements $(\forall N \exists n \dots)$ with explicitly relaxed compatibility requirements, while the goal is Π^0_1. The extra quantifier buys nothing. The mechanism is a single exact identity plus a single unconditional envelope.

Lemma 2.4 (Collapse mechanism; proved). *Let $A \subseteq \mathbb{N}$ with $1 \notin A$, put $\delta := 1/2 - x_A$, and let* $\mathrm{ihc}(A, N)$ *denote the integer half-carry at level N. Then*

$$\mathrm{ihc}(A, N) \;=\; 2^{N+1}\,\delta \;+\; \mathrm{T}(N+1), \qquad 0 \;\le\; \mathrm{T}(N+1) \;\le\; 2\sqrt{N+1} + 4,$$

where T *is the binary coefficient tail of* sc_A. *coord:half-carry scale:uniform [Lean]*

Sites: *Lean source*; *Lean source*; *Lean source*.

Proposition 2.5 (Route collapse; proved). *Any hypothesis asserting, for cofinally many N, that* $\mathrm{ihc}(A, N)$ *lies inside a strip of width $O(\sqrt{N})$ forces $\delta = 0$.*

Proof. By Lemma 2.4, $2^{N+1}\delta = \mathrm{ihc}(A,N) - \mathrm{T}(N+1)$. For the greedy support of $1/2$ one has $x_G \le 1/2$, hence $\delta \ge 0$ (Lean source). If $\mathrm{ihc}(A,N) \le c\sqrt{N} + c'$ for cofinally many N, then $2^{N+1}\delta \le c\sqrt{N} + c'$ along that sequence; since 2^{N+1} outgrows any $\sqrt{N}$ envelope and $\delta \ge 0$ is a fixed real, $\delta = 0$. Conversely $\delta = 0$ makes $\mathrm{ihc}(A,N) = \mathrm{T}(N+1)$ exactly, and the unconditional envelope of Lemma 2.4 then *supplies* the strip for free. And $\delta = 0$ is exactly $1/2 \in \mathcal{A}$, by Lean source. □

Lemma 2.6 (Perfect-square strip witness; PROVED, this pass)**.** *Suppose $x_A = 1/2$ and $1 \notin A$. Then for every $k \ge 1$,*

$$\mathrm{ihc}(A, k^2 - 1) \;=\; \mathrm{T}(k^2) \;\le\; 2k + 4 \;=\; \texttt{halfStripBound}(k^2).$$

coord:half-carry scale:cofinal [Math]

Proof. By Lemma 2.4 with $\delta = 0$, $\mathrm{ihc}(A, k^2 - 1) = \mathrm{T}(k^2) \le 2\sqrt{k^2} + 4$. At a perfect square the real and truncated square roots agree exactly, $\sqrt{k^2} = k = \lfloor \sqrt{k^2} \rfloor$, so the real envelope coincides with the discrete strip $\texttt{halfStripBound}(n) = 2\lfloor \sqrt{n} \rfloor + 4$ (Lean source) with no slack required. □

Lemma 2.6 matters because it removes the last technical obstruction previously recorded against two of the sockets below — the mismatch between $\lfloor \sqrt{\cdot} \rfloor$ and $\sqrt{\cdot}$, for which the corpus carries only the wrong-direction inequality. Restricting the cofinal witness depths to perfect squares dissolves it, and no widening of the strip constant from 4 to 6 is needed. For the terminal-only socket, take a to be the indicator of an achieving support truncated to $\{0, \dots, M\}$ at $M = k^2$, using Lean source and $1 \notin A$ (which holds since $x_1 = 1 > 1/2$).

Proposition 2.7 (The collapsed sockets; PROVED, composite)**.** *Each of the following is logically* equivalent *to $1/2 \in \mathcal{A}$, not strictly weaker:*

(a) *the* hbound *hypothesis of Lean source and of Lean source;*
(b) *Lean source;*
(c) *Lean source, including its advertised weakening that* no *compatibility is required between witness supports at different endpoints;*
(d) *Lean source;*
(e) *Lean source;*
(f) *the negation of Lean source, and all seven forms of the associated classification chain.*

Forward directions are landed Lean (e.g. Lean source, Lean source, Lean source). The reverse directions of (d) and (e) are derived here from Lemma 2.6 and are not *themselves landed Lean theorems.* *coord:half-carry / exact-row scale:cofinal [Lean][Math]*

Rules out. Every strategy of the form "the full problem is hard, so prove this weaker sufficient condition instead". In particular the specific hope that mutually *incompatible* witnesses at different endpoints buy flexibility is void, and reading the Π^0_2 shape of these sockets as evidence of extra room is an error.

Scope limit, preserved exactly. Item (a) does *not* extend to every A with $1 \notin A$. Dropping the greedy-specific fact $x_G \le 1/2$ yields only the one-sided statement hbound $\iff$ $(1/2 - x_A) \le 0$. This limit must be stated verbatim wherever the equivalence is used.

2.4 Barrier 3: scale mismatch — the available information is at the wrong resolution

The input the half-value branch actually needs is a per-index, growing-precision, *short-window* near-integer anti-concentration bound: at reset row r, a lower bound on the distance from $\mathbb{Z}$ of $\sum_{i\le L}\delta(M+i)\,2^{-i}$ with $\delta(M)=\tau(M)-1$, where the precision demanded grows like $2^{-r/2}$ while the window length is only $L\approx r/2$. Every information source tried lives at one of two extremes, and provably cannot reach that intermediate scale.

Class (a): fixed precision or bounded state. Three landed theorems, all problem-agnostic — they mention no Mersenne structure and bind #249 identically.

Proposition 2.8 (Local information is void; PROVED)**.** *(i) After L common steps the endpoint residue mod 2^L of an affine binary orbit is* independent *of the initial carry: $u(L)-v(L)=2^L(u_0-v_0)$ (LEAN SOURCE). (ii) No bounded or autonomous state summarising the history before position m determines the tail after m: for a balanced-pulse family no decoder recovers the radius, and any finite state type needs cardinality at least $\lfloor m/2\rfloor+2$, unbounded in m (LEAN SOURCE). (iii) For any starting carry and any finite word of odd valuation-unit symbols at fixed precision $u>0$, a compatible centred carry completion* exists *(LEAN SOURCE), so fixed-precision local data never excludes anything.* coord:carry scale:uniform [Lean]

The parity route is the same fact by hand: "$K(M)$ even $\iff$ M a perfect square" is true, but $J(M+1)=2J(M)-\delta(M+1)$ makes $J(M+1)\equiv\delta(M+1)\pmod 2$ an identity for *any* digit stream whatsoever. It is a tautology of the recursion shape; it forbids no residue and no run. [Math]

Class (b): global averages and Diophantine data.

Proposition 2.9 (Exponent-budget mismatch; PROVED, given a cited μ)**.** *The skip set at row n has $|\mathrm{Skip}_n|\approx n^2/4$, so an irrationality-measure exponent μ for the target constant yields a separation of order only $2^{-\mu n^2/4}$, while the required precision is $2^{-3n/2}$. The shortfall is a factor $\approx 0.42\,n$ in the exponent, and it grows linearly in n. Since $\mu\ge 2$ for every irrational, no improvement in μ can close it: Zudilin's $\mu(E)\le 2.42343562\ldots$ is irrelevant to the gap, and so is any future bound.* coord:diophantine scale:cofinal [Math][Cited]

Rules out. Any argument deriving the run-length or magnitude bound from (a) arithmetic information at fixed finite precision or a bounded automaton state, or (b) a global average, density, or Diophantine-approximation quantity attached to $C=E-3/2$. Concretely: no proposal to bound the R-run via an irrationality measure of a Mersenne or Erdős–Borwein-type constant can work, however good the measure becomes; and no proposal defining a bounded carry state summarising pre-index history can determine the post-index tail.

Explains. Prime-doubling fails for a recorded reason rather than for want of effort: N prime $\ge n$ gives $b(N)=1$ so J doubles at that step, but not cumulatively, because composites subtract. The explicit witness is the quiet run at $n=607$, spanning $M\in(607,617]$ with $J\in[1,5]$ for eleven steps *while three of those indices* — 607, 613, 617 — *are themselves prime*; and 26 of the 40 longest runs with $n\le 1500$ contain a prime. No run-length bound follows from prime counting. [Cert]

Remark 2.10 (Status of the exhaustion claim). The seven named families are each closed with mechanism. The claim that classes (a) and (b) *exhaust* the available tools is an OBSERVED PATTERN, not a theorem, and is labelled as such wherever it is used below.

2.5 Barrier 4: decision asymmetry — the certificates point the wrong way

Theorem 2.11 (One-sidedness; PROVED)**.** $1/2 \notin \mathcal{A}$ *is* Σ^0_1*: a single finite object — a fatal gap — certifies it.* $1/2 \in \mathcal{A}$ *is* Π^0_1*, and no finite object certifies it. Consequently no finite computation, at any depth, can establish* $1/2 \in \mathcal{A}$*.* *coord:meta scale:uniform [Lean]*

Proof. The dichotomy $1/2 \in \mathcal{A}$ or $\exists$ a fatal gap is unconditional (LEAN SOURCE), and the two alternatives are exactly complementary (LEAN SOURCE), with LEAN SOURCE a finite witness type. In the other direction, certified death implies non-membership (LEAN SOURCE, over LEAN SOURCE; 3/4 is certified dead at level 1, lookahead 0, by LEAN SOURCE), whereas the absence of such a certificate at any depth D is compatible with a certificate at depth $D+1$ and therefore proves nothing. Membership asserts survival at *every* rank (Theorem 1.6), an infinite conjunction with no finite sub-conjunction implying it. □

Rules out. Any programme whose plan is "extend the certified depth". Pushing the truncation rung to $J = 23$, the seam to row 10^6, or the macro receipt to depth 10^8 cannot in principle decide the branch — each raises a lower bound on a death rank which, if the supported branch is correct, does not exist.

Explains. Every certificate the programme owns is of this one type. The certified depths grew by roughly two orders of magnitude across the recorded history with no change whatsoever in decision status; that is not a coincidence to be pushed through, it is the theorem above being observed empirically. Accumulated survival is not convergence toward proof.

2.6 Barrier 5: measure and category point in opposite directions

Proposition 2.12 (The two soft heuristics contradict each other; PROVED)**.** $\mathcal{A}$ *has Lebesgue measure exactly* 1 *inside an ambient span of* $E \approx 1.6067$*, a relative density of about* 62%*, and is simultaneously nowhere dense, perfect and totally disconnected* (*Theorem 1.4*)*. Hence the measure heuristic — "a random point lies in* $\mathcal{A}$ *with probability* ≈ 0.62*" — predicts #257* false*, while the category heuristic — "*$\mathcal{A}$ *is meagre, so a generic point misses it" — predicts #257* true*.* *coord:achievement-set scale:uniform [Lean]*

Rules out. Every soft argument: measure-theoretic, Baire-category, metric number theory (almost-all / almost-none), and probabilistic digit models. The two standard notions of "typical" disagree on this object, so neither can be the mechanism. Worse, by injectivity of the digit coding (Observation 2.1) there is no ensemble of candidate representations to average over: each rational has exactly one candidate support, so counting, second-moment and pigeonhole arguments have nothing to range over. It also rules out deciding (U) by showing $\mathcal{A}$ is "small": it is not small — it is 62% of its ambient interval by measure.

Explains. The programme's own probabilistic estimate sits squarely on the losing side of this split: under a Bernoulli$(1/2)$ digit model the expected number of counterexamples with $n \geq 32$ is about 2×10^{-4}, i.e. the model predicts $1/2 \in \mathcal{A}$. That estimate is

worth exactly as much as the model, and the model is a measure-side object that cannot see a specific point of a nowhere-dense set having Lebesgue measure exactly 1. [Cert]

2.7 Barrier 6: the two-axis certificate box (universal direction)

Observation 2.13 (The box; OBSERVED PATTERN). Every landed unconditional theorem in the table of §1 requires the support to sit inside a two-axis box: **(i) divisor independence**, so that a CRT construction can prescribe the incidence sc_A across a block — pairwise coprimality, periodicity, or dilation of a full support; **and (ii) a global tail budget**, so that the middle window's mass satisfies the height inequality $q(C + N + L + 2) < b^L$ — summable reciprocals, or sparsity that makes this trivial. No landed theorem sits outside the box. This is a complete audit of the landed inventory, not a theorem that no such support can exist. coord:certificate-engine scale:uniform [Lean][Math]

Rules out. Any plan to obtain (U) by instantiating the existing engine at a general infinite A. For a general support the divisibility conditions $a \mid n$, $a \in A$, are *entangled* through pairwise gcds, so no CRT construction can prescribe the incidence in a block, and the engine's first hypothesis is unsatisfiable by construction rather than by search failure. It also rules out closing the gap by "removing the summability bound" from Erdős 1968: the bound is consumed at exactly one step, `CertificateKernel.lean:9586--9598`, as a global tail budget. Removing it is not weakening a hypothesis; it requires a different, windowed argument.

Explains only a method boundary. The primes are infinite and pairwise coprime — axis (i) is satisfied in the cleanest possible way, and LEAN SOURCE makes the dilation algebra exact there — yet this certificate theorem does not apply because $\sum 1/p$ diverges by Mertens and axis (ii) fails. This explains the absence of a *Lean theorem from this engine*; it does not make the base-2 value open, since Tao–Teräväinen settle it by a different analytic method. The sunflower generalisation inherits both axes and therefore does not recover that external theorem.

2.8 What the six barriers converge on

Both halves of #257 terminate at the same missing object, stated twice in two coordinates. On the half-value side it is a near-integer anti-concentration bound for $\sum_{i \le L} (\tau(M+i) - 1)\, 2^{-i}$, a geometrically weighted divisor sum in a short interval. On the universal side it is the dense-branch producer: for a divisor-dense A, a block of K consecutive indices with $2^K \mid \sum_{r=1}^{K} \mathrm{sc}_A(N + r)\, 2^{K-r}$ and a bounded middle window. These are the same statement in the incidence coordinate. No result of that shape exists in the literature in any form, and the only rigorous results in that direction — lower bounds on runs, and guarantees that patterns occur — have the opposite logical type. That convergence, from two independent directions, is the strongest evidence the corpus offers that there is exactly one wall here rather than several.

3 What the wall does not block

Six barriers eliminate a great deal. What they do not eliminate is a short, specific list, and the point of the elimination is that this list is short. Each route below is given with

the exact statement it needs, the barrier it evades, and *why* it evades it. Where a route evades every proved barrier but sits inside the class indicted by an observed pattern, that is said plainly.

3.1 Route 1: the reset $\surd$-escape / R-run anti-concentration target

What it needs. Any one of three forms, in decreasing strength.

(i) *Run-length form.* For every reset row $r \geq 31$ the maximal pure-R run starting at r satisfies $L_r < (r-3)/2$; equivalently $|\mathrm{rem}(r+1) - 2^{r+1}| > 2^{(r+5)/2}$. Any sublinear bound $L_r = o(r)$ suffices.

(ii) *Lean-facing form, already fanned in.* `SeamUpperResetDyadicBandEscape`: for every $d \geq 13$ with `successorCarries` and every $j \leq d$, the reset charge $4 \cdot$ overshoot + abovePulse avoids the linear-width band $(2^{d-j+1} - 2(d+j),\ 2^{d-j+1}]$.

(iii) *Sign law alone.* At every reset, M-resets have dev > 0 and U-resets have dev < 0. Strictly weaker than either inequality above, and flagged in its own source as worth attempting first.

Sites. Lean source with consumer Lean source; quantifier collapse Lean source, which reduces the $\forall j \in [0,d]$ band check to one nearest-boundary index per row; finite verification Lean source; alternative fan-in Lean source with consumer at `:167` and base case Lean source.

Why it evades Barrier 2.3. This is the one reduction in the corpus that is genuinely *sufficient* rather than equivalent. It is a quantitative bound at every reset row, while $1/2 \in \mathcal{A}$ is qualitative; it can be false while the goal is true. There is no implication $1/2 \in \mathcal{A} \Rightarrow \surd$-escape anywhere in the corpus, and by the shape of the two statements there cannot be. Proposition 2.5's mechanism needs a strip hypothesis of width $O(\sqrt{N})$ around the half-carry; this route asserts a lower bound on a deviation, which the collapse identity does not manufacture.

Why it evades Barrier 2.5. It is Π^0_1 with certificates aligned with the branch being proved. The margins *grow*: +1.1119 bits at row 14, about 1246 bits by row 2500, mean 627, never re-approaching the threshold. So for once the finite computation is evidence for the target rather than for its unprovable complement. [Cert]

Why it evades Barrier 2.2. It is a statement about the one canonical orbit — exactly the object canonicalisation says everything reduces to. Here the singleton witness space is help, not obstruction: there is no adversary to beat, only one sequence to bound.

The trap it avoids. A *uniform constant* run-length bound is empirically false: runs grow like $\log_2(\mathrm{row})$. The route survives because any $L_r = o(r)$ suffices, and the observed slack is enormous — maximum R-run 19 at row 158,096 against a requirement of $(r-3)/2 = 79{,}046$, a margin of roughly 4000×. [Cert]

What it does not evade. Barrier 2.4. Its residual difficulty *is* the missing object of §2.7. This route and that barrier are the same frontier viewed from the two sides, which is the honest reason it is listed first: it is where the real mathematics has to be done, not a way around it. [Open]

3.2 Route 2: strict endpoint progress for exact rows

What it needs. Rule out the recycle branch returning to a bounded endpoint forever. Any one of:

(a) `SkippedCoreCriticalQuotientSupply`: for every below-half core D bounded by $[2, c)$ with genuine crossing deficit, $2^{(2c-2)-1} \le$ `localPrefixQuotient` (insert c D) $(2c-2)$;

(b) a direct proof that the crossing rank c produced by the recycle branch cannot repeat the value 4, or any bounded value, infinitely often;

(c) a growth law for c as a function of the incoming endpoint n.

By Proposition 2.2(ii), form (a) need only be checked along the *one* canonical orbit `halfGreedyPrefixSupport`, not searched over D; and by the landed sharp-capacity biconditional it is exactly the $c-2$ bit capacity test `localBinarySuffix` $D\,1\,(2c-2) < 2^{c-2}$. It must *not* be attacked via the fractional-mass route, which the fixture $D = \{2,3\}$, $c = 5$ refutes as non-necessary.

Sites. Lean source; the formal falsifier Lean source (the model $n \mapsto (n = 6)$ satisfies the seed and the exact transition shape yet is not cofinal, so the schema plus an endpoint-six seed provably does *not* give cofinality); conditional strict progress Lean source; unconditional recycling producer Lean source; socket Lean source; capacity iff Lean source; seed fixture $\{2,3,6\}$ at Lean source.

Why it evades Barrier 2.2. This is the single place where canonicalisation provably does *not* bite, and the reason is exact. Proposition 2.2(ii) requires both a real-valued below-half condition *and* the deficit clause. The exact-row predicate `ExactLocalMersenneHalfRow` imposes neither: it asks only that a floor quotient hit $2^{n-1}-1$ in $\mathbb{N}$. Floor division is not injective. So the per-endpoint witness space is an exponential search over $D \subseteq [2,n]$, is decidable for each n, and is not forced to be the greedy prefix. It is the only genuine degree of freedom left anywhere in the half-value branch.

Its relation to Barriers 2.3 and 2.5. The *cofinal* statement is still equivalent to the goal, by Proposition 2.7(c). But equivalence of truth value is not equivalence of proof difficulty when the existential has an exponentially large witness space rather than a singleton. And decidability of `ExactLocalMersenneHalfRow`(n) makes the exact-row set computable, so the data here is a two-sided observable rather than a one-sided survival lower bound — a real softening of Barrier 2.5, though not an escape from it.

3.3 Route 3 audit: the prime-support theorem is known outside this formal corpus

The base-2 prime-support value is not an open target:

$$\sum_{p \text{ prime}} \frac{1}{2^p - 1} = \sum_{n \ge 1} \frac{\omega(n)}{2^n}$$

by expanding each reciprocal Mersenne term geometrically and collecting divisors. Tao and Teräväinen prove the right-hand series irrational in Theorem 1.3 of *Quantitative correlations and some problems on prime factors of consecutive integers*. Their proof is analytic and is cited here; it is not formalised in the pinned Lean corpus.

The local pairwise-coprime certificate theorem still does not cover the primes because it assumes summability of $\sum_p 1/p$. That is a boundary of this engine, not a boundary of the mathematics. The previously proposed windowed-tail repair is therefore withdrawn as a research contribution. It may remain useful as an audit question about whether the existing Lean engine can reproduce the known theorem, but reproducing a cited result is not new progress on #257. The authors mention broader base and prime-power extensions with modifications omitted in the cited version; this paper records only the fully proved base-2 theorem.

3.4 Route 4: land the route-collapse equivalences

What it needs. Five biconditionals, each with $1/2 \in \mathcal{A}$ on the right-hand side: the `hbound` hypothesis of the two `mobiusCenteredHalfCarry` consumers; `CofinalPositiveHalfGreedySkips`; `CofinalExactLocalMersenneHalfRows`; `GreedyHalfCarryCofinalStripReturn`; `HalfCarryCofinalTerminalOnlyStrip`. The forward directions are landed. Every reverse direction runs through Lemma 2.6, stated once, whose only inputs are Lean source, Lean source, Lean source, and the evaluation of the square root at a perfect square. The scope limit of Proposition 2.7 must be stated verbatim in the docstrings.

Why it is here. This route does not attempt the wall; it *measures* it, which is exactly what makes it bankable. Each equivalence is a theorem — a biconditional between a named `Prop` and half-membership — not a reduction, and each one permanently retires a socket from the list of things worth chasing directly. Systematic elimination of this kind is a result about the problem in a way that a reformulation is not. The perfect-square normalisation is what makes it available now: it dissolves the $\lfloor \sqrt{\cdot} \rfloor$ versus $\sqrt{\cdot}$ mismatch that previously blocked two of the five, and it needs no widening of the strip constant.

3.5 Route 5: the sparse/dense trichotomy closure

What it needs. An unconditional

$$A \text{ infinite} \implies (\text{cofinally many super-logarithmic zero windows in } \mathrm{sc}_A) \vee (\forall q > 0, \text{ a weighted block certificate exists for } \mathrm{sc}_A).$$

The first disjunct contradicts Lean source whenever the series is rational; the second closes via the certificate engine. The sparse branch is done — that theorem, and independently Lean source, close it. What is required is the *dense-branch producer*: for A whose divisor incidence is bounded below infinitely often, construct a block of K consecutive indices with $2^K \mid \sum_{r=1}^{K} \mathrm{sc}_A(N+r)\, 2^{K-r}$, a bounded middle window C, and $q(C+N+L+2) < 2^L$.

Why it evades Barrier 2.7. It attacks the box from outside rather than instantiating inside it. It does not ask a general support to satisfy axis (i) or (ii); it asks for a dichotomy on the support's own divisor density, and then uses a different tool on each side. It evades Barrier 2.6 because it is a structural statement about the incidence function, not a measure or category claim.

The honest flag. The dense-branch producer is the same missing object as §2.7, restated in the incidence coordinate. That is not a reason to discard the route — the re-

statement is genuinely a different coordinate, and a proof might be easier there — but it is a reason not to count it as an independent line of attack. [Open]

3.6 Ranking, stated plainly

Route 4 is the only one that can be completed with the mathematics currently in hand; it produces theorems, not progress on the wall. Route 3 is the only one that would produce a new irrationality theorem for a named constant, and its missing input is a windowed Mertens estimate, which is ordinary analytic number theory rather than a new phenomenon. Route 2 is the only place in the half-value branch with a non-singleton witness space, which makes it the best target for search. Route 1 is where the actual obstruction lives, and Route 5 is Route 1 in another coordinate. Neither problem is close to being solved, and no route above is a proof of anything.

4 How to use this document

Read it in full before starting work. That instruction is not a courtesy. The recurring practical failure of this programme has not been difficulty; it has been rediscovery — the same reformulations, the same dead attack families, arrived at independently three times because no single document held the whole picture. This document is that picture for #257. Anything proposed without reading it has a substantial prior probability of already appearing above with its mechanism of death recorded.

Evidence bands are never blurred. [Lean] means kernel-checked, with the declaration name as the authority and the file:line as where it sat at the pinned checkpoint. [Cert] means exact finite computation, integer arithmetic, no floating point. [Math] means ordinary mathematics, proved but not formalised. [Cited] means published, with the reference given. [Open] means not proved by anyone. A composite tag means the statement has parts at different bands. Declarations in the per-problem tree carry an `Erdos257/` path prefix; everything else lives in the main package.

The scale axis is the load-bearing one. Statements carry scale:fixed (a specific index or depth), scale:bounded (all indices up to a parameter), scale:cofinal (arbitrarily large indices), or scale:uniform (all indices, or independent of the parameter). The single most important structural fact about the corpus is that its certified content sits on the fixed/bounded side — rungs $J \leq 22$, seam rows $\leq 200{,}000$, ranks ≤ 3000, macro depth $\leq 13{,}548{,}057$ — while both open obligations are cofinal. By Theorem 2.11 that gap is a gap in kind, not in degree: no amount of extension on the fixed/bounded side becomes a proof. Read a scale tag as a hard type, not as a measure of how far along something is.

Obstructions are coordinate-relative. Every obstruction below carries a coordinate tag, such as coord:half-carry or coord:seam-row, because an obstruction is always an obstruction *in a representation*, never a property of the object. coord:half-carry facts say nothing about coord:seam-row arguments; a wall in the coord:diophantine coordinate (Proposition 2.9) is silent about the coord:certificate-engine coordinate. Before recording any new wall, name the coordinate it lives in and test at least one re-representation. Several of the barriers above are strong precisely because they are coordinate-free — Proposition 2.8 mentions no Mersenne structure at all and binds #249 identically — and that is worth stating explicitly when it is true, and not otherwise.

Map of what follows. The premise catalogue comes next: every landed statement with exact hypotheses, conclusion, coordinate, scale and site, organised by lane — the certificate kernel and universal-direction inventory; the greedy/achievement-set geometry; the seam model, its master identity and its dynamics; the half-carry and Boolean–Möbius coordinates; and the support rigidity results. After the catalogue come the typed implication tables, which record what chains with what, and the near-miss index, which records, for each open obligation, the statements that come closest and the exact mismatch in each case. The catalogue is a reference surface and may be read by lane. Sections 2 and 3 are not: they are an argument, and they should be read in order.

5 The object, the two questions, and the exact record

5.1 The object

Fix the Mersenne weight function $x_n := 1/(2^n - 1)$ for integers $n \geq 2$. For a set $A \subseteq \{2, 3, 4, \dots\}$, write $x_A := \sum_{n \in A} x_n$ when the sum converges (it always does, since $x_n = O(2^{-n})$). Erdős's Problem #257 asks about the family of series $\{x_A : A \text{ infinite}\}$.

Definition 5.1 (Universal Problem #257)**. Is $x_A = \sum_{n \in A} 1/(2^n - 1)$ irrational for *every* infinite $A \subseteq \mathbb{N}$?** This is a single universally-quantified statement over all infinite subsets of $\mathbb{N}$. It is **OPEN**. [Open]

The record in this paper does *not* decide universal #257. Everything that follows is either (i) a proof of irrationality for specific, named infinite-support families – which narrows but does not close the universal statement, since infinitely many other supports remain unchecked – or (ii) an attack on one specific, maximally symmetric candidate counterexample, described next.

Definition 5.2 (The distinguished half-value question)**.** Define the *Mersenne achievement set* $\mathcal{A} := \{\, y \in \mathbb{R} : \exists A \subseteq \mathbb{N},\ 0 \notin A,\ y = \sum_{n \in A} x_n \,\}$, the set of reals realizable as a subset sum of the positive-index Mersenne weights $\{x_n\}_{n \geq 1}$. **Is $1/2 \in \mathcal{A}$?** This is **OPEN**. [Open]

Observation 5.3 (Why the half-value question controls universal #257). Any finite subset sum of $\{x_n\}$ has odd reduced denominator, since every $2^n - 1$ is odd; hence no finite A can sum to $1/2$ (Lean source). Consequently, *if* $1/2 \in \mathcal{A}$, the witnessing support A is necessarily infinite, and $x_A = 1/2$ is rational. Such an A is then a single infinite set for which the universal-#257 sum is rational – an outright refutation of universal #257. Conversely, universal #257 being true forces $1/2 \notin \mathcal{A}$ as one instance among uncountably many. The half-value question is thus the maximally symmetric, most heavily studied test case for universal #257's negation, not a restatement of it: proving $1/2 \notin \mathcal{A}$ leaves universal #257 completely open (all other infinite A remain unchecked), while proving $1/2 \in \mathcal{A}$ closes universal #257 immediately, in the negative. scale:n/a [Math]

5.2 What is proved

5.2.1 Full-support irrationality, for every base

Theorem 5.4 (Full-support irrationality, unconditional, every base $b \geq 2$)**.** *For every integer $b \geq 2$,*

$$\sum_{k=0}^{\infty} \frac{1}{b^{k+1}-1} \quad \textit{is irrational.}$$

Hypotheses: $b \in \mathbb{N}$, $b \geq 2$. ***Conclusion:*** *unconditional irrationality of the full-support Erdős–Borwein-type series at every base, not merely base* 2. *coord:full-support scale:uniform [Lean]*
Site: LEAN SOURCE (*'b : ℕ, hb : 2 ≤ b'*), *landed via a certificate machine (bounded Bertrand/CRT first-block frame, middle-window divisor-pair average with pigeonhole selection, explicit parameter closure). The base-2 instance recovering the classical Erdős–Borwein constant's irrationality is* LEAN SOURCE, *a one-line corollary.*

This is the base-b generalization of Erdős's original 1948 irrationality theorem for $\sum 1/(2^n-1)$ (cited, in its base-2 form, as an input at Remark 5.12 below); its Lean form removes the base restriction entirely and is landed with no open hypotheses.

5.2.2 Named infinite-support families

Beyond full support ($A = \{1, 2, 3, \dots\}$, i.e. $m = 1$ below), the following named infinite families are each proved irrational unconditionally, for every base $b \geq 2$:

Theorem 5.5 (Purely periodic support)**.** *For every $b \geq 2$, every modulus $m \geq 1$, and every m-periodic $A \subseteq \mathbb{N}$ (i.e. $n+m \in A \Leftrightarrow n \in A$ for all n) containing a positive element, the support series $\sum_{a\in A} 1/(b^a-1)$ is irrational.* ***Hypotheses:*** *$b \geq 2$, $m \geq 1$, A m-periodic, $\exists a > 0$, $a \in A$.* *coord:periodic-support scale:uniform [Lean]*
Site: LEAN SOURCE.

Theorem 5.6 (Eventually periodic support)**.** *For every $b \geq 2$: if $A \subseteq \mathbb{N}$ is infinite and its membership is m-periodic from some threshold N_0 on (i.e. $n+m \in A \Leftrightarrow n \in A$ for all $n \geq N_0$), then $\sum_{a\in A} 1/(b^a-1)$ is irrational – proved by transferring irrationality across the finite symmetric difference from the shifted purely periodic set $A_{\text{pure}} := \{n : n+mN_0 \in A\}$.* ***Hypotheses:*** *$b \geq 2$, $m \geq 1$, $N_0 \in \mathbb{N}$, eventual m-periodicity from N_0, A infinite.* *coord:eventually-periodic-support scale:uniform [Lean]*
Site: LEAN SOURCE.

Theorem 5.7 (Residue-class support)**.** *For every $b \geq 2$, modulus $m \geq 1$, and residue c: $\sum_{n\equiv c \pmod m} 1/(b^n-1)$ is irrational.* *coord:residue-class-support scale:uniform [Lean]*
Site: LEAN SOURCE (*specializes purely periodic support*).

Theorem 5.8 (Odd support – first density-strictly-between-0-and-1 class)**.** *For every $b \geq 2$: $\sum_{n \text{ odd}} 1/(b^n-1)$ is irrational. This is the first landed support class with natural density strictly between 0 and 1 (odd n has density $1/2$ inside $\mathbb{N}$), as opposed to the density-1 full-support case or density-$1/m$ residue classes for large m.* *coord:odd-support scale:uniform [Lean]*
Site: LEAN SOURCE (*specializes residue-class support at $m = 2, c = 1$*).

None of these families contains A of the shape produced by a candidate half-value witness (the greedy support for $1/2$, described in §5.3, is not eventually periodic in any landed sense) – so Theorem 5.4 and its periodic-family corollaries neither prove nor disprove $1/2 \in \mathcal{A}$; they exhaust only the structured-support corner of universal #257's infinite search space.

5.2.3 Topology and measure of the achievement set

Theorem 5.9 (Achievement-set topology and measure). $\mathcal{A}$ (*Definition 5.2*) *is compact, closed, perfect, totally disconnected, and nowhere dense; its Lebesgue measure is exactly* 1: $\text{volume}(\mathcal{A}) = 1$. ***Hypotheses:*** *none.* ***Conclusion:*** $\mathcal{A}$ *is a Cantor-like set of Lebesgue measure exactly* 1 *– topologically thin (nowhere dense and totally disconnected), but neither full measure in* $\mathbb{R}$ *nor full measure in its ambient interval* $[0, E]$*, whose length is* $E \approx 1.6067$. *coord:achievement-set-topology scale:n/a [Lean]*
Site: Lean source, Lean source, Lean source, Lean source, Lean source, Lean source. *The proofs use only superincreasingness of* $\{x_n\}$ *and summability, so the argument transfers verbatim to any strictly superincreasing weight series.*
Chains with: *Theorem 5.10 below (compactness is exactly what powers every "limit of achieved points is achieved" argument used downstream, including the seam-limit route of Part 2).*

Theorem 5.10 (Membership equals greedy survival at every level). $x \in \mathcal{A} \iff 0 \le x \wedge \forall n,\ \rho(x,n) \le T_{n+1}$, *where* $\rho(x,n)$ *is the canonical greedy (take-if-possible) residual after rank* n *and* $T_{n+1} := \sum_{j>n} x_j$ *is the full remaining tail.* ***Hypotheses:*** *none (unconditional biconditional).* ***Conclusion:*** *membership in* $\mathcal{A}$ *is exactly the statement that the greedy process never enters a fatal state at any rank – this is the master theorem underlying every sufficient-condition result for* $1/2 \in \mathcal{A}$ *used later in this paper. coord:greedy-survival scale:uniform [Lean]*
Site: Lean source *(forward direction* Lean source*, reverse* Lean source*).*

5.3 The greedy orbit, exactly

Fix the target $1/2$ and the weights $x_n = 1/(2^n - 1)$, $n \ge 2$ (rank 1 contributes the undefined x_1 and is excluded by convention; all indexing below starts at rank 2). The greedy orbit is the unique deterministic process forced by superincreasingness (Theorem 5.10):

Definition 5.11 (The half-value greedy orbit). Let $\rho_1 := 1/2$ (the residual entering rank 2). At rank $k \ge 2$, given residual ρ_k:

- **Take** iff $\rho_k \ge x_k$; if taken, $\rho_{k+1} := \rho_k - x_k$.
- **Skip** otherwise; $\rho_{k+1} := \rho_k$.

A skip at rank k is:

- **safe** iff $\rho_k \le T_{k+1} := \sum_{j>k} x_j$ (the residual still fits inside the remaining tail, so future ranks can still reach it);
- **fatal** iff $\rho_k \in (T_{k+1}, x_k)$ (the residual is too large for the remaining tail to ever reach, yet too small to have been taken at k – an impossible gap).

By Theorem 5.10, $1/2 \in \mathcal{A}$ iff no rank of this orbit is ever fatal. Two boundary degeneracies are excluded unconditionally and independently of any open hypothesis:

Remark 5.12 (No-ties: both greedy-orbit boundaries are strict). (a) $\rho_k \ne x_k$ at any rank: a finite Mersenne sum has odd reduced denominator (each $2^n - 1$ is odd), so no finite prefix mass can equal x_k exactly, sharpening the general denominator-parity fact used in Observation 5.3. (b) $\rho_k \ne T_{k+1}$ at any rank: the tail $T_{k+1} = E - 1 - \sum_{2 \le j \le k} x_j$, where $E := \sum_{n \ge 1} x_n$ is the Erdős–Borwein constant, is irrational by Erdős's own 1948 theorem (the base-2 case of Theorem 5.4), while ρ_k is always rational; equality is impossible. **Conclusion:** both greedy-orbit boundaries are provably strict at every rank, so the orbit's

decisions stabilize after finitely many strict inequalities at any fixed truncation depth. coord:no-ties scale:uniform [Math]

5.4 The seam model and the master identity

The seam model is a base-4 rescaled integer reformulation of Definition 5.11, built so that every quantity involved is an exact natural number (no real-number or limiting arithmetic). Fix a row $n \geq 6$. The seam weights at row n are

$$w(n,d) := \left\lfloor \frac{4^n}{2^d - 1} \right\rfloor, \qquad d = 2, \dots, n-1,$$

and the seam target is $T_n := 2^{2n-1} - 2^n$. The seam greedy process takes ranks $d = 2, \dots, n-1$ in ascending order against T_n exactly as in Definition 5.11, producing a take set $D_n \subseteq \{2, \dots, n-1\}$ and skip set $\mathrm{Skip}_n := \{2, \dots, n-1\} \setminus D_n$. Define the *seam remainder* $\mathrm{rem}(n) := T_n - \sum_{d \in D_n} w(n,d)$ and the *seam deviation* $\Delta_n := \mathrm{rem}(n) - 2^n$.

Definition 5.13 (The universal recursion $K(M)$)**.** For $M \geq 2$, define

$$K(M) := 2^{M-1} - \sum_{d=2}^{M} \left\lfloor \frac{2^M}{2^d - 1} \right\rfloor.$$

K is *row-independent*: it depends only on M, not on any seam row n. It satisfies the recursion $K(2) = 1$, $K(M+1) = 2K(M) - (\tau(M+1) - 1)$, where τ is the number-of-divisors function (so $\tau(M+1)-1$ is the count of proper divisors of $M+1$). coord:universal-recursion scale:n/a [Cert]

Theorem 5.14 (Master seam identity)**.** *For every seam row $n \geq 6$,*

$$\Delta_n \;=\; K(2n) \;+\; \sum_{d \in \mathrm{Skip}_n} \left\lfloor \frac{4^n}{2^d - 1} \right\rfloor. \tag{I}$$

Hypotheses: *$n \geq 6$; $\mathrm{rem}(n)$, D_n, Skip_n as in the seam greedy process for target T_n over weights $w(n,d) = \lfloor 4^n/(2^d-1) \rfloor$, $d = 2, \dots, n-1$, ascending greedy.* ***Conclusion:*** *the entire $2n$-bit seam deviation collapses to one universal, row-independent recursion value $K(2n)$ plus a short sum restricted to the skipped ranks only – every rank the greedy* takes *drops out of the identity entirely.* *coord:seam-row / integer-greedy (Mersenne weights, binary digits of K) scale:fixed [Math][Cert]*

Proof sketch* (shape preserved from the source record; the full two-line derivation lives in the source document and is not independently re-derived here):** *unfold $\mathrm{rem}(n) = T_n - \sum_{d \in D_n} w(n,d)$ and compare the full-range sum $\sum_{d=2}^{2n} \lfloor 4^n/(2^d-1) \rfloor$ (the sum defining $K(2n)$, evaluated at $M = 2n$ so that $2^M = 4^n$) against the seam row's own sum over $D_n \cup \mathrm{Skip}_n = \{2, \dots, n-1\}$; the ranks $d \geq n$ present in $K(2n)$'s defining sum but absent from the row's range, together with the taken ranks D_n, absorb exactly into $T_n - 2^n = 2^{2n-1} - 2^{n+1}$ against $K(2n)$'s leading 2^{2n-1} term, leaving only the skipped-rank sum as residue.* ***Certification: *verified rows $6 \leq n \leq 200$, zero mismatches, independently reproduced by two subagents (§5.7).*

Chains with: *any statement about $\tau(M+1)$ parity/growth, or about the binary digit stream of $K(M)$, composes directly through (I).*

5.5 The real form and the master identity's analytic reading

Theorem 5.15 (Real (non-integer) form of the master identity)**.** *Let $x_d := 1/(2^d - 1)$ and define the constant*

$$C := \sum_{d \geq 2} x_d - \tfrac{1}{2} = E - \tfrac{3}{2} = 0.1066951524152917\ldots,$$

where E is the Erdős–Borwein constant. Then, for every seam row $n \geq 6$,

$$\Delta_n = 4^n\Big(\sum_{d \in \mathrm{Skip}_n} x_d - C\Big) + \eta_n, \qquad |\eta_n| < 2n + 2. \tag{II}$$

Hypotheses: *as Theorem 5.14.* ***Conclusion:*** *escaping the wall (making Δ_n large, at scale $2^{(n+5)/2}$ or beyond – the exact target of Part 2's Theorem A) is exactly the statement that the greedy skip-set's real-valued partial sum $\sum_{\mathrm{Skip}_n} x_d$ never approximates the fixed constant C to within roughly $2^{-1.5n}$, at the natural scale 2^{-n} set by $4^n \cdot 2^{-n} = 2^n$.* coord:real/analytic (Erdős–Borwein tail); bridges the integer seam coordinate to a continuous Diophantine-approximation statement scale:fixed [Math][Cert]

Note: *$C = E - 3/2$ is Mersenne-specific, but the shape "scaled deviation $= 4^n \cdot$(finite skip-sum $-$ target constant) $+O(n)$" is a template, not yet matched to any analogous constant on the #249 side.*

5.6 Dynamics

Theorem 5.16 (Deviation dynamics and the R/M/U branch structure)**.** *Let $\lambda_n := \Delta_n/2^n$. Then*

$$\lambda_{n+1} = 2\lambda_n + 2e_n + \delta_n, \qquad e_n \in \{0, +1, -1\} \text{ for branches } R/M/U \text{ respectively}, \qquad |\delta_n| \leq \frac{2n+6}{2^n}. \tag{III}$$

Hypotheses: *as Theorem 5.14.* ***Conclusion:*** *the deviation dynamics is a doubling map on λ_n with a forced signed digit e_n (recording which of three branches – Right/Middle/Upper – fires at each row) plus a vanishing correction $\delta_n = O(n2^{-n})$.* ***Measured bands*** *(rows 14 through 699, certified computation): branch R $\lambda \in [-0.499189,\ 0.499846]$; branch M $\lambda \in [-0.998378,\ -0.501224]$; branch U $\lambda \in [0.501652,\ 0.999692]$ – the three bands are disjoint and together nearly cover $[-1, 1]$.* coord:seam-affine dynamics (doubling map with signed digit e_n) scale:bounded [Math][Cert]

Consequence: *the target statement (escaping the wall) is equivalent to: the R-run (maximal run of consecutive R branches) after each reset row $r \geq 31$ has length $L_r < (r-3)/2$; a long R-run is exactly what is dangerous, and any $L_r = o(r)$ closes it, while a* uniform *constant bound on run length is empirically* ***false*** *(observed run lengths grow like $\log_2(\text{row})$, §5.7).*

Chains with: *this is the exact object whose long-run/anti-concentration behavior is the paper's central open producer (Part 2 develops it in full; the shape recurs, independently derived, under the coordinate name $w_n = \mathrm{rem}(n) - 2^n$ in the reset-crossing unification record).*

5.7 Every anchor

The following are exact, certified numerical facts about the seam orbit of §5.4–§5.4, stated literally so they can be used as regression fixtures without re-running any computation.

Observation 5.17 (Exact low-row values). At row $n = 14$: $\mathrm{rem}(14) = 392$, with take set $D_{14} = \{2, 3, 6, 7\}$. At row $n = 15$: $\mathrm{rem}(15) = 34333$. At row $n = 32$: $\mathrm{rem}(32) = 3865046005$ and $\Delta_{32} = -429921291$. coord:seam row, exact integer scale:fixed [Cert]
Site: the row-14 through row-31 range, including $\mathrm{rem}(14) = 392$, additionally carries independent [Lean] confirmation via the kernel-computed certificate Lean source, which lists the exact seam remainder at every row 13–31 by 'decide'-checked kernel computation and finds none violating the dyadic band-escape condition.

Observation 5.18 (Certified computation depth, all four coordinates). The master identity (Theorem 5.14) is verified rows $6 \le n \le 200$ with zero mismatches, independently reproduced by two subagents, and the seam orbit itself is extended and certified to row 2500 (and separately, in a longer single run, to row 200,000 – §SE-7 of the underlying record, 207.7s runtime, branch counts R:M:U = 100197:49899:49898, TARGET failing at exactly one reset row $\le$ 200,000, namely $r = 7$). coord:multiple: seam row, binary digit stream of C, greedy set G scale:fixed [Cert]

Observation 5.19 (The four coordinate depths, side by side). Definition 5.11's single orbit is independently certified in four distinct integerizations, to the following depths: truncation rung $J = 3, \ldots, 22$ (all 20 rungs proved to survive, not merely computed – Part 3 of the truncation-rung record); seam row $6 \le n \le 2500$ (Theorem 5.14/5.16); integer margin scan to $m_c = 5 \times 10^5$ (advisory Type B scan, not yet promoted to a certified theorem); sharp tail margin, ranks 2 through 3000 (no fatal skip at any rank; take fraction 0.4992; tightest certified take-margin 0.0340 bits at rank 7, the near-tie $1/126 \ge 1/127$). coord:meta: all four coordinates scale:fixed [Cert]
Chains with: these four depths measure the *same* greedy decision stream in different coordinates. They do not make the finite certificate families equivalent and do not give a semi-decision procedure for the universal problem. A certificate at one depth in one coordinate does not, by itself, extend to another coordinate without an explicit transport or stabilization theorem.

The record above is exhaustive at the level of exact statements and sites available in the premise bank for this Part; no numeral in this section was recomputed independently by the present writer – each is either directly Lean-kernel-checked (as flagged) or reported as a certified-computation result from the source record, and every scale/evidence tag reflects only what its cited source claims for it.

6 Catalogue of usable premises

This section catalogues machinery for Erdős #257 (is $\sum_{n \in A} 1/(2^n - 1)$ irrational for every infinite $A \subseteq \mathbb{N}$?) drawn from the half-membership route (is 1/2 in the Mersenne achievement set?) and the support-rigidity / full-support route (irrationality of $\sum_{a \in A} 1/(2^a - 1)$ for named families A). Every entry states exact hypotheses, exact conclusion, coordinate, scale, evidence band, and Lean site. Nothing below decides #257; the problem remains open. Entries are grouped **Producers** (hypothesis $\Rightarrow$ conclusion, including conditional sockets whose hypothesis is itself an open sub-obligation, always flagged) and **Converters and exact identities** (iff-reformulations, structural invariants,

definitions, exact numerals). Within each subsection entries are ordered by scale: fixed < bounded < cofinal < uniform < n/a.

6.1 Producers

Theorem 6.1 (Cofinal skip forces half-membership)**.** *Fatality is absorbing for the greedy Mersenne recursion at target $x \geq 0$: once* GreedyMersenneFatalAt x n *holds it holds at every $n + k$, and forces every later exponent selected, so a fatal state leaves only finitely many omitted exponents. Contrapositively: if the greedy skipped support (the set of ranks the greedy process does* not *select) is infinite, then x lies in the Mersenne achievement set.*

[Lean]scale:cofinalcoord:greedy-orbit LEAN SOURCE

Theorem 6.2 (Seam-limit convergence route to 1/2)**.** *Let* seamGreedyFiniteValue s *be the real point coded by the finite integer seam word at row s (always achieved) and* seamGreedyNormalizedRemainder s *the seam integer remainder rescaled by 4^{-s}. If there is a cofinal sequence of rows* rows *along which the normalized remainder tends to* 0 (SeamGreedyRemainderSubquadraticAlong rows)*, then* seamGreedyFiniteValue(rows j)→ 1/2 *and, by closedness of the achievement set,* 1/2 *is achieved.*

[Lean]scale:cofinalcoord:seam-integer LEAN SOURCE

Theorem 6.3 (Two-channel and dyadic cap sufficiency)**.** *If, at every rank n the greedy process for target* 1/2 *actually skips, the residual* greedyMersenneRemainder(1/2) n *is bounded above by the explicit rational envelope* halfTwoChannelCap(n+1) *(resp. the coarser* halfDyadicCap(n+1))*, then* 1/2 ∈ mersenneAchievementSet*. It therefore suffices to control the residual only at actually-skipped ranks, not at every rank.*

[Lean]scale:uniformcoord:two-channel-dyadic-cap LEAN SOURCE LEAN SOURCE

Theorem 6.4 (Second-channel phase separation sufficiency)**.** *If the greedy second-channel phase* (greedyMersenneSecondChannelPhaseRat) *stays separated* (HalfSecondChannelSeparatedRat) *at every skipped rank, then* 1/2 ∈ mersenneAchievementSet*; a rational specialization discharges the hypothesis from row 7 on.*

[Lean]scale:uniformcoord:second-channel-phase LEAN SOURCE LEAN SOURCE

Theorem 6.5 (Straddle-prefix closed-set criterion)**.** *If for every depth d there exists a finite word u (all exponents $\leq d$) straddling target t — meaning* value$(u) \leq t \leq$ value$(u) +$ mersenneTail(d) *— then $t \in$* mersenneAchievementSet*. The proof needs only compactness of the achievement set, no coinductive machinery.*

[Lean]scale:uniformcoord:dedekind-cut LEAN SOURCE

Theorem 6.6 (Largest-skip-late induction: single-socket reduction)**.** *Define* LargestSkipLateAt s *as "row s's largest false rank d satisfies $2s < 3d$" and* LargestSkipLateStepSocket *as the universally quantified one-step arithmetic implication* $\forall s \geq 14, d,$ IsLargestFalseRank(seamGreedyWord s) $d \to 2s < 3d \to (2(s+1) < 3d \vee$ SeamGreedyUpperOrMiddleAt s _$)$*. Given this socket, plus the base case at row* 14 *(kernel-verified by* `decide`*), the late-largest-skip invariant propagates to every row* $\geq$ 14 *by induction, supplying a cofinal sequence of unboundedly skipped ranks, hence*

$1/2 \in$ mersenneAchievementSet. *This is the single cleanest remaining reduction of the whole half-branch: closing it factors entirely through one local arithmetic step.*

[Lean]scale:uniformcoord:seam-integer LEAN SOURCE LEAN SOURCE

Theorem 6.7 (Middle-producer card/row escape sockets). *If, at every late middle-branch row,* card(support) + belowPulse + 5 $<$ 4 $\cdot$ remainder (SeamMiddleProducerCardEscape), *then* $1/2 \in$ mersenneAchievementSet. *The strictly weaker-hypothesis, stronger-conclusion socket* SeamMiddleProducerRowEscape ($s \leq$ remainder *at every late middle row*) *implies the card escape and hence also closes the branch.*

[Lean]scale:uniformcoord:seam-integer LEAN SOURCE LEAN SOURCE

Theorem 6.8 (Middle all-right tail forces strict defect). *If a middle transition at row* $D \geq 13$ (*with* ¬carries) *is followed by an all-right* (*eventually every rank selected*) *tail, then the producer carry is strictly below its complete future incidence tail: equivalently* seamGreedyFloorZ D $<$ seamTakeThreshold D, *equivalently a strict rational-skip defect. This overconstrains the "last transition is middle, then forever right" scenario, feeding the final-middle-cell exclusion* (*Theorem 6.44*).

[Lean]scale:uniformcoord:seam-integer LEAN SOURCE LEAN SOURCE

Theorem 6.9 (Two-sided dyadic cell escape: global remainder control). SeamTwoSidedDyadicCellEscape *is a minimal local socket* (*it excludes exactly three exceptional middle cells at signed coordinate* $-3, -2, -1$, *plus one right-pulse-leak bound*). *Granted this socket, plus a kernel-verified base case at row* 5, *induction gives the universal two-sided bound* $\forall s \geq 5$, seamIntegerGreedyRemainder $s \leq 2^s \vee$ overshoot $s \leq 2^s$ *— a global dyadic-scale control on the seam orbit independent of which branch fires at each row.*

[Lean]scale:uniformcoord:seam-integer LEAN SOURCE LEAN SOURCE

Theorem 6.10 (Upper-reset dyadic-band escape, with rows 13–31 certified). SeamUpperResetDyadicBandEscape *says: for every* $d \geq 13$ *and every* $j \leq d$, *the reset charge avoids a linear-width band immediately below the dyadic power* 2^{d-j+1} ($2^{d-j+1} <$ resetCharge *or* resetCharge $+ 2(d + j) \leq 2^{d-j+1}$). *Granted this,* $1/2 \in$ mersenneAchievementSet. *Rows* 13 *through* 31 *are unconditionally certified by kernel* `decide` *computation of the exact remainder at each row* (*e.g. row* $14 \to 392$, *row* $31 \to 4187487147$; *none failing*) *— the current computational frontier of this sub-route.*

[Lean]scale:uniformcoord:dyadic-boundary LEAN SOURCE LEAN SOURCE

Theorem 6.11 (Möbius-centred carry nonnegativity below 1/2). *For a support* A *with* $1 \notin A$: *if* erdosSupportSeries $2\,A < 1/2$, *then* $0 \leq$ mobiusCenteredHalfCarry $A\,N$ *for every* N. *Proof: the identity* integerHalfCarry $A\,N = 2^{N+1}(1/2 -$ seriesValue$) +$ binaryCoeffTail$(\dots)$ *writes the carry as a sum of two manifestly nonnegative terms. Stated over the generic divisor-incidence function* supportCoeff, *not over Mersenne weights directly — since* $\varphi = \mu * \mathrm{id}$ *is itself Möbius-built, this is a strong candidate for adaptation to #249.*

[Lean]scale:uniformcoord:mobius-mersenne LEAN SOURCE LEAN SOURCE

Theorem 6.12 (Square-root carry-growth route: current open numeric-analytic frontier). *Nonnegativity of the Möbius-centred carry of the actual canonical greedy support for*

$1/2$ *is unconditional* (*Theorem 6.11 plus two supporting lemmas*). If in addition *the pointwise upper bound* mobiusCenteredHalfCarry(greedySupport$(1/2))\,N \le 2\sqrt{N}+4$ *holds for every* N*, then the greedy skipped support is infinite and the greedy value equals exactly* $1/2$*. Only the upper bound remains open; a distinct socket from the dyadic-band* (*Theorem 6.10*) *and largest-skip-late* (*Theorem 6.6*) *routes. A separate project-memory note records a* different *sqrt-escape wall on the unconditional* ($\forall$ *infinite* A) *#257 track; whether that wall bears on this half-membership socket is unverified — treat any connection as* `hypothesis`*, not established.*

[Lean]scale:uniformcoord:greedy-orbit LEAN SOURCE LEAN SOURCE

Theorem 6.13 (Governed frozen-margin producer family: three graded sockets)**.** *Three sufficient sockets, strongest to weakest hypothesis, all funnel into* $1/2 \in$ mersenneAchievementSet*:* (*i*) HalfGreedySkippedFullShellNonnegative (*at every skipped rank,* $0 \le$ greedyHalfFrozenMargin$(n-1)\,n$*), direct and seam-free;* (*ii*) HalfGreedySkippedSeamAlignmentZero (*the seam remainder is exactly* 0 *whenever the actual word equals the seam-greedy word*)*, proved* equivalent *to* (*i*)*;* (*iii*) HalfGreedySkippedSeamEscape (halfStripBound$(2n) <$ seamIntegerGreedyRemainder n *at every skipped rank*)*, strictly sufficient for* (*i*) *but not equivalent.*

[Lean]scale:uniformcoord:seam-integer LEAN SOURCE LEAN SOURCE LEAN SOURCE

Theorem 6.14 (Erdős–Borwein full-support irrationality, unconditional)**.** *For every base* $b \ge 2$*:* $\sum'_{n\ge1} 1/(b^n-1)$ *is irrational, unconditionally. No open obligations. Base* 2 *is the classical Erdős–Borwein constant. This is the terminal target every other rigidity/support theorem in this catalogue is measured against.*

[Lean]scale:uniformcoord:mobius-mersenne LEAN SOURCE LEAN SOURCE

Theorem 6.15 (Pairwise-coprime support irrationality, Erdős 1968)**.** *For every base* $b \ge 2$ *and every infinite pairwise-coprime support* $A \subseteq \mathbb{N}$ *with summable reciprocals,* $\sum_{a\in A} 1/(b^a-1)$ *is irrational.*

[Lean]scale:uniformcoord:mobius-mersenne LEAN SOURCE

Theorem 6.16 (Generic weighted-coefficient block-certificate engine)**.** *For base* $b \ge 2$ *and coefficient* $c : \mathbb{N} \to \mathbb{N}$ *with* $c(m) \le m$*: if a weighted block certificate exists at every precision* q (*an explicit* $\exists N, K, L, C$ *package: block-divisibility, bounded middle window, far-tail positivity, height inequality* $q(C+N+L+2) < b^L$)*, then* $\sum'_m c(m+1)/b^{m+1}$ *is irrational. This is the single engine both the full-support theorem* (*Theorem 6.14, via* $c = \tau$*, the divisor-count function*) *and the #249 sockets* ($c = \varphi$) *instantiate; certificate supply is proved for* τ *and left open for* φ *— exactly why #249 stays open where #257 full-support closes.*

[Lean]scale:uniformcoord:binary-digit LEAN SOURCE

Theorem 6.17 (LCM-gap irrationality engine)**.** *For base* $b \ge 2$ *and strictly monotone support* $a : \mathbb{N} \to \mathbb{N}$ *with* $a(0) \ge 1$*: if* $a(k) - \mathrm{lcm}(a(0), \dots, a(k-1)) \to \infty$*, then* $\sum'_k 1/(b^{a(k)}-1)$ *is irrational. Pure statement about any strictly increasing exponent sequence and any base; no Mersenne-specific content.*

[Lean]scale:uniformcoord:lcm-gap LEAN SOURCE

Theorem 6.18 (Factorial-support and 2^k-support instances). *For every $b \geq 2$: $\sum'_k 1/(b^{(k+1)!}-1)$ is irrational, and $\sum'_k 1/(b^{2^k}-1)$ is irrational. Explicitly flagged in-source as* instances *of the #257 statement family, not the universal theorem — the universal problem (every infinite A) remains open.*

[Lean]scale:uniformcoord:lcm-gap LEAN SOURCE LEAN SOURCE

Theorem 6.19 (Multiples-support irrationality via dilation). *For every $b \geq 2$, $d \geq 1$: $\sum_{a \in d\mathbb{N}} 1/(b^a-1)$ is irrational, via the exact identity* erdosSupportSeries $b\,\{n : d \mid n\} = \sum'_k 1/((b^d)^{k+1}-1)$ *— dilation to base b^d full support, reducing directly to Theorem 6.14.*

[Lean]scale:uniformcoord:dilation LEAN SOURCE LEAN SOURCE

Theorem 6.20 (Periodic-support irrationality: the periodic divisor-orbit sieve). *For every $b \geq 2$ and every m-periodic support A ($\forall n,\ n+m \in A \leftrightarrow n \in A$) containing a positive element:* erdosSupportSeries $b\,A$ *is irrational. At $m = 1$ recovers full support; residue classes, unions of residue classes, and gcd-pattern supports are all instances. Mechanism: a prime $p \nmid m$ makes the divisor ray $d, pd, p^2d, \ldots$ sweep residues mod m periodically; a prime with exact valuation $b\varphi(m) - 1$ splits every divisor ray into complete residue cycles repeated a multiple of b times, reusing the same bounded Bertrand/CRT frame as full support with exponent $b-1 \to b\varphi(m)-1$.*

[Lean]scale:uniformcoord:periodic-sieve LEAN SOURCE

Theorem 6.21 (Eventually-periodic support irrationality). *For every $b \geq 2$: an infinite support whose membership is m-periodic from some threshold N_0 on has* erdosSupportSeries $b\,A$ *irrational.*

[Lean]scale:uniformcoord:periodic-sieve LEAN SOURCE

Theorem 6.22 (Residue-class and odd-support irrationality). *For every $b \geq 2$, modulus $m \geq 1$, residue c: $\sum_{n \equiv c\,(m)} 1/(b^n-1)$ is irrational. Specializing $m = 2, c = 1$: $\sum_{n\ odd} 1/(b^n-1)$ is irrational for every $b \geq 2$ — the first support class of density strictly between 0 and 1 to land.*

[Lean]scale:uniformcoord:periodic-sieve LEAN SOURCE LEAN SOURCE

Theorem 6.23 (Signed periodic-weight dichotomy and its unconditional one-sided closures). *For base $b \geq 2$ and any m-periodic integer-valued signed weight w: the weighted series $\sum w(a)/(b^a-1)$ is either irrational or base-b terminating ($b^k \cdot x = z$ for some integer k, z), unconditionally, mixed signs allowed. If additionally the divisor-coefficient* intWeightedCoeff w *is one-sided (all nonnegative, or all nonpositive) and not eventually zero, the terminating branch is excluded and the series is unconditionally irrational.*

[Lean]scale:uniformcoord:signed-divisor-calculus LEAN SOURCE LEAN SOURCE LEAN SOURCE

Theorem 6.24 (Mersenne-channel denominator survival, T4). *For a finite family P of "upper-half" primes (each p with $t < 2p$ for scale t, each dividing a squarefree radical r all of whose prime factors are $\leq t$): the pairwise-coprime Mersenne product $C = \prod_{p \in P}(2^p-1)$ survives fraction reduction — $C/\gcd(C,h) \mid (h \cdot \mathrm{mobiusNumerator}(r)/\mathrm{mersenne}(r)).\mathrm{den}$ — exactly when C is coprime to scale h (or, weaker, when every prime factor of h is $\leq t$). This is the load-bearing quantitative step behind Theorem 6.14.*

[Lean]scale:uniformcoord:cyclotomic LEAN SOURCE LEAN SOURCE

Theorem 6.25 (Mersenne-channel exponential denominator growth). *Via Bertrand's postulate, at every scale $t \geq 5$ there is an upper-half prime p with*

$2^{t/2} \le 2^{p-1} \le \text{mersenne}(p)$, *giving the explicit exponential lower bound* $2^{t/2} \le$ (*reduced denominator of the scaled Möbius shadow at LCM-height* t). *The exact reduced denominator at every LCM height is also computed in closed form.*

[Lean]scale:uniformcoord:cyclotomic Lean source Lean source

6.2 Converters and exact identities

Observation 6.26 (Half-greedy regression fixtures: exact rational checkpoints). Three closed numeric facts, `norm_num`/`decide`-checked: $1/2 - W(\{2,3,6,7\}) = 1/16002$ exactly (*finiteErdosSum*); $\text{mersenneCorrectionTail}(1) \le 5/42$; $3/8 < 1/2 - \text{mersenneCorrectionTail}(1)$. Anchors for validating any new automated fatal-gap search at small depth.

[Cert]scale:fixedcoord:regression-fixture Lean source Lean source

Observation 6.27 (Period-4 sign weight vanishes on a residue class). The explicit period-4 sign pattern $1, 0, -1, 0$ (*periodFourSignWeight*) has $\text{intWeightedCoeff}\, w\, n = 0$ identically for every $n \equiv 3 \pmod 4$, via the divisor-involution $d \mapsto n/d$. This is exactly why the general signed-periodic case (Theorem 6.23) only reaches a dichotomy rather than an unconditional theorem: no residue-blind selection can force a nonzero protected residue for every periodic signed weight.

[Lean]scale:fixedcoord:signed-divisor-calculus Lean source

Theorem 6.28 (Möbius–Mersenne Lambert identity). $\sum'_{d:\mathbb{N}^+} \mu(d)/(2^d - 1) = 1/2$ *exactly* (*imported as* MersenneLambertLadder.tsum_moebius_div_two_pow_sub_one_eq_half). *A signed* (*not Boolean*) *series over Mersenne denominators hitting the target value exactly — a landmark calibration point for rigidity arguments bounding how close a Boolean infinite support can get to* $1/2$.

[Cited]scale:fixedcoord:mobius-mersenne Lean source

Corollary 6.29 (Negative-Möbius Boolean support overshoots $1/2$). *Consequently the negative-Möbius Boolean support* $N = \{d \ge 2 : \mu(d) = -1\}$ overshoots $1/2$: $\sum_{d \in N} 1/(2^d - 1) = 1/2 +$ (*positive-Möbius tail*) $\ge 1/2 + 1/63$ (*using the* $d = 6$ *term,* $\mu(6) = 1$). *A route-sufficiency no-go only: it rules out this one natural candidate construction, not the general question of whether some infinite Boolean support sums to exactly* $1/2$.

[Lean]scale:fixedcoord:mobius-mersenne Lean source Lean source

Observation 6.30 (Skip-block invariant failure: two explicit fixtures). Two explicit witness states show "safe bracket + numerator monotone" is *not* an inductive invariant across skip-blocks in the dyadic-prefix reformulation: the transition $7/17 \to 15/119$, and a full-LCM-saturated state $12149/1240155$. A negative calibration fact for anyone attempting to induct directly on that pair of properties.

[Cert]scale:fixedcoord:seam-integer Lean source Lean source

Theorem 6.31 (The half-skip dichotomy via Erdős–Borwein irrationality). $1/2 \in \text{mersenneAchievementSet} \iff (\text{greedyMersenneSkippedSupport}(1/2))$ *is infinite. Forward direction uses irrationality of the Erdős–Borwein constant* (erdosBorweinMersenneConstant $= \text{mersenneTail}\, 0$, *unconditional*): *if only finitely many exponents were skipped, the constant would equal* $1/2$ *plus a finite rational, contradicting irrationality. The mechanism*

(*rational/finite-representable target, irrational full-series baseline ⇒ complement support infinite*) *is fully general and would restate for #249's own baseline constant if one exists.*

[Lean]scale:cofinalcoord:greedy-orbit LEAN SOURCE LEAN SOURCE

Theorem 6.32 (The nine-way half-membership classification hub)**.** $1/2 \in$ mersenneAchievementSet *is equivalent to* all *of:* (*1*) ¬SeamGreedyEventuallyRight*;* (*2*) SeamGreedyUnboundedTerminalFalse (*false successor terminal bits beyond every bound*)*;* (*3*) SeamGreedyUnboundedUpperOrMiddle (*U/M events beyond every bound*)*;* (*4*) SeamGreedyCofinalTerminalFalse (*one cofinal sequence of false terminals*)*;* (*5*) ∃ rows, SeamGreedyUnboundedSkippedRanksAlong rows*;* (*6*) *the greedy skipped support is infinite* (*Theorem 6.31*)*;* (*7*) *no last half-greedy skip exists* (*Theorem 6.40*)*. The central hub of the whole half-branch: every sufficient-condition theorem elsewhere in this catalogue is secretly proving one of these seven equivalent forms; a disproof of any one* (*a bounded/eventually-right certificate*) *closes the branch.*

[Lean]scale:cofinalcoord:seam-integer LEAN SOURCE LEAN SOURCE LEAN SOURCE

Lemma 6.33 (Eventually-right seam tail is analytically impossible for 1/2)**.** *If the seam eventually always extends "true"* (*right branch*) *from some row S on with a fixed lower prefix u, the resulting cofinite-support value stays strictly below* 1/2 (prefix_add_mersenneTail_lt_half_of_eventually_right)*; the matching alternative "upper competitor" word gives a strict excess* above 1/2 (half_lt_upper_competitor_of_eventually_right)*. An eventually-right tail can therefore never land exactly on* 1/2 *— it always undershoots or overshoots. The analytic engine behind the final-middle-cell exclusion* (*Theorem 6.44*)*.*

[Lean]scale:cofinalcoord:seam-integer LEAN SOURCE LEAN SOURCE

Lemma 6.34 (Mersenne weight/tail: strict superincreasingness)**.** mersenneWeight n = $1/(2^n - 1)$*;* mersenneTail n = $\sum_k'$ mersenneWeight$(n+k+1)$*;* mersenneTail n = mersenneWeight$(n+1)$ + mersenneTail$(n+1)$*; and for* $n > 0$, mersenneTail $n <$ mersenneWeight n (*strict superincreasingness*)*, with the two-sided envelope* mersenneTail $n \le 2 \cdot$ mersenneWeight$(n+1)$*. The whole superincreasing-series machinery downstream is built on the strict inequality. The proof pattern* (*remaining tail after rank n strictly smaller than weight n*) *is a fully generic superincreasing-series fact, directly relevant to #249's own* $\varphi(n)/2^n$ *series.*

[Lean]scale:uniformcoord:mobius-mersenne LEAN SOURCE LEAN SOURCE

Theorem 6.35 (Greedy survival criterion — the master theorem)**.** $x \in$ mersenneAchievementSet $\iff 0 \le x \wedge \forall n$, greedyMersenneRemainder $x\, n \le$ mersenneTail n*, where* greedyMersenneRemainder *is the canonical take-if-possible greedy recursion. Membership is exactly "every greedy residual survives* (*stays ≤ the full remaining tail*)*". The proof pattern* (*fatal-state absorption plus convergence-from-survival*) *is a completely generic superincreasing-series fact; only* mersenneWeight/ mersenneTail *are Mersenne-specific inputs. The canonical entry point for every half-membership sufficient-condition theorem in this catalogue.*

[Lean]scale:uniformcoord:greedy-orbit LEAN SOURCE

Lemma 6.36 (Rank-step trichotomy of the Dedekind cut)**.** *From a depth-d straddle of target t, exactly one of:* (*a*) *still straddles at* $d+1$ *without taking* $d+1$*;* (*b*) *straddles at* $d+1$

after taking $d{+}1$*;* (c) t *sits in the fatal gap* $(\text{value}(u) + \text{mersenneTail}(d{+}1), \text{value}(u) + \text{mersenneWeight}(d{+}1))$*. Specialized to* $t = 1/2$*, branches* (a) *and* (b) *are proved mutually exclusive* (half_step_forced)*, since both boundary equalities are impossible for* $1/2$ (*Lemma 6.38*)*. A pure consequence of superincreasingness plus* $\text{mersenneTail}\, n = \text{mersenneWeight}(n{+}1) + \text{mersenneTail}(n{+}1)$*.*

[Lean]scale:uniformcoord:dedekind-cut LEAN SOURCE LEAN SOURCE

Lemma 6.37 (Fatal-gap exclusion: kills every support, not just greedy)**.** *If* t *lies strictly inside the rank-*$(d{+}1)$ *gap* $(\text{mersenneSupportPrefix}\, A\, d{+}\text{mersenneTail}(d{+}1), \text{mersenneSupportPrefix}\, A\, d{+}\text{mersenneWeight}(d{+}1))$ *for a fixed prefix mass, then* $\text{positiveMersenneSupportValue}\, A \neq t$ *for* every *support* A *achieving that prefix, not merely the greedy orbit. Pure consequence of superincreasingness, no Mersenne-specific arithmetic beyond the weight function. The "kill certificate" primitive: any concrete numeric interval computation landing inside a fatal gap immediately falsifies half-membership via this lemma.*

[Lean]scale:uniformcoord:dedekind-cut LEAN SOURCE

Lemma 6.38 (Half endpoint kills: no finite word hits $1/2$ exactly)**.** *No finite positive-support word has value exactly* $1/2$ (positiveMersenneSupportValue_coe_finset_ne_half*, via the odd reduced denominator of the finite Erdős sum*)*. No finite word plus a complete tail equals* $1/2$ *either* (half_ne_coe_finset_add_mersenneTail*, via irrationality of the Erdős–Borwein constant, same mechanism as Theorem 6.31*)*. Discharges the two boundary branches of the rank-step trichotomy* (*Lemma 6.36*)*, forcing every macrostep to be a strict trichotomy with no equality case.*

[Lean]scale:uniformcoord:parity-irrationality LEAN SOURCE LEAN SOURCE

Lemma 6.39 (Straddle words are canonical: they agree with the greedy prefix)**.** *Every half-straddling word agrees with the actual canonical greedy word through the same depth:* $\forall n \le d,\ n > 0,\ (n \in u \leftrightarrow n \in \text{greedyMersenneSupport}(1/2))$*, by induction on* d *using the backward step* erase_top*. Straddle words are not an independent search tree — they are forced to* be *the greedy prefix. This is what lets the word-level fatal-gap exclusion* (*Lemma 6.37, valid for any support*) *be applied specifically to rule out the actual greedy orbit's half-representation.*

[Lean]scale:uniformcoord:greedy-orbit LEAN SOURCE LEAN SOURCE

Theorem 6.40 (Last-skip iff local fatality: half-membership as an infinite pointwise conjunction)**.** $\text{IsLastHalfGreedySkip}\, M \iff M \in \text{greedyMersenneSkippedSupport}(1/2) \wedge \text{GreedyMersenneFatalAt}(1/2)\, M$*. Equivalently: no last skip exists iff every actual skip survives, giving* $1/2 \in \text{mersenneAchievementSet} \iff \forall M \in \text{greedyMersenneSkippedSupport}(1/2), \text{greedyMersenneRemainder}(1/2)\, M \le \text{mersenneTail}\, M$*. Turns half-membership into an infinite conjunction over* actually skipped *ranks only — the reduction from a global fatal-state statement to a rank-local pointwise inequality is a generic greedy-recursion fact.*

[Lean]scale:uniformcoord:greedy-orbit LEAN SOURCE LEAN SOURCE

Lemma 6.41 (Seam upper-or-middle classification)**.** *For row* $s \ge 5$*: the empirical "false successor terminal bit" event is exactly the arithmetic disjunction* $\text{SeamGreedyUpperOrMiddleAt}\, s\, hs := (\text{seamAdjacentCut}\, s\, hs).\text{successorCarries} \vee (\neg\text{carries} \wedge 4 \cdot \text{remainder} + \text{gap} - \text{belowPulse} < \text{terminalWeight})$*. The empirical* $U/M/R$

label at each seam row is removable in favor of pure arithmetic on seamAdjacentCut*; the seam-arithmetic analogue of the rank-step trichotomy* (*Lemma 6.36*).

[Lean]scale:uniformcoord:seam-integer LEAN SOURCE

Lemma 6.42 (Largest-false-rank algebra: exact weight identity and branch propagation)**.** *Exact identity linking the largest-false-rank's lower vs. upper word weights:* $3 \cdot \text{lowerWeight} + \text{exactLateGap} = \text{upperWeight}$*. Successor propagation: the right branch* preserves *the same largest-false rank* d (IsLargestFalseRank.seamGreedyWord_succ_of_rightBranch)*; the upper/middle branches* reset *it to the new terminal rank* s (seamGreedyWord_succ_isLargestFalseRank_terminal_of_upperOrMiddle)*. The exact bookkeeping consumed inside the largest-skip-late induction* (*Theorem 6.6*).

[Lean]scale:uniformcoord:seam-integer LEAN SOURCE LEAN SOURCE

Theorem 6.43 (Critical dyadic-band index: a problem-agnostic quantifier collapse)**.** $\text{CriticalDyadicBandIndex}\, d\, E\, j$ *picks out the unique nearest dyadic power* $2^{d-j+1} \geq E$*. Fully abstract, pure* $\mathbb{N}$ *arithmetic, zero seam content:* $\text{DyadicBandEscape}\, d\, E \iff \exists j, \text{CriticalDyadicBandIndex}\, d\, E\, j \wedge E + 2(d+j) \leq 2^{d-j+1}$ *— the* $\forall j$ *band-avoidance condition collapses to checking exactly* one *critical index, reducing a check of* $d+1$ *separate inequalities to one nearest-boundary check per row. Specialized to the seam reset charge* (*with* $E \leq 2^{d+1}$ *automatically supplied by Theorem 6.9*)*, the reduced socket* SeamUpperResetCriticalBandEscape *is proved logically equivalent to Theorem 6.10's band-avoidance hypothesis.*

[Lean]scale:uniformcoord:dyadic-boundary LEAN SOURCE LEAN SOURCE

Theorem 6.44 (Final-middle-cell -3 excluded)**.** *The first of three exceptional final-middle integer cells* (*coordinate* $4 \cdot \text{remainder} - \text{belowPulse} - 4 = -3$) *cannot be the actual last transition before an all-right tail. Combines the analytic strict inequality of Lemma 6.33 with a two-step Möbius-carry recurrence forcing the centred carry to zero then strictly negative within two more coefficient rows, contradicting the nonnegativity of Theorem 6.11. One of three exceptional cells* $(-3, -2, -1)$*; the companion file records the analogous* binaryCoeffTail *inequalities for cells* -2 *and* -1 *— closing all three fully resolves the final-middle producer.*

[Lean]scale:uniformcoord:seam-integer LEAN SOURCE LEAN SOURCE

Lemma 6.45 (Skipped-endpoint trichotomy: frozen-margin sign classification)**.** *At every genuinely skipped rank* $s \geq 5$*, the frozen margin* $\text{greedyHalfFrozenMargin}(s-1)\, s$ *is in exactly one of three signed cells: negative* (*actual word = seam-greedy word, remainder* ≥ 1*, margin* $= -\text{remainder}$)*; zero* (*actual word = seam-greedy word, remainder* $= 0$)*; positive* (*actual word = the seam "above" word, margin* $=$ overshoot)*. The zero cell is provably safe; only the negative cell needs excluding, and it is exactly the case-split that Theorem 6.44's three exceptional cells live inside.*

[Lean]scale:uniformcoord:seam-integer LEAN SOURCE

Lemma 6.46 (Reverse carry-word overlap spacing: zero Mersenne content)**.** *A* ReverseCarryWord *is* $(\text{coeff}, \text{bit}, \text{carry} : \mathbb{N} \to \mathbb{Z})$ *with* $\text{bit}(m) + 2 \cdot \text{carry}(m) = \text{coeff}(m) + \text{carry}(m+1)$*. If two such words share a* $1/0$ *seam and then agree on both coefficients and bits for* length *further positions, their carry difference at the far end is exactly* 2^{length} *times an odd integer. Combined with a common Archimedean bound on both terminal carries, this forces*

$2^{\text{length}} \leq \text{bound}$. *Fully abstract: no Mersenne, base-4, or seam structure at all — pure integer-sequence recurrence, directly portable to #249's own digit-carry structure if given an analogous reverse-carry presentation.*

[Lean]scale:uniformcoord:binary-digit Lean source Lean source

Lemma 6.47 (Relation-invariant linear channels force a rank-one determinant collapse)**.** *If a finite family of linear channels* channel $: \iota \to V \to_{\mathbb{Q}} \mathbb{Q}$ *all vanish on* $\ker(\text{ev})$ *and* $\exists e$, ev $e = 1$, *then the evaluation matrix* $(i, j) \mapsto$ channel j (row i) *has rank* ≤ 1 *for* every *finite index type* ι *with* [Nontrivial ι] *— hence every square minor of size* ≥ 2 *has determinant zero, at all ranks simultaneously (previously checked only experimentally at ranks 2–4). Pure linear algebra over* $\mathbb{Q}$*-vector spaces, zero arithmetic content specific to Mersenne or #257; closes an entire family of "maybe there's a determinant obstruction at some higher rank" hopes in one shot.*

[Lean]scale:uniformcoord:linear-algebra Lean source

Theorem 6.48 (Two-thirds band: exact localisation of post-take skip-unsafety)**.** *Writing a greedy residual as* rem $= 1/R$*: a skip at rank* k *is dyadically safe iff* $R \geq 2^k$*. For a take-skip-take run with pre-take reciprocal* R*, take at* b $(q = 2^b - 1)$*, next take at* c $(m = 2^{c-1})$*, the exact localisation* PostTakeUnsafeAt $R\, q\, m \iff$ Band $R\, q\, m$ *holds, with band width exactly* $q^2/((q+m)(q+m-1))$*. Specializing to a single-skip run* $(m = 2q+2)$ *gives the two-thirds band, width exactly* $q^2/((3q+1)(3q+2)) < 1/9$*, pinning* $2q < 3R < 2q + 2/3$*. Two unconditional corollaries: an integral* R *is never unsafe; for* rem $= p/(2D)$ *in lowest terms with* p, D, q *odd, an unsafe run forces* $p \geq 7$ *(via a 2-adic parity argument on the band defect,* $4 \mid \Delta$*). Entirely general to any greedy expansion using the dyadic-safety test, not specific to Mersenne weights, but dyadic safety is sufficient, not proved necessary, for the true greedy process to survive:* $(p, D, b) = (17, 41, 3)$ *is an explicit odd-coprime example that* is *dyadically unsafe, and nothing here asserts the actual half-greedy orbit avoids the band.*

[Lean]scale:uniformcoord:rational-band Lean source Lean source Lean source

Theorem 6.49 (Sharp two-thirds fatal gap: the true-tail-safe criterion)**.** *For a skipped residual* $\rho = u/(2L)$ *at rank* k*, writing the skip margin as* $2^k u + a = 2L + u$ *(so* $0 < a \Leftrightarrow$ *skip): the dyadic safety test (Theorem 6.48) is* $u \leq a$*, but the* true *tail-based safety test only needs* $2u \leq 3a$*, proved via the explicit three-channel Lambert lower bound* mersenneTailLB3 $k = 1/2^k + 1/(3 \cdot 4^k) + 1/(7 \cdot 8^k) <$ mersenneTail k*. Unconditional corollaries: unit numerators* $(u = 1)$ *are always safe; fatality forces* $3a < 2u$*, hence* $u \geq 2$ *unconditionally and* $u \geq 3$ *once* u *is odd. Dyadic-safe* $\Rightarrow$ *sharp-safe strictly:* $(u, a) = (3, 2)$ *is sharp-safe but not dyadic-safe. One-third tighter than the dyadic test, discharged against the actual Mersenne tail, not just the three-channel lower bound. A separate project-memory note records matching "*$2u \leq 3a$*" language on the* unconditional $(\forall$ *infinite* $A)$ *#257 track; whether that is the same theorem or an independently proved analogue on this half-membership track is unverified — flag as* `reported_prior`*.*

[Lean]scale:uniformcoord:mobius-mersenne Lean source Lean source Lean source

Lemma 6.50 (Total gap-mass summability beyond any level)**.** mersenneGap $n :=$ mersenneWeight $n -$ mersenneTail $n > 0$*. The total gap mass strictly beyond level* N *is summable and explicitly bounded:* $\sum'_k$ mersenneGap $(N+k+1) \leq (2/9)(1/4)^N + (3/7)(1/8)^N \to 0$*, obtained by summing the pre-existing per-level asymptotic bound geometrically.* ***Scope caveat*** *(explicit in the source file header): bounds gap* mass *only — says nothing*

about which reals the greedy run reaches, does not certify membership or nonmembership of any point, and in particular says nothing about $1/2$ *directly.*

[Lean]scale:uniformcoord:mobius-mersenne LEAN SOURCE LEAN SOURCE

Lemma 6.51 (Half-divisor unit drop: smallest single-bit coefficient producer)**.** *At row* $2(N+1)$*, changing only exponent* $N+1$ *from* false *to* true *in a* HalfWord N *adds exactly the half-divisor incidence to* supportCoeff*:* supportCoeff(wordSupport(extend a true))$(2(N+1))$ = supportCoeff(wordSupport(extend a false))$(2(N$ 1*. The smallest possible "flip one bit, coefficient changes by exactly* 1*" producer. The mechanism (flipping a single word-bit changes a divisor-incidence sum by exactly the number of new divisor relations created) is a generic digit/divisor-incidence fact, restating for any base-*2 *divisor-coefficient bookkeeping, including #249's own divisor sums.*

[Lean]scale:uniformcoord:divisor-incidence LEAN SOURCE

Theorem 6.52 (Tempered-orbit rigidity: the shared #249/#257 trunk)**.** *For any nonnegative-integer coefficient sequence* $c : \mathbb{N} \to \mathbb{N}$ *with* $c(n) \le n$*: the binary coefficient series* $X_c = \sum'_{n\ge1} c(n)/2^n$ *is rational iff there exists a positive integer multiplier* v *and an integer orbit* $u : \mathbb{N} \to \mathbb{Z}$ *satisfying the exact carry recurrence* $u(N+1) = 2u(N) - v \cdot c(N+1)$ *together with the tempered (subexponential) growth condition* $u(N)/2^N \to 0$*. Every such tempered orbit is* rigid*:* $u(N) = v \cdot T_c(N)$ *exactly, where* $T_c(N) = \sum_{j\ge1} c(N+j)/2^j$ *is the scaled tail — there is at most one tempered orbit up to the multiplier* v*, and it equals the analytic tail exactly, not just asymptotically. Positivity of the orbit alone is deliberately* not *used as an equivalent criterion: a homogeneous* 2^N*-scaled perturbation can be added to any orbit without breaking the recurrence, so only temperedness is load-bearing. Explicitly asserted in-source as the shared binary carry trunk for #249 and #257:* $c = \tau$ *restricted to support (#257's* supportCoeff*) and* $c = \varphi$ *(#249) are both literal instances with no adaptation needed.*

[Lean]scale:uniformcoord:binary-digit LEAN SOURCE LEAN SOURCE

Lemma 6.53 (Finite symmetric-difference tail-transfer of irrationality)**.** *If* erdosSupportSeries b A *is irrational and* A, B *differ by a finite symmetric difference, then* erdosSupportSeries b B *is irrational too. Irrationality of a support series is invariant under changing finitely many support elements. Fully general; likely reusable for #249 tails as well. Feeds the eventually-periodic theorem (Theorem 6.21).*

[Lean]scale:uniformcoord:tail-transfer LEAN SOURCE

Lemma 6.54 (Dyadic-vs-Mersenne exact excess reformulation)**.** *The greedy-toward-*$1/2$ *decision at rank* $n+1$ *(take the next Mersenne weight* $1/(2^{n+1}-1)$ *or skip it) is exactly equivalent, with no real-analytic approximation, to an integer sign test on the excess numerator* nextDyadicExcessIntNumerator p n $L := 2^n p - L$ *for a displayed residual* $p/(2L)$*:* $p/(2L) \le 1/2^{n+1} \iff$ nextDyadicExcessIntNumerator p n $L \le 0$*. The forbidden "Mersenne–dyadic sliver" (skip is dyadically unsafe) is exactly* $0 < E \wedge 2E < p$*.*

[Lean]scale:uniformcoord:seam-integer LEAN SOURCE LEAN SOURCE

Lemma 6.55 (Denominator sandwich: odd-part survival, domain-neutral)**.** *Subtracting an odd reduced fraction* r/D *from a dyadic* $p/2^c$ *leaves a residual whose reduced denominator sandwiches the original:* $D \mid$ (residual).den $\mid 2^c \cdot D$*. Pure* `Rat.divInt` *denominator-reduction*

algebra, independent of the greedy/Mersenne setting; directly reusable anywhere a base-2 rational minus an odd-denominator rational needs its reduced denominator tracked, including #249's own carry bookkeeping.

[Lean]scale:uniformcoord:denominator-algebra LEAN SOURCE

Lemma 6.56 (Rational denominator survival: general-purpose fraction-reduction algebra)**.** *If $m \mid D$ and m is coprime to numerator a, then $m \mid (a/D \text{ as } \mathrm{Rat}).\mathrm{den}$. Scaled variant: if $C \mid D$ and C is coprime to a, then $C/\gcd(C,h) \mid ((h \cdot a)/D \text{ as } \mathrm{Rat}).\mathrm{den}$ — a divisor of the displayed denominator survives rational reduction up to exactly the part the scale multiplier can cancel. Zero domain content: the general-purpose "which prime power survives fraction reduction" lemma, feeding Theorem 6.24 directly, and equally applicable to any future #249 denominator-tracking argument.*

[Lean]scale:uniformcoord:denominator-algebra LEAN SOURCE LEAN SOURCE

Lemma 6.57 (Two-prime exact-valuation layer commutation)**.** *For an integer-valued support-coefficient function and two distinct primes $p \neq q$ with exponents $e, f > 0$: the iterated exact-level-difference operator* (primePowerLayer) *commutes across the two primes, and two nested layers extract exactly the iterated exact-valuation pullback of the support:* $\mathrm{mixedPrimePowerLayerTwo}\, p\, e\, q\, f\, (\mathrm{supportCoeffInt}\, A)\, n = \mathrm{supportCoeffInt}\, (\mathrm{pullback}_{q,f}(\mathrm{pullback}_{p,e} A))\, n$ *for n coprime to pq. Extends the single-prime layer theory to two primes; purely algebraic, supplies no transport from rationality through multiplicative decimation and asserts no bounded-Ω endpoint on its own.*

[Lean]scale:uniformcoord:p-adic LEAN SOURCE LEAN SOURCE

Definition 6.58 (The Mersenne achievement set and the support-series object)**.** $\mathrm{mersenneAchievementSet} := \{x \in \mathbb{R} \mid \exists A \subseteq \mathbb{N}, 0 \notin A \wedge x = \mathrm{positiveMersenneSupportValue}\, A\}$, where $\mathrm{positiveMersenneSupportValue}\, A = \sum'_k \mathrm{indicator}\, A\, \mathrm{mersenneWeight}(k+1)$ — the set of reals representable as a sum of a *subset* of positive-index Mersenne weights, ranging over all supports A, not only infinite ones. The general-b, general-support object used throughout the support-family route is $\mathrm{erdosSupportSeries}\, b\, A := \sum'_n \mathrm{indicator}\, A\, (1/(b^n - 1))\, n$, with $\mathrm{supportCoeff}\, A\, n := \#\{d \mid n : d \in A\}$ (the Dirichlet incidence $1_A * 1$, #257's analogue of τ); $A = \mathrm{Set.univ}$ recovers the full-support series of Theorem 6.14.

[Lean]scale:n/acoord:binary-digit LEAN SOURCE LEAN SOURCE LEAN SOURCE

Proposition 6.59 (Achievement-set topology: compact, closed, perfect, measure exactly one)**.** mersenneAchievementSet *is compact* (*continuous image of the binary-sequence Cantor space $\mathbb{N} \to \mathrm{Fin}\, 2$ under the product topology, via* positiveMersenneDigitValue)*, hence closed; it is also perfect, totally disconnected, and nowhere dense, with Lebesgue measure exactly* 1*. The compactness/closedness argument* (*binary coding → Cantor space → continuous image*) *is a fully generic technique for characterizing the achievement set of* any *absolutely convergent digit-weighted series, not specific to Mersenne denominators — directly reusable for a #249-style $\varphi(n)/2^n$ series after checking summability and superincreasingness. Closedness alone is what powers every "limit of a sequence of achieved points is achieved" argument in this catalogue* (*e.g. Theorem 6.2, Theorem 6.5*)*.*

[Lean]scale:n/acoord:binary-digit LEAN SOURCE LEAN SOURCE LEAN SOURCE

Theorem 6.60 (The master dichotomy: fatal-gap existence iff half-nonmembership)**.** ExistsFatalHalfGap := $\exists u, d$, (*bounds*) $\wedge$ value(u) + mersenneTail$(d{+}1)$ $<$ $1/2$ $\wedge$ $1/2$ $<$ value(u) + mersenneWeight$(d{+}1)$. ***Theorem:*** ExistsFatalHalfGap $\iff$ $1/2 \notin$ mersenneAchievementSet, *equivalently* $1/2 \in$ mersenneAchievementSet $\iff$ $\neg$ExistsFatalHalfGap — *the complete unconditional dichotomy: one of the two disjuncts always holds. The whole of the half-branch of #257 reduces to a search for one finite, checkable, numerically verifiable fatal-gap witness* (u, d) *— a completely finite existence statement, versus its negation, which by the straddle machinery (Lemma 6.36–6.39) is equivalent to membership. The shape of this dichotomy (finite fatal-gap certificate exists, or the closed achievement set contains the target) is the generic Dedekind-cut-in-a- superincreasing-series pattern, directly portable to #249 given its own analogues of Lemma 6.38 and Proposition 6.59.*

[Lean]scale:n/acoord:dedekind-cut Lean source Lean source

Definition 6.61 (Perturbed family and adjacent cut: the abstract seam gadget)**.** structure PerturbedFamily(α) where oldSum : $\alpha \to \mathbb{N}$; pulse : $\alpha \to \mathbb{N}$; gap, pulseCap : $\mathbb{N}$; with axioms gap_pos, pulse_le, oldSum_injective, a superincreasing-gap *separated* condition, and pulseCap $<$ $3 \cdot$ gap. Derived: newSum x := $4 \cdot$ oldSum x + pulse x; structure AdjacentCut(C) packages a "best subset $\le C$, best subset $> C$" adjacent pair with admissibility/maximality/strictness axioms. No Mersenne content whatsoever: α, oldSum, pulse are all free parameters — this is the general theory that seamPerturbedFamily/ seamAdjacentCut (used throughout the seam files) are specific instances of.

[Lean]scale:n/acoord:abstract-perturbed-greedy Lean source Lean source

Theorem 6.62 (Perturbed-family maximality and the three-branch recurrence)**.** successorCarries := $4 \cdot$ overshoot + abovePulse $\le$ gap; prefixChoice := if successorCarries then above else below *is provably maximal among all* x *with* newSum $x \le$ newCapacity (prefixChoice_maximal); nextRemainder_trichotomy *gives the exact three-branch (carry / middle / right) recurrence for the next remainder. Proved once and for all at the abstract* PerturbedFamily *level (Definition 6.61); the single most reusable abstract result in this catalogue — any future formalization of a different base-b greedy digit process needing a maximality-of-greedy-choice theorem can instantiate this structure directly instead of re-proving maximality from scratch.*

[Lean]scale:n/acoord:abstract-perturbed-greedy Lean source Lean source

6.3 Consumers

The BooleanMöbius lane (repo root `Erdos257PeriodNoncollapse/`, mirrored at `Erdos249257/` for the modules that have been made public) is where the reduction of Erdős #257 to a single arithmetic socket actually lives. Sixteen of its modules have never been published: `BooleanMobiusCofinalExactRows`, `BooleanMobiusCriticalCapacityCofinal`, `BooleanMobiusCriticalCapacityGeometric`, `BooleanMobiusExactRowCrossing`, `BooleanMobiusExactRowDichotomy`, `BooleanMobiusExactRowDoubling`, `BooleanMobiusExactRowRankTwo`, `BooleanMobiusExactRowSeed`, `BooleanMobiusExactTransition`, `BooleanMobiusGlobalRepair`, `BooleanMobiusGreedyReduction`, `BooleanMobiusLocalRepair`, `BooleanMobiusSkipRow`, `BooleanMobiusSkipRowCofinal`, `BooleanMobiusSkippedCoreCriticalCapacity`, `BooleanMobiusSkippedCoreExactRow`; plus

three sibling modules on the half-membership track, also never published: `HalfGapMass`, `HalfGreedyFatalGap`, `HalfUpperResetCriticalBand`. Every one of the nineteen is covered below. Consumers in this subsection are theorems whose conclusion is a genuine payoff ($(1/2 : \mathbb{R}) \in$ mersenneAchievementSet, or an infinite counterexample support) but whose hypothesis is an unsupplied predicate. Every such hypothesis is displayed, not buried in prose: that predicate is the actual target.

The load-bearing definitions

Definition 6.63 (257bm:d1 — `ExactLocalMersenneHalfRow`)**.** For $n : \mathbb{N}$:

$$\text{ExactLocalMersenneHalfRow}(n) \;:=\; \exists D : \text{Finset}\,\mathbb{N},\ (\forall d \in D,\ 2 \le d \le n) \ \wedge\ \text{localPrefixQuotient}(D, n) = 2^{n-1} - 1.$$

A finite Boolean support D bounded by n exactly hits the integer target $2^{n-1}-1$ under floor-division Mersenne quotients at scale n — the natural-number carry accounting is exact, not merely close.
[Lean] scale:n/a coord:mobius-mersenne LEAN SOURCE

Definition 6.64 (257bm:d2 — `localPrefixQuotient`, `localMersenneQuotient`, `localRepairInteger`)**.**

$$\text{localMersenneQuotient}(M, d) := \left\lfloor \frac{2^M}{2^d - 1} \right\rfloor, \qquad \text{localPrefixQuotient}(D, M) := \sum_{d \in D} \text{localMersenneQuotient}(M, d),$$

both natural-number floor division, and localRepairInteger(D, k) := the signed one-step carry defect of localPrefixQuotient at rank k (see 257bm:i8 below for the transition identity it satisfies). These are the three primitive accounting objects every other definition and theorem in this lane is built from.
[Lean] scale:n/a coord:mobius-mersenne LEAN SOURCE LEAN SOURCE LEAN SOURCE

Definition 6.65 (257bm:d3 — `localBinarySuffix`)**.**

$$\text{localBinarySuffix}(D, k, M) \;:=\; 2^{M-k} - \text{localPrefixQuotient}(D, M) - 1,$$

the leftover carry after subtracting D's integer quotient contributions from the binary expansion of 2^{-k} truncated at place M. Every capacity test in the lane is a bound on this one quantity.
[Lean] scale:n/a coord:binary-digit LEAN SOURCE

Definition 6.66 (257bm:d4 — `CofinalExactLocalMersenneHalfRows`, THE MASTER PREMISE)**.**

$$\text{CofinalExactLocalMersenneHalfRows} \;:=\; \forall N : \mathbb{N},\ \exists n : \mathbb{N},\ N \le n \ \wedge\ \text{ExactLocalMersenneHalfRow}(n).$$

Flag prominently: this requires *no compatibility whatsoever* between the witness supports D_n at different endpoints n. It is strictly weaker than a coherent global trajectory (257bm:c1 below) — mutually *incompatible* witnesses at each endpoint are entirely sufficient. This weakening is simultaneously the whole reason the lane is tractable and the exact place an attacker must be careful: a proof that looks like it supplies coherent rows is doing strictly more work than is needed, and a proof that looks like it refutes coherent rows says nothing about this predicate.
[Open] scale:cofinal coord:mobius-mersenne LEAN SOURCE

The two payoff consumers

Theorem 6.67 (257bm:c1 — `half_mem_mersenneAchievementSet_of_cofinalExactLocalRows`, THE PAYOFF THEOREM)**.**

$$\text{CofinalExactLocalMersenneHalfRows} \longrightarrow (1/2 : \mathbb{R}) \in \text{mersenneAchievementSet},$$

where $\text{mersenneAchievementSet} := \{x \mid \exists A : \text{Set}\,\mathbb{N},\ 0 \notin A \wedge x = \text{positiveMersenneSupportValue}(A)\}$. *Proof: for each N pick $n \geq N$ with exact row D_N; the real value y_N of D_N satisfies $|y_N - 1/2| \leq (n+1)/2^n \to 0$ (257bm:i9);* mersenneAchievementSet *is compact (continuous image of $2^{\mathbb{N}}$ under the digit-coding map), hence closed, so the limit $1/2$ is achieved by an actual point. This is the theorem 257bm:d4, and hence ultimately every producer feeding it, must reach.*
[Lean] scale:cofinal coord:other:topological-achievement-set LEAN SOURCE

Theorem 6.68 (257bm:c2 — `cofinalExactLocalMersenneHalfRows_of_positiveHalfGreedySkips`, `half_mem_mersenneAchievementSet_of_positiveHalfGreedySkips`)**.**

$$\text{CofinalPositiveHalfGreedySkips} \longrightarrow \text{CofinalExactLocalMersenneHalfRows} \longrightarrow (1/2 : \mathbb{R}) \in \text{mersenneAchievementSet},$$

where $\text{CofinalPositiveHalfGreedySkips} := \forall N,\ \exists c,\ \max(N,4) \leq c \wedge 0 < \text{greedyMersenneRemainderRat}(1/2, c-1) < \text{mersenneWeightRat}(c)$ *— the canonical rational half-greedy orbit itself has infinitely many positive skip events. A single skip at c already gives* $\text{ExactLocalMersenneHalfRow}(2c-2)$ *unconditionally* (`exactLocalMersenneHalfRow_of_positiveHalfGreedySkip`, *no capacity hypothesis at that step*)*. This is the SHORTEST known path to the disproof: cofinally many such skips close the chain with zero further combinatorial search.*
[Lean] scale:cofinal coord:greedy-orbit LEAN SOURCE LEAN SOURCE LEAN SOURCE

The coherent-repair consumer (superseded, kept as a fallback target)

Theorem 6.69 (257bm:c3 — `GlobalBooleanMobiusRepairFeasible`)**.** *A repair T (a bit-assignment* $\text{bit} : \mathbb{N} \to \mathbb{N} \to \text{Bool}$*, endpoint×rank→chosen bit, with* `bit_stable`*: once endpoint $2d$ is reached coordinate d is frozen forever) is feasible when*

$$\text{GlobalEndpointExponentialBound}(T) \wedge (\textit{word-value clause}) \wedge (\textit{word-capacity clause}) \wedge \forall n \geq 2,\ \text{localPrefixQuotient}(\text{globalRepairStageSupport}(T.\text{bit}, n), n) = 2^{n-1} - 1,$$

where the named open producer is

$$\text{GlobalEndpointExponentialBound}(T) := \forall n \geq 2,\ \text{let } D := \text{globalRepairLowerSupport}(T.\text{bit}, n);\ 2^{\text{endpointDivisorContribution}(D,n)-1} - 1 \leq \text{localBinarySuffix}(D, 1, n-1).$$

This is the STRONGER, harder consumer — a genuinely coherent (frozen-diagonal) repair, unlike 257bm:d4's cofinal-exact-row supply which needs no coherence at all. Kept as a fallback target only; 257bm:c1/c2 supersede it.
[Open] scale:uniform coord:binary-digit LEAN SOURCE LEAN SOURCE LEAN SOURCE

The critical-quotient-supply consumer chain

Theorem 6.70 (257bm:c4 — `SkippedCoreCriticalQuotientSupply`, THE REMAINING OPEN SOCKET)**.**

SkippedCoreCriticalQuotientSupply := $\forall D\ c,\ 4 \le c \to (\forall d \in D,\ 2 \le d < c) \to \mathrm{value}(D) < \frac{1}{2} \to \left(\frac{1}{2} - \mathrm{value}(D) < \mathrm{mersenneWeightRat}(c)\right) \to 2^{(2c-2)-1} \le \mathrm{localPrefixQuotient}(\mathrm{insert}\ c\ D,\ 2c-2)$.

Whenever a below-half core is genuinely crossed by rank c, adjoining c already reaches the integral half-target. By 257bm:i5 (the sharp capacity iff) this is EXACTLY the sharp $(c-2)$-bit capacity test, universally quantified over all crossing cores. If proved, this Prop closes the whole cofinal-exact-row route and hence disproves Erdős #257.
[Open] scale:uniform coord:binary-digit LEAN SOURCE

Theorem 6.71 (257bm:c5 — `cofinalExactLocalMersenneHalfRows_of_criticalQuotientSupply`)**.** *An already-formalised induction from the endpoint-six seed (257bm:i7) consumes 257bm:c4 at every step: each* `ProtectedExactLocalMersenneRow` *either doubles below half (unconditional) or recycles at its first crossing rank $e >$ cutoff, giving endpoint $2e-2 >$ previous endpoint;* `protection` *(endpoint $< 2\cdot$cutoff, new ranks $>$ cutoff) is exactly what converts the non-growing recycle endpoint of the bare dichotomy (257bm:k1) into strict progress. The below-half branch never fires twice from the seed arithmetic, so the supply is needed at essentially every step. Ends literally at* CofinalExactLocalMersenneHalfRows*.*
[Lean] scale:uniform coord:binary-digit LEAN SOURCE LEAN SOURCE LEAN SOURCE

Theorem 6.72 (257bm:c6 — `HalfGreedySkippedCriticalQuotientSupply`, the already-reduced canonical form)**.**

HalfGreedySkippedCriticalQuotientSupply := $\forall c \ge 4$ *skipped by the rational half-greedy orbit*, $2^{(2c-2)-1} \le \mathrm{localPrefixQuotient}(\mathrm{insert}\ c\ (\mathrm{halfGreedyPrefixSupport}(c-1)),\ 2c-2)$.

The sharp inequality at one crossing rank per step, in the canonical form the induction of 257bm:c5 actually consumes — this is 257bm:c4 with the universal quantifier already collapsed onto the one orbit by 257bm:i3 (the wall).
[Open] scale:uniform coord:greedy-orbit LEAN SOURCE

Theorem 6.73 (257bm:c6a — `halfGreedy_precriticalSuffix_lt_of_next_skip`, unconditional two-consecutive-skips producer)**.** *UNCONDITIONAL: if $c \ge 6$ and BOTH rank c and rank $c+1$ are skipped by the rational half-greedy orbit, then*

$$\mathrm{localBinarySuffix}(\mathrm{halfGreedyPrefixSupport}(c-1),\ 1,\ 2c-3) < 2^{c-3},$$

which doubles into sharp capacity at c and hence an exact row at $2c-2$ with strict-upper separation. No supply hypothesis anywhere. The residual open set is exactly the skip-then-take rows: `halfGreedySkippedPrecriticalSuffixSupply_iff_preTake` *proves the whole socket equivalent to* `HalfGreedyPreTakePrecriticalSuffixSupply`*, with $c = 4$ and $c = 5$ already discharged by* `decide`*. What would close it: cofinally many $c \ge 6$ at which two consecutive ranks $c, c+1$ are both skipped by the rational half-greedy orbit — or, dually, the pre-take residual at cofinally many skipped-then-taken ranks.*
[Lean] scale:uniform coord:greedy-orbit LEAN SOURCE LEAN SOURCE LEAN SOURCE

Theorem 6.74 (257bm:c6b — `halfGreedy_precriticalSuffix_lt_of_future_skip_after_takenBlock`, gap-uniform generalisation)**.** *UNCONDITIONAL and uniform in BOTH parameters, strictly generalising 257bm:c6a: skip at c, takes at $c+1,\dots,c+t-1$, skip at $c+t$, with $0 < t \le c-3$ and $c-2 \le 2^{c-t-3}$ $\implies$ the precritical suffix bound at c $\Rightarrow$ sharp capacity $\Rightarrow$ exact row at $2c-2$. Handles any skip gap $t \lesssim c-3-\log_2(c-2)$, not just $t=1$. Fully uniform proof — no finite table, every constant explicit ($|D| \le c-2$, dyadic room 2^{c-t-3}, gap lemma `dyadic_lt_forcedBlock_of_tail_lt`) — so any orbit-level SKIP-GAP BOUND immediately yields cofinal exact rows with no further work. What would close it: for cofinally many skipped ranks c of the rational half-greedy orbit, the next skipped rank lies at most $c + (c-3-\lceil\log_2(c-2)\rceil)$. Empirically the skips have roughly constant density (certified computationally to row 200,000 on the sqrt-escape track), so this is a slack requirement never previously isolated as a target.*
[Lean] scale:uniform coord:greedy-orbit LEAN SOURCE LEAN SOURCE LEAN SOURCE

Two unconditional fill consumers (`BooleanMobiusSkipRow`)

Theorem 6.75 (257bm:c7 — `exists_exactRowStrictUpperFill_of_skippedCoreSharpCapacity`)**.** *For any $c \ge 4$ and any finite $D \subseteq [2,c)$ with $\mathrm{value}(D) < 1/2$ and the sharp capacity*

$$\mathrm{localBinarySuffix}(D,1,2c-2) < 2^{c-2},$$

there is an exact row E at endpoint $2c-2$ with $D \subseteq E$ and every new rank strictly above c. Unconditional, no deficit/crossing hypothesis, no reference to the greedy orbit at all — this consumer is strictly more general than every route that feeds it: the crossing-core deficit hypothesis used in 257bm:c4/c6 is a self-imposed restriction on the corpus's own attack, not a requirement of this theorem.
[Lean] scale:uniform coord:binary-digit LEAN SOURCE

Theorem 6.76 (257bm:c8 — `exactLocalMersenneHalfRow_two_mul_sub_two_of_skippedCoreSharpCapacity`)**.** *Bare-Prop corollary of 257bm:c7: the same hypotheses give* $\mathrm{ExactLocalMersenneHalfRow}(2c-2)$ *directly, dropping the support-extension data. The consumer form actually chained into 257bm:c1/c5 when the strict-upper separation is not needed downstream.*
[Lean] scale:uniform coord:binary-digit LEAN SOURCE

Proposition 6.77 (257bm:c9 — `localMersenneQuotient_eq_two_pow_sub_of_half_lt`)**.** *In the pure upper window used by 257bm:c7's fill, the Mersenne quotient is exactly* $\mathrm{localMersenneQuotient}(M,d) = 2^{M-d}$ *once the value condition places d strictly above half the depth. This one identity is what lets the fill encode the missing residue as a literal Boolean word rather than a recursive repair — the arithmetic reason 257bm:c7 is possible at all.*
[Lean] scale:uniform coord:binary-digit LEAN SOURCE

Theorem 6.78 (257bm:c10 — `exactLocalMersenneHalfRow_two_mul_sub_two_of_skippedCore`)**.** *Companion unconditional consumer at* `BooleanMobiusSkipRow.lean`*: given a first-crossing witness produced by* `exists_skippedCoreExactRow_of_value_above` / `exactLocalMersenneHalfRow_of_first_localMersenne_crossing`

(`BooleanMobiusExactRowCrossing.lean:155, 191`), *builds the exact row at* $2c - 2$ *from the skipped core without any capacity input. This is the base construction 257bm:c7/c8 sharpen with the strict-upper separation clause.*
[Lean] scale:bounded coord:binary-digit LEAN SOURCE

The dyadic-band quantifier-collapse consumer (`HalfUpperResetCriticalBand`)

Theorem 6.79 (257bm:c11 — `DyadicBandEscape` ⇔ `CriticalDyadicBandIndex`, PROBLEM-AGNOSTIC quantifier collapse)**.** *Pure* $\mathbb{N}$*-arithmetic, zero seam or Mersenne content: for* d, E, j *with* CriticalDyadicBandIndex(d, E, j) *the unique nearest dyadic power* $2^{d-j+1} \geq E$,

$$\text{DyadicBandEscape}(d, E) \iff \exists j,\ \text{CriticalDyadicBandIndex}(d, E, j) \wedge E + 2(d + j) \leq 2^{d-j+1}.$$

The $\forall j \in [0, d]$ *band-avoidance condition* ($d + 1$ *separate inequalities*) *collapses to checking exactly ONE nearest-boundary index per row. Specialised to the concrete seam reset charge,* SeamUpperResetCriticalBandEscape *is proved logically EQUIVALENT to* SeamUpperResetDyadicBandEscape (*the F5-style band-avoidance producer*)*, hence also consumes into* half_mem_mersenneAchievementSet_of_upperResetCriticalBandEscape.
[Lean] scale:uniform coord:other:dyadic-boundary LEAN SOURCE LEAN SOURCE LEAN SOURCE LEAN SOURCE

Proposition 6.80 (257bm:c12 — prerequisite `seamUpperResetCharge_le`)**.** *The reduction 257bm:c11 needs* $E \leq 2^{d+1}$ *as a side condition; this is supplied automatically for the concrete seam reset charge by the two-sided dyadic bound of* `HalfCylinderMiddleCarryLowerBound.lean` (*the* min(remainder, overshoot) $\leq 2^s$ *global invariant*)*, so 257bm:c11's reduction is never blocked on this side condition in practice.*
[Lean] scale:uniform coord:other:dyadic-boundary LEAN SOURCE

The frozen-margin consumer family (`HalfCylinderFullShellSeamBridge`, half-membership track)

Three graded sockets, strongest hypothesis to weakest, all funnel into the same endpoint half_mem_mersenneAchievementSet_of_governedFrozenMarginProducer.

Theorem 6.81 (257bm:c13 — (i) `HalfGreedySkippedFullShellNonnegative`)**.**

$$\forall \text{ skipped rank } n, \quad 0 \leq \text{greedyHalfFrozenMargin}(n - 1, n).$$

Direct, seam-free. The plainest sufficient socket in the family.
[Open] scale:uniform coord:other:frozen-margin LEAN SOURCE LEAN SOURCE

Theorem 6.82 (257bm:c14 — (ii) `HalfGreedySkippedSeamAlignmentZero`)**.**

$$\forall \text{ skipped rank whose actual word equals the seam-greedy word}, \quad \text{seam remainder} = 0,$$

proved EQUIVALENT to (*i*) *via* `skippedSeamAlignmentZero_iff_skippedFullShellNonnegative`.
[Open] scale:uniform coord:other:frozen-margin LEAN SOURCE LEAN SOURCE

Theorem 6.83 (257bm:c15 — (iii) `HalfGreedySkippedSeamEscape`)**.**

$$\forall \textit{ skipped rank}, \quad \text{halfStripBound}(2n) < \text{seamIntegerGreedyRemainder}(n),$$

a strictly SUFFICIENT (not equivalent) condition for (i). The layered graded-socket pattern (equivalent core + one strictly-sufficient stronger variant) recurs across this whole lane as a proof-engineering idiom.
[Open] scale:uniform coord:other:frozen-margin LEAN SOURCE LEAN SOURCE

Two square-root-scale consumers (half-membership track)

Theorem 6.84 (257rig:c16 — `infinite_support_half_of_mobiusCenteredHalfCarry_sqrtBound`, `greedy_half_infinite_of_mobiusCenteredHalfCarry_sqrtBound`)**.** *If*

$$\forall N, \quad \text{mobiusCenteredHalfCarry}(\text{greedySupport}(1/2), N) \;\le\; 2\sqrt{N} + 4,$$

then the greedy skipped support is infinite and $\text{erdosSupportSeries}(2, \text{greedyMersenneSupport}(1/2)) = 1/2$. *Nonnegativity of the left side is unconditional (see 257bm:i-carry below); ONLY the upper $2\sqrt{N} + 4$ bound remains open. Distinct from 257bm:c11's dyadic-band route and from 257bm:c7's largest-skip socket — the corpus records this as a genuinely independent square-root-scale frontier.*
[Open] scale:cofinal coord:other:sqrt-carry-growth LEAN SOURCE LEAN SOURCE

Theorem 6.85 (257rig:c17 — `exists_infinite_support_half_of_cofinalTerminalOnlyStrip`)**.**

$\forall N\ \exists$ *depth* $M \ge N$, $\exists$ *normalized word* a, $\quad |\text{integerHalfCarry}(\text{wordSupport}(a), M-1)| \le \text{halfStripBound}(M) \implies \exists A$ *infinite*, $\text{erdosSupportSeries}(2, A) = 1/2$.

Same architecture as 257bm:d4: cofinal, no coherence, mutually incompatible witnesses allowed. Trades off against 257bm:c7 in the OPPOSITE direction — it tolerates any support inside the depth-M window but demands the carry inside a $\sim 2\sqrt{M}$ strip, whereas 257bm:c7 tolerates a carry up to 2^{c-2} (exponentially looser) but demands the support be confined to ranks $< c$. Neither socket dominates the other; this has never been compared to 257bm:c7 on band width in any prior bank.
[Open] scale:cofinal coord:other:sqrt-carry-growth LEAN SOURCE LEAN SOURCE

Theorem 6.86 (257rig:c18 — `greedy_half_infinite_of_cofinalStripReturn`)**.** *Cofinal RETURNS of the actual greedy carry to the square-root strip give an infinite support with value exactly 1/2, an alternate coordinate to 257rig:c16's pointwise bound. The exact bridge between the two is already landed:* `halfGreedy_precriticalSuffix_lt_iff_futureSkipCoverage` *shows sharp capacity* $\Leftrightarrow \text{mobiusCenteredHalfCarry}(G, 2c-4) \le \text{futureSkipCapacity}(G, c, c-3)$, *so transport between the frozen-prefix carry (needed by 257bm:c7's fill) and the live-orbit carry (bounded here) costs exactly the future-skip capacity term.*
[Open] scale:cofinal coord:other:sqrt-carry-growth LEAN SOURCE LEAN SOURCE LEAN SOURCE

The non-effective horizon consumer

Theorem 6.87 (257bm:c19 — `exists_greedyHalfFrozenMargin_nonneg_iff_excess_neg`)**.** *At every dyadically safe rank $k > 0$, the frozen margin* greedyHalfFrozenMargin(k, J) *is nonnegative for SOME horizon J (monotone up in J). The proof is a bare limit argument* (`tendsto_finiteCoeffWindow_atTop` *then* `(tendsto_order.1 hlim).1 ... .exists`)*, so J is produced non-effectively — the exact-row route needs the specific horizon $J = c - 3$ (consumed at* `halfGreedy_precriticalSuffix_lt_iff_frozenMargin_nonneg`, `BooleanMobiusCriticalCapacityCofinal.lean:784`*). Everything to make it effective is already on disk:* binaryCoeffTail_eq_finiteCoeffWindow_add_shiftedTail *gives* gap $=$ tail$(k+1+J)/2^J$*, and 257bm:i2 (the tail growth bound) gives* tail$(m) \le m + 2$*. What would close it:* margin$(k, J) \ge 0$ *as soon as* $(k + J + 3)/2^J <$ |halfGreedyNextDyadicExcessNumerator(k)|/halfGreedyPrefixDenominator(k) *— pure bookkeeping over existing lemmas, no new idea.*
[Lean] scale:uniform coord:other:frozen-margin LEAN SOURCE LEAN SOURCE

The equivalence `CofinalPositiveHalfGreedySkips` reduces to

Theorem 6.88 (257bm:c20 — `half_mem_mersenneAchievementSet_iff_greedySkippedSupport_infinite` (A5/A6, public))**.**

$(1/2 : \mathbb{R}) \in$ mersenneAchievementSet $\iff$ (greedyMersenneSkippedSupport$(1/2)$).Infinite.

The positivity conjunct of 257bm:c2's hypothesis $(0 <$ greedyMersenneRemainderRat$(1/2, c-1))$ *is unconditionally dischargeable (odd-denominator parity, 257bm:i9-style), so* CofinalPositiveHalfGreedySkips $\Leftrightarrow$ (greedyMersenneSkippedSupport$(1/2)$).Infinite, *which this theorem proves equivalent to* $1/2 \in$ mersenneAchievementSet *directly — the consumer 257bm:c2 chains into is, once unwound, this exact equivalence. See 257bm:k5 for why this makes 257bm:c2 a trap, not a shortcut.*
[Lean] scale:cofinal coord:other:topological-achievement-set LEAN SOURCE LEAN SOURCE

6.4 Invariants

Invariants are unconditional identities, monotonicity facts, and quantitative bounds that hold regardless of which branch of the lane's case splits fires. They are the load-bearing algebra every producer and consumer above is built from.

Endpoint transition algebra (`BooleanMobiusExactTransition`)

Theorem 6.89 (257bm:i1a — `localMersenneQuotient_endpoint_succ`)**.** *Exact one-step binary transition formula for* localMersenneQuotient *from endpoint n to $n+1$: each rank's quotient contribution either stays or doubles-plus-carries depending on divisibility. Unconditional identity, no hypothesis beyond the endpoint arithmetic itself.*
[Lean] scale:uniform coord:binary-digit LEAN SOURCE

Theorem 6.90 (257bm:i1b — `localPrefixQuotient_succ`)**.** *The rank-wise identity of 257bm:i1a summed over a fixed support D: exact recurrence for* localPrefixQuotient$(D, n) \to$ localPrefixQuotient$(D, n+1)$.

[Lean] scale:uniform coord:binary-digit LEAN SOURCE

Theorem 6.91 (257bm:i1c — `localEndpointDefect_succ`, the $H = 2A+1-S$ recurrence)**.** *The signed defect version gives the literal transition $H = 2A + 1 - S$ referenced throughout the module docstrings, where A is the suffix carry and $S =$* endpointDivisorContribution *(how many selected ranks divide the new endpoint). This is the SAME recurrence shape as the generic tempered-orbit recurrence $u(N+1) = 2u(N) - v \cdot c(N+1)$ from the public* `GenericTailOrbitRigidity.lean` *T7 criterion (257bm:i-t7 below), specialised to Mersenne local repair. Whether* localRepairInteger *literally instantiates a tempered orbit for $c :=$* supportCoeff *A is UNVERIFIED but high-value to check: it would let the T7 criterion feed rationality of* erdosSupportSeries 2 *A directly, a different route into the achievement-set argument than the cofinal-exact-row route of this catalogue.*
[Lean] scale:uniform coord:binary-digit LEAN SOURCE LEAN SOURCE

Capacity identities (`BooleanMobiusSkippedCoreCriticalCapacity`, `BooleanMobiusSkippedCoreExactRow`)

Definition 6.92 (257bm:i-cap — the sharp capacity test)**.** For $c \geq 4$, $\forall d \in D,\ 2 \leq d < c$, value$(D) < 1/2$:

$$\text{localBinarySuffix}(D, 1, 2c-2) < 2^{c-2}$$

is the sharp skipped-core capacity — the residual carry after inserting rank c fits inside exactly $c-2$ bits, making a below-half core repairable into an exact row at endpoint $2c-2$.
[Lean] scale:n/a coord:binary-digit LEAN SOURCE

Theorem 6.93 (257bm:i5 — `localBinarySuffix_two_mul_sub_two_lt_criticalCapacity_iff`)**.** *Under the hypotheses of 257bm:i-cap:*

$$\text{localBinarySuffix}(D, 1, 2c-2) < 2^{c-2} \iff 2^{(2c-2)-1} \leq \text{localPrefixQuotient}(\text{insert } c\ D,\ 2c-2).$$

The sharp capacity bound is equivalent to "adjoining rank c already reaches the integer half-target" — capacity and integral-crossing are the same fact viewed two ways. Rewrites the analytic capacity question into a purely combinatorial (integer-quotient) one.
[Lean] scale:uniform coord:binary-digit LEAN SOURCE

Theorem 6.94 (257bm:i6 — `localBinarySuffix_two_mul_sub_two_lt_upperHalfCapacity`, loose (one-bit) capacity, `BooleanMobiusSkippedCoreExactRow`)**.** *Unconditionally, for every $c \geq 4$ and every below-half core $D \subseteq [2, c)$ with deficit $<$* mersenneWeightRat(c)*:*

$$\text{localBinarySuffix}(D, 1, 2c-2) < 2^{c-1}.$$

Uniform in c, no table, no case split — but ONE BIT looser than the sharp 257bm:i5 test the fill (257bm:c7) actually needs. Reading the proof, the loss is purely additive, not multiplicative: it derives $A < 2^{c-2} + |D|$ and then discards $|D| \leq c - 2 \leq 2^{c-2}$. The true unproved statement is only that A avoids the LINEAR-WIDTH band $[2^{c-2}, 2^{c-2} + c - 3]$ — width $c-2$ inside a

range of size 2^{c-2}. This is the same band shape as 257bm:c11's dyadic-band route (F5/F6), and EXPONENTIALLY weaker than the $\sqrt{}$-scale target of 257rig:c16 (width c versus width $2^{(r+5)/2}$ on the unconditional track) — the two have never been directly compared on band width elsewhere in this corpus.
[Lean] scale:uniform coord:binary-digit LEAN SOURCE LEAN SOURCE

Doubling-branch invariants (`BooleanMobiusExactRowRankTwo`)

Theorem 6.95 (257bm:i7 — `localBinarySuffix_two_mul_sub_one_lt_upperWindow_of_exact_below`). *For $n \geq 6$, D bounded $[2, n]$, $2 \in D$, exact quotient at n, value $< 1/2$:*

$$\text{localBinarySuffix}(D, 1, 2n-1) < 2^n.$$

The missing quotient at $2n-1$ fits the pure upper window $\{n+1, \ldots, 2n-1\}$ — enough carry slack to encode the missing digits as a literal Boolean word with no recursion into already-used ranks. Uses `one_third_le_localMersenneFraction_two` *(rank-2 residue always $\geq 1/3$) and* `three_mul_sub_two_lt_two_pow_pred` *($3(n-2) < 2^{n-1}$, the elementary slack estimate).*
[Lean] scale:uniform coord:binary-digit LEAN SOURCE

Proposition 6.96 (257bm:i-rank2 — `two_mem_of_exact_localMersenneQuotient`). *The side condition $2 \in D$ used throughout the doubling branch is FREE: it holds uniformly for $n \geq 3$ at any exact quotient, so it never needs separate discharge downstream.*
[Lean] scale:uniform coord:binary-digit LEAN SOURCE

Base-case fixture (`BooleanMobiusExactRowSeed`)

Definition 6.97 (257bm:i-seed — `exactRowSixSupport`, the smallest exact row). exactRowSixSupport := $\{2, 3, 6\}$ satisfies

$$\text{localPrefixQuotient}(\{2,3,6\}, 6) = 2^5 - 1 = 31, \qquad \text{localMersennePrefixValue}(\{2,3,6\}) < 1/2,$$

hence ExactLocalMersenneHalfRow(6). Closed numeric fixture, `norm_num`-checked. Every $6 \leq n$ hypothesis in the doubling/dichotomy files traces back to this witness being the first valid seed.
[Cert] scale:fixed coord:mobius-mersenne LEAN SOURCE LEAN SOURCE

Global-repair invariants (`BooleanMobiusGlobalRepair`)

Proposition 6.98 (257bm:i9 — `abs_localMersennePrefixValue_sub_half_le`). *An exact row at endpoint n has real value within $O(n/2^n)$ of $1/2$: the quantitative bound $|y_n - 1/2| \leq (n+1)/2^n$ consumed by 257bm:c1 to turn a sequence of exact rows tending to depth infinity into a sequence of values tending to $1/2$.*
[Lean] scale:uniform coord:mobius-mersenne LEAN SOURCE

Proposition 6.99 (257bm:i10 — `finite_boolSupport_ne_half`). *Cites the odd-denominator fact (`finiteErdosSum_den_odd`, 257bm:k11) to conclude directly: no finite Boolean support ever equals $1/2$ exactly. The corollary consumed at the very end of 257bm:c1's chain to certify the achieved witness is infinite.*
[Cited] scale:uniform coord:p-adic LEAN SOURCE

Greedy-window uniqueness (`BooleanMobiusGreedyReduction`)

Theorem 6.100 (257bm:i11a — `remainder_lt_gap_iff_eq_integerGreedyBits`)**.** *Under gap dominance (the next digit's weight strictly exceeds the sum of all lower digits' slack), a Boolean word matches the unique greedy-bit word IFF its remainder is below the separation gap. Whenever the separation gap dominates there is exactly ONE admissible Boolean word achieving a given target defect — the greedy word. The "superincreasing weight sequence + below-gap remainder ⇒ unique greedy word" technique is standard and fully general, independent of the Mersenne instantiation.*
[Lean] scale:uniform coord:greedy-orbit LEAN SOURCE

Proposition 6.101 (257bm:i11b — gap-dominance facts)**.** `localMersenneWeightsFrom_gapDominates` *and its even/odd endpoint specializations* `_gapDominates_even` / `_gapDominates_odd`*, established from* `three_mul_sub_two_lt_two_pow_pred`*-style elementary slack bounds — the concrete hypotheses 257bm:i11a needs, discharged unconditionally at every relevant endpoint parity.*
[Lean] scale:uniform coord:greedy-orbit LEAN SOURCE LEAN SOURCE LEAN SOURCE

Proposition 6.102 (257bm:i11c — `localMersenneHalfTarget_lower_word_eq_greedy_and_remainder_eq`)**.** *Consumer form of 257bm:i11a specialised to the half-target: the lower word of any exact half-target construction is forced to equal the greedy word, remainder identity included. This is the general uniqueness principle instantiated by 257rig:i3 (the wall).*
[Lean] scale:uniform coord:greedy-orbit LEAN SOURCE

Geometric normal form (`BooleanMobiusCriticalCapacityGeometric`)

Theorem 6.103 (257bm:i12 — division-free re-coordinatisation of the sharp capacity iff)**.** localMersenneQuotient_eq_geometric*: Euclidean division* $2^M/(2^d-1)$ *equals a shifted finite geometric sum (division-free "descending multiples" normal form); the iff-forms* `localBinarySuffix_two_mul_sub_two_lt_criticalCapacity_iff_geometric` *and* `_iff_geometricCore` *restate 257bm:i5 entirely in this coordinate, with the "crossing coin" contribution* localMersenneGeometricQuotient_critical_self *visible as an explicit finite sum rather than a floor division. Same content as 257bm:i5, alternate coordinate; the division-free re-expression technique for* $N/(b^d-1)$*-type quotients is fully general to any base b.*
[Lean] scale:uniform coord:other:geometric-series LEAN SOURCE LEAN SOURCE LEAN SOURCE

Growth-bound invariants for the generic layer

Theorem 6.104 (257bm:i2 — `binaryCoeffTail_le`, public)**.** *For* $c : \mathbb{N} \to \mathbb{N}$ *with* $\forall n,\ c\,n \le n$*:* binaryCoeffTail$(c,N) \le N+2$ *for all* N *— ANY linearly-bounded natural coefficient sequence has scaled tail* $O(N)$*, hence* $o(2^N)$*. Makes the tempered-orbit criterion T7 (257bm:i-t7) non-vacuous, and confirms* $c := \varphi$ *is safely inside its domain, since* $\varphi(n) \le n$ *is standard.*
[Cited] scale:uniform coord:binary-digit LEAN SOURCE

Theorem 6.105 (257bm:i-t7 — `binaryCoeffSeries_rational_iff_exists_temperedBinaryOrbit` (T7), public, THE cross-problem bridge)**.** *For* $c : \mathbb{N} \to \mathbb{N}$ *with* $\forall n,\ c\,n \le n$*:*

$$\text{HasRationalValue}(\text{binaryCoeffSeries}\, c) \iff \exists v > 0,\ \exists u : \mathbb{N} \to \mathbb{Z},\ \text{IsTemperedBinaryOrbit}(c, v, u),$$

where $\text{binaryCoeffSeries}\, c := \sum'_n c(n{+}1)/2^{n+1}$ *and* $\text{IsTemperedBinaryOrbit}(c, v, u) := (\forall N,\ u(N{+}1) = 2u(N) - v \cdot c(N{+}1)) \wedge u(N)/2^N \to 0$. *ONLY the linear growth bound* $c\, n \le n$ *is required — no divisor, Mersenne, or #257-specific structure at all. Fully problem-agnostic; instantiating* $c := \varphi$ *gives a candidate #249 reduction to the SAME criterion used for #257, modulo an index-shift caveat that remains unverified.*
[Lean] scale:uniform coord:binary-digit Lean source Lean source Lean source

Theorem 6.106 (257bm:i-mob — `moebius_mul_supportCoeffAF`, public, Möbius layer)**.** $\text{ArithmeticFunction.moebius} * \text{supportCoeffAF}(A) = \text{positiveSupportBitAF}(A)$ *for ANY* $A : \text{Set}\,\mathbb{N}$ *— exact Möbius recovery of a 0/1 support indicator from its divisor-count coefficient. Fully base-free; only downstream specialisations* $(\text{erdosSupportSeries}\, 2\, A)$ *fix the base to 2.*
[Lean] scale:uniform coord:other:mobius-inversion Lean source Lean source Lean source

Theorem 6.107 (257bm:i-bridge — `erdosSupportSeries_two_eq_binaryCoeffSeries`, public)**.** $\text{erdosSupportSeries}(2, A) = \text{binaryCoeffSeries}(\text{supportCoeff}\, A)$ *— identifies the #257 support series at base 2 exactly with T7's generic machinery, closing the loop that makes T7 the actual engine behind #257's rationality criterion too.*
[Lean] scale:uniform coord:mobius-mersenne Lean source Lean source

Rigidity corollaries directly constraining a rational #257 support

These five invariants (`RationalSupportCarrySkeleton.lean`, `SublogDivisorCoverage.lean`, `MaximalOmegaLayer.lean`) are the corollaries the task names explicitly. G1 is cross-referenced, not duplicated: it is 257bm:i10 above.

Theorem 6.108 (257rig:i2 — `dyadic_support_fraction_reciprocalMass_diverges_or_gt_one`)**.** *If an INFINITE support* A *has* $\text{erdosSupportSeries}(2, A)$ *equal to a dyadic rational* $p/2^c$ *(any* c *— the easiest terminating case to rule out), then*

$$\neg\text{Summable}(1/a \text{ on } A) \ \vee\ 1 < \text{reciprocalMass}(A).$$

Proved by a Cesàro-mean identification of the shifted-tail-state average with the reciprocal mass, plus an explicit two-spike construction at $\text{lcm}(a, b)$ *for any two distinct positive support elements. Directly excludes any candidate infinite support with SMALL (convergent,* ≤ 1*) reciprocal mass — e.g. any support sparse enough for the lcm-gap engine's irrationality instances — from also being a dyadic-rational counterexample.*
[Lean] scale:uniform coord:other:cesaro-tail Lean source Lean source

Theorem 6.109 (257rig:i3 — `shifted_state_unbounded_of_infinite_support` (T13))**.** *For ANY infinite support* A*, any* $v > 0$*, any positive-valued* u *satisfying* $u(n{+}1) + v \cdot \text{supportCoeff}\, A\, (c{+}n{+}1) = 2u(n)$*: u is UNBOUNDED,* $\forall B\ \exists n,\ B < u(n)$*. Contrast with T7 (257bm:i-t7): T7's orbit is additionally required TEMPERED* $(o(2^n))$ *— boundedness and temperedness are different conditions; an unbounded-but-tempered orbit is possible, and is exactly what T7 constructs when the series IS rational. The proof selects* $2B + 1$ *positive support*

elements and uses their product as an explicit common multiple forcing supportCoeff *mass* $\geq 2B+1$ *there — a technique fully portable to any Dirichlet-incidence-style coefficient forced by an infinite index set.*
[Lean] scale:uniform coord:other:common-multiple-forcing Lean source Lean source Lean source

Theorem 6.110 (257rig:i4a — `supportCoeffZeroWindow_length_le_eps_logb` (T11)). *If an infinite-support series (with a positive element) equals a rational* $p/(2^c \cdot v)$ (v *odd): any run of* h *consecutive zero support-coefficients starting right after index* N *is bounded,* $\forall \varepsilon > 0\ \exists B,\ \forall N{\geq}1\ \forall h,\ \mathrm{SupportCoeffZeroWindow}(A, c{+}N, h) \to h \leq \varepsilon \log_2 N + B$ *— a rational-valued support series cannot have super-logarithmic gaps in its Dirichlet-incidence coefficient. Uses the elementary divisor-count bound* $\tau(n)^k \leq (k^{2^k})^k \cdot n$ (*257rig:i4b*) *feeding a support-tail envelope combined with the doubling-tail recurrence.*
[Lean] scale:uniform coord:other:divisor-envelope Lean source

Proposition 6.111 (257rig:i4b — `card_divisors_pow_le_divisorSubpowerConst_pow_mul`). $\tau(n)^k \leq (k^{2^k})^k \cdot n$ *for all* n, k *— a pure number-theory fact with zero support/problem content, directly reusable anywhere a* $1/k$*-power divisor-count envelope is needed. Would plausibly transfer to a #249 argument bounding zero-runs of* φ *itself, though* φ*'s growth is linear not polylog, so the same zero-window CONCLUSION would not transfer — only this underlying divisor-bound lemma might, via a different route.*
[Lean] scale:uniform coord:other:divisor-envelope Lean source

Proposition 6.112 (257rig:i5 — `mixedPrimePowerLayerTwo_supportCoeffInt`, `primePowerLayer_comm`). *For an integer-valued support-coefficient function and distinct primes* $p \neq q$ *with exponents* e, f*: the iterated exact-level-difference operator* primePowerLayer *commutes across the two primes, and two nested layers extract exactly the iterated exact-valuation pullback of the support,* $\mathrm{mixedPrimePowerLayerTwo}(p,e,q,f,\mathrm{supportCoeffInt}\,A,n) = \mathrm{supportCoeffInt}(\mathrm{pullback}_{q,f}(\mathrm{pullback}_{p,e}A))(n)$ *for* n *coprime to* pq*. Self-contained (81 lines); the module's own docstring is explicit it supplies no transport from rationality through multiplicative decimation and asserts no bounded-*Ω *endpoint — genuinely just algebra, not a rigidity result on its own.*
[Lean] scale:uniform coord:other:mobius-inversion Lean source Lean source

First-crossing machinery (`BooleanMobiusExactRowCrossing`)

Definition 6.113 (257bm:i-cross — `localMersenneCrossingRanks`).

$$\mathrm{localMersenneCrossingRanks}(E) \;:=\; E.\mathrm{filter}(c \mapsto \tfrac{1}{2} < \mathrm{localMersennePrefixValue}(E.\mathrm{filter}(\cdot \leq c))),$$

the ranks at which E's inclusive ordered prefix is already above one half.
[Lean] scale:n/a coord:binary-digit Lean source

Theorem 6.114 (257bm:i-cross2 — `exists_first_localMersenne_crossing`). *For* E *with* $\forall d \in E,\ 2 \leq d$ *and* $1/2 < \mathrm{localMersennePrefixValue}(E)$*:*

$$\exists c \in E,\ 4 \leq c \ \wedge\ \mathrm{localMersennePrefixValue}(E.\mathrm{filter}(\cdot < c)) < \tfrac{1}{2} \ \wedge\ \tfrac{1}{2} < \mathrm{localMersennePrefixValue}(\mathrm{insert}\ c\ (E.\mathrm{filter}(\cdot < c))).$$

There is a first rank $c \geq 4$ at which the running prefix strictly crosses one half; the crossing rank c feeds directly into 257bm:k4/257bm:c7's skipped-core constructors. The SHAPE (a monotone-ish accumulation crossing a threshold has a first crossing index) is a completely generic real-analysis fact; the specific localMersennePrefixValue *instantiation is #257-specific, but the same "first crossing rank" pattern would recur verbatim for any accumulating series — including a #249 φ-digit accumulation — if restated with a different weight function.*
[Lean] scale:uniform coord:binary-digit LEAN SOURCE

Möbius-centred carry nonnegativity and the skipped-rank trichotomy (half-membership track)

Theorem 6.115 (257hg:i6 — `mobiusCenteredHalfCarry_nonneg_of_supportSeries_lt_half` (G1)). *For* mobiusCenteredHalfCarry(A, N) *(an integer, defined via* integerHalfCarry*): if* $1 \notin A$ *and* erdosSupportSeries$(2, A) < 1/2$*, then*

$$\forall N, \quad 0 \leq \text{mobiusCenteredHalfCarry}(A, N),$$

via the identity integerHalfCarry$(A, N) = 2^{N+1}(1/2 - \text{seriesValue}) + \text{binaryCoeffTail}(\dots)$*, both terms nonnegative. Notably problem-agnostic in STATEMENT SHAPE:* supportCoeff A/integerHalfCarry *are generic divisor-incidence constructs, and since #249 ($\varphi/2^n$) is itself built from a Möbius/divisor-sum structure ($\varphi = \mu * \mathrm{id}$), this carry-nonnegativity mechanism is a strong candidate for direct reuse or close adaptation on the #249 side. The analytic half of 257bm:k12's contradiction machine below; also the base case consumed by 257rig:c16's sqrt-bound route.*
[Lean] scale:uniform coord:other:mobius-centred-carry LEAN SOURCE LEAN SOURCE

Theorem 6.116 (257hg:i7 — `halfGreedy_skipped_endpoint_trichotomy` (G5)). *At every genuinely skipped rank $s \geq 5$, the frozen margin* greedyHalfFrozenMargin$(s-1, s)$ *is in exactly one of three signed cells: NEGATIVE (actual word = seam-greedy word,* seamIntegerGreedyRemainder$(s) \geq 1$*, margin = −*remainder*); ZERO (actual word = seam-greedy word,* remainder $= 0$*); POSITIVE (actual word = the seam "above" word, margin =* overshoot*). A complete case split with exact algebraic identities for the margin in each branch — the zero-cell is provably SAFE/ALLOWED; only the negative cell needs excluding for 257bm:c13's socket (i) to close, and G2 (257bm:k12 below) is exactly the case split inside the negative cell that the corpus has managed to exclude so far.*
[Lean] scale:uniform coord:other:frozen-margin LEAN SOURCE

Total gap mass (`HalfGapMass`)

Definition 6.117 (257hg:i1 — `mersenneGap`). mersenneGap(n) := mersenneWeight(n) − mersenneTail$(n) > 0$, the per-level shortfall between a single weight and the true tail sum beyond it.
[Lean] scale:n/a coord:mobius-mersenne LEAN SOURCE LEAN SOURCE

Theorem 6.118 (257hg:i2 — `mersenneGap_tail_le`, `tendsto_mersenneGap_tail_zero`). *The TOTAL gap mass strictly beyond level N is summable and explicitly bounded:*

$$\sum_k{}' \,\mathrm{mersenneGap}(N+k+1) \;\le\; \frac{2}{9}\Big(\frac{1}{4}\Big)^N + \frac{3}{7}\Big(\frac{1}{8}\Big)^N \;\longrightarrow\; 0,$$

obtained purely by summing the pre-existing per-level asymptotic bound geometrically — no new arithmetic input. ***Scope caveat, explicit in the file header:*** *this bounds gap MASS only. It says nothing about which reals the greedy run reaches, does not certify membership or non-membership of any point, and in particular says nothing about* $1/2$ *directly. It is a budget statement: useful if a future argument needs "total measure of all fatal gaps beyond level* N *is small", e.g. a measure-theoretic rather than single-point argument.*
[Lean] scale:uniform coord:other:summed-gap-mass Lean source Lean source Lean source Lean source

The sharp fatal-gap criterion (`HalfGreedyFatalGap`)

Definition 6.119 (257hg:i3 — `mersenneTailLB3`, the three-channel Lambert lower bound)**.** $\mathrm{mersenneTailLB3}(k) := 1/2^k + 1/(3\cdot 4^k) + 1/(7\cdot 8^k) < \mathrm{mersenneTail}(k)$ — a strict rational lower bound obtained by truncating the Lambert-type series at three terms.
[Lean] scale:n/a coord:other:lambert-bound Lean source Lean source Lean source

Theorem 6.120 (257hg:i4 — `skipSafe_of_two_mul_le_three_mul`, the sharp $2u \le 3a$ criterion)**.** *For a skipped residual* $\rho = u/(2L)$ *at rank* k*, writing the skip-margin as* $2^k u + a = 2L + u$ *(so* $0 < a \Leftrightarrow$ *skip): the DYADIC safety test is* $u \le a$*, but the TRUE (tail-based) safety test only needs*

$$2u \le 3a.$$

One-third tighter than the dyadic test, discharged for the ACTUAL Mersenne tail (not just the three-channel lower bound) in `skipSafe_actualTail_of_two_mul_le_three_mul`*.* `sharp_of_dyadic : u ≤ a → 2u ≤ 3a` *shows dyadic-safe ⇒ sharp-safe (containment), and* `sharp_strictly_stronger` *exhibits* $(u,a) = (3,2)$ *as sharp-safe but NOT dyadic-safe — the containment is strict. This matches the "*$2u \le 3a$ *fatal band supersedes two-thirds" / "frontier* $3a \ge 2u$*" language on the UNCONDITIONAL (*$\forall$ *infinite* A*) #257 track recorded in project memory; whether it is the same theorem or an independently-proved analogue on this half-membership track is UNVERIFIED, flagged* `reported_prior`*, not re-verified equivalence.*
[Lean] scale:uniform coord:other:lambert-bound Lean source Lean source Lean source Lean source

Proposition 6.121 (257hg:i5 — unconditional safety corollaries)**.** `unitNumerator_skipSafe`*: unit numerators (*$u = 1$*) are always safe. The fatal converse forces* $3a < 2u$ (`three_mul_lt_two_mul_of_fatal`)*; fatality forces* $u \ge 2$ *unconditionally* (`two_le_of_fatal`) *and* $u \ge 3$ *once* u *is odd* (`three_le_of_fatal_of_odd`)*. Mirror statements* `unitNumerator_skipSafe_actualTail`, `three_mul_lt_two_mul_of_actualTail_fatal`, `three_le_of_actualTail_fatal_of_odd` *hold for the actual (not merely three-channel-bounded) Mersenne tail.*
[Lean] scale:uniform coord:other:lambert-bound Lean source Lean source Lean source Lean source

Two-sided dyadic-scale control (`HalfCylinderMiddleCarryLowerBound`, feeds `HalfUpperResetCriticalBand`)

Theorem 6.122 (257bm:i13 — `SeamTwoSidedDyadicCellEscape.twoSided`)**.** *Given a minimal local socket excluding exactly three middle "cells" plus one right-pulse-leak bound, induction (base case at row 5 verified by* `decide`*) propagates it to the universal two-sided bound*

$$\forall s \geq 5, \quad \min(\text{seamIntegerGreedyRemainder}(s),\ \text{overshoot}(s)) \ \leq\ 2^s.$$

A global dyadic-scale control on the seam orbit, independent of which branch fires. Exactly the ingredient 257bm:c11/c12 needs to make its forbidden-band inequalities well-posed. The "exclude finitely many exceptional local cells, get a global dyadic bound by induction" pattern is generic to digit-DP dynamics.
[Lean] scale:uniform coord:other:dyadic-scale LEAN SOURCE LEAN SOURCE LEAN SOURCE

Proposition 6.123 (257bm:i14 — `seamUpperResetDyadicBandEscape_through_thirty`, finite certificate to row 31)**.** *Explicit* `decide+kernel` *computation of* seamIntegerGreedyRemainder *for every row* 13–31*: row* 14 → 392, *...,* *row* 31 → 4187487147*, NONE failing the band-avoidance condition of* SeamUpperResetDyadicBandEscape ($\forall d \geq 13$, $\forall$*carries*, $\forall j \leq d$, $2^{d-j+1} <$ resetCharge $\vee$ resetCharge $+\ 2(d+j) \leq 2^{d-j+1}$)*. This is the current COMPUTATIONAL FRONTIER of the band-escape route — a finite verified list to a concrete row,* not *a cofinal supply. The open task is either extending the same finite-computation technique cofinally, or proving the general band-avoidance arithmetically (which is what 257bm:c11's quantifier collapse reduces the burden of).*
[Cert] scale:fixed coord:other:dyadic-boundary LEAN SOURCE LEAN SOURCE LEAN SOURCE

Abstract, fully problem-agnostic invariants read alongside this lane

Theorem 6.124 (257bm:i15 — `PerturbedFamily`, `AdjacentCut`, the abstract perturbed--greedy machine)**.** `structure PerturbedFamily`(α) *where* oldSum : $\alpha \to \mathbb{N}$; pulse : $\alpha \to \mathbb{N}$; gap : $\mathbb{N}$; pulseCap : $\mathbb{N}$, *with axioms* gap_pos, pulse_le, *injectivity of* oldSum, *a superincreasing separation axiom, and* pulseCap $< 3 \cdot$ gap*; derived* newSum$(x) :=$ oldSum$(x) +$ pulse(x)*. An abstract* `AdjacentCut`(C) *structure (best subset* $\leq C$*, best subset* $> C$*) yields* successorCarries $:= 4 \cdot$ overshoot $+$ abovePulse $\leq$ gap*; the resulting* prefixChoice *is provably MAXIMAL among all* newSum$(x) \leq$ newCapacity (`prefixChoice_maximal`)*, and* `nextRemainder_trichotomy` *gives the exact 3-branch (carry/middle/right) recurrence. Fully general — no Mersenne content whatsoever;* α, oldSum, pulse *are free parameters. THE single most reusable abstract result in this entire lane: any future base-b greedy digit process (a #249 analogue included) needing "maximality of the greedy choice under an admissible-capacity constraint" can instantiate this structure directly instead of re-proving maximality.*
[Lean] scale:uniform coord:other:abstract-perturbed-greedy LEAN SOURCE LEAN SOURCE LEAN SOURCE

Theorem 6.125 (257bm:i16 — `ReverseCarryWord`, abstract carry-overlap spacing)**.** `structure ReverseCarryWord` *where* coeff, bit, carry : $\mathbb{N} \to \mathbb{Z}$; normalized :

$\forall m$, $\mathrm{bit}(m) + 2 \cdot \mathrm{carry}(m) = \mathrm{coeff}(m) + \mathrm{carry}(m{+}1)$. *If two such words share a coefficient at a "seam" position with output bits* $1/0$*, then agree on both coefficients and bits for* length *further positions, their carry difference at the far end is EXACTLY* 2^{length} *times an odd integer* (`overlappingReverseCarryWords_carryDifference_eq_twoPow_mul_odd`). *Combined with a common Archimedean bound on both carries, this forces* $2^{\mathrm{length}} \le \mathrm{bound}$. *Pure integer-sequence recurrence, no Mersenne/base-4/seam structure at all — directly applicable to #249's own digit-carry structure if it is given an analogous reverse-carry presentation.*
[Lean] scale:uniform coord:other:reverse-carry-word LEAN SOURCE LEAN SOURCE LEAN SOURCE LEAN SOURCE

Proposition 6.126 (257bm:i17 — `supportCoeff_extend_true_eq_false_add_one_at_double`). *At row* $2(N{+}1)$*, changing only exponent* $N{+}1$ *from* `false` *to* `true` *in a* HalfWord(N) *adds EXACTLY the half-divisor incidence to* supportCoeff*: a one-bit flip changes a divisor-incidence sum by exactly the number of new divisor relations created — the smallest possible object-level coefficient-drop producer for* RewindBaseUnitDropAt*. A generic digit/divisor-incidence fact, restatable for any base-2 divisor-coefficient bookkeeping, e.g. #249's own divisor sums.*
[Lean] scale:uniform coord:other:half-divisor LEAN SOURCE LEAN SOURCE

6.5 Killers and countermodels

Every record below is a HAZARD RECORD: a proof that some plausible-looking route to cofinality, either does not work, or provably cannot work as stated. These are load-bearing warnings, not merely negative curiosities — several are the exact reason the corpus's live attack is aimed where it is aimed, and re-deriving what they already kill wastes a proof attempt.

Theorem 6.127 (257bm:k1 — `boundedDoubleOrRecycleModel_not_cofinal`, THE DICHOTOMY-SHAPE FALSIFIER — read before re-deriving cofinality from the dichotomy alone). *The dichotomy* ExactLocalMersenneHalfRow$(2n{-}1) \vee \exists c,\ 4 \le c \le n \wedge$ ExactLocalMersenneHalfRow$(2c{-}2)$ (*proved for* $n \ge 6$ *at 257bm:k-dich below*) *is not by itself enough for cofinality. Countermodel:* boundedDoubleOrRecycleModel$(n) := (n = 6)$ *satisfies EXACTLY the same two-branch transition shape* (*seed at 6, and* `boundedDoubleOrRecycleModel_transition` *reproduces the* $\vee$ *shape by always taking the recycle branch with* $c = 4$*, landing back at* $2 \cdot 4 - 2 = 6$), *yet*

$$\neg(\forall N,\ \exists n \ge N,\ \mathrm{boundedDoubleOrRecycleModel}(n)) \qquad (\textit{NOT cofinal — only ever true at } n = 6).$$

Packaged existentially as `exists_seeded_bounded_double_or_recycle_model`. *The double-or-recycle transition shape, EVEN TOGETHER with an endpoint-six seed, is logically insufficient to conclude cofinal exact rows. A genuinely new progress input* (*strict endpoint growth, or an independent cofinality argument*) *is required. The TECHNIQUE — build a bounded one-point fixed model with the same recursion shape to show a transition schema doesn't imply cofinality — is fully general and transfers to any other double-or-recycle-shaped inductive scheme, including one built for #249.*
[Lean] scale:fixed coord:other:meta-logical LEAN SOURCE LEAN SOURCE LEAN SOURCE LEAN SOURCE

Theorem 6.128 (257bm:k-dich — `exactLocalMersenneHalfRow_double_or_recycle`, the dichotomy 257bm:k1 attacks)**.** *For $n \geq 6$,* ExactLocalMersenneHalfRow(n)*:*

$$\text{ExactLocalMersenneHalfRow}(2n-1) \;\vee\; \exists c,\ 4 \leq c \leq n \wedge \text{ExactLocalMersenneHalfRow}(2c-2).$$

Every exact row of endpoint ≥ 6 has exactly one of two continuations — literal doubling (below-half) or recycling through a first-crossing to some $2c-2$ with $c \leq n$ (the recycled endpoint is not *claimed to exceed n). Equality is excluded by 257bm:k11 (`finiteErdosSum_den_odd`). What is missing, per 257bm:k1, is not another dichotomy but STRICT ENDPOINT PROGRESS in the recycle branch, or a proof the below-half branch recurs — any future attack that merely re-derives a double-or-recycle transition is already refuted.*
[Lean] scale:bounded coord:mobius-mersenne LEAN SOURCE

Proposition 6.129 (257bm:k2 — `splitFractionMass_one_bound_not_necessary_fixture`, sufficient-but-not-necessary fixture)**.** *Concrete fixture $D = \{2,3\}$, $c = 5$:*

$$\text{value}(D) < \tfrac{1}{2} \;\wedge\; \tfrac{1}{2} < \text{value}(\text{insert}\,5\,D) \;\wedge\; 1 < \text{localFractionMass}(D,8) + \text{localMersenneFraction}(8,5) \;\wedge\; \text{localBinarySuffix}(D,1,8) < 2^3.$$

`norm_num`*-checked. The "combined residue mass ≤ 1" sufficient condition for sharp capacity (*`localBinarySuffix_two_mul_sub_two_lt_criticalCapacity_of_splitFractionMass`*) is genuinely NOT necessary — here the mass exceeds 1 yet sharp capacity still holds. DO NOT attack* SkippedCoreCriticalQuotientSupply *(257bm:c4) via the fractional-mass route alone; it will fail on real crossing cores. Use 257rig:i3-wall (*`eq_halfGreedyPrefixSupport_of_critical_crossing`*) to reduce to the one canonical orbit instead, and attack the orbit's own binary digits directly.*
[Cert] scale:fixed coord:mobius-mersenne LEAN SOURCE

Observation 6.130 (257bm:k3 — doubling immediately produces an above-half row, numeric countermodel). 257bm:i7's doubling extension is NOT preserved iteratively: below-halfness is not preserved under one application of doubling, and demonstrably fails at the very first step. From the identity $\text{value}(E) - \frac{1}{2} = (\text{localFractionMass}(E, 2n-1) - 1)/2^{2n-1}$, at the seed $D = \{2,3,6\}$ (257bm:i-seed) with $M = 11$ the core alone contributes

$$\frac{2^{11} \bmod 3}{3} + \frac{2^{11} \bmod 7}{7} + \frac{2^{11} \bmod 63}{63} = \frac{2}{3} + \frac{4}{7} + \frac{32}{63} \approx 1.746 > 1,$$

so the doubled row is STRICTLY ABOVE half. Iterating the cheap (doubling) arm is therefore impossible from the only landed seed; the chain is forced into the capacity-gated recycle arm at step one. What would close the doubling route instead: cofinally many BELOW-HALF exact rows, $\forall N\ \exists n \geq N\ \exists D$ exact at n with $\text{localFractionMass}(D,n) < 1$ — a single such row at each of cofinally many n would make doubling unnecessary; a self-reproducing one would close everything.
[Cert] scale:fixed coord:mobius-mersenne LEAN SOURCE

Theorem 6.131 (257bm:k4 — the recycling witness is capped by input n, not guaranteed to grow)**.** *257bm:c10 (*`exists_skippedCoreExactRow_of_value_above`*) is UNCONDITIONAL, but its witness is $\exists c \leq n$, not $\exists c$ large: $2c-2$ may be $\leq n$, so the endpoint need*

not grow. 257bm:k1's `exists_seeded_bounded_double_or_recycle_model` *is the explicit falsifier of the naive hope that growth comes for free:* boundedDoubleOrRecycleModel$(n) :=$ $(n = 6)$ *satisfies the same transition schema plus a seed and is NOT cofinal. Growth is recovered only inside* ProtectedExactLocalMersenneRow *(257bm:c5), whose invariants* endpoint $<$ $2 \cdot$ cutoff *and* new_above_cutoff *force* $c >$ cutoff *hence* $2c - 2 >$ endpoint *— and maintaining those invariants is precisely what needs the strict-upper (sharp capacity) fill of 257bm:c7 rather than this general recycling theorem.*
[Lean] scale:bounded coord:binary-digit LEAN SOURCE

Observation 6.132 (257bm:k5 — `CofinalPositiveHalfGreedySkips` is not a reduction, it IS the problem — a trap for attackers). The positivity conjunct of 257bm:c2's hypothesis, $0 <$ greedyMersenneRemainderRat$(1/2, c-1)$, is unconditionally dischargeable (`localMersennePrefixValue_halfGreedy_lt_half`, itself just the odd-denominator parity fact of 257bm:k11, together with `greedyMersenneRemainderRat_eq_sub_finiteErdosSum`). Hence

CofinalPositiveHalfGreedySkips $\iff$ (greedyMersenneSkippedSupport$(1/2)$).Infinite,

which 257bm:c20 (A6) proves EQUIVALENT to $1/2 \in$ mersenneAchievementSet. Anyone attacking 257bm:c2's hypothesis head-on is attacking #257-false directly, with NO reduction whatsoever, despite the corpus's own docstrings advertising it as "the shortest known path". The genuine escape is to abandon the greedy-orbit arm entirely (where 257rig:i3-wall forces canonicity and removes all freedom) and target the form-$(\star)$ socket of 257bm:c7 instead, whose cores are NOT forced to be greedy prefixes.
[Lean] scale:cofinal coord:greedy-orbit LEAN SOURCE

Theorem 6.133 (257rig:k6 — `eq_halfGreedyPrefixSupport_of_critical_crossing`, THE WALL — the reduction to one universal-carry orbit)**.** *For $c \geq 4$, D bounded $[2, c)$, below-half, with genuine crossing deficit* $\frac{1}{2} -$ value$(D) <$ mersenneWeightRat(c)*:*

$$D \;=\; \text{halfGreedyPrefixSupport}(c-1).$$

EVERY crossing core is forced, uniquely, to be the canonical rational half-greedy prefix. There is no freedom left in D at all — 257bm:c4's universal quantifier collapses onto a single orbit (the greedy digit-by-digit construction of $1/2$ *in Mersenne weights), proved via* `IsStraddlePrefix.half_agrees_greedy` *(superincreasing straddle-prefix rigidity, 257bm:i11a) plus a primitive-prefix bridge. This is the killer of the idea that* SkippedCoreCriticalQuotientSupply *is a search over arbitrary D: it is a check against ONE orbit. A future attack should target the concrete arithmetic of* halfGreedyPrefixSupport *directly. The straddle-prefix/superincreasing-sequence rigidity technique is general; whether a #249-side analogue exists (is the greedy φ-digit prefix similarly rigid?) is UNVERIFIED, worth checking.*
[Lean] scale:uniform coord:greedy-orbit LEAN SOURCE

Observation 6.134 (257rig:k7 — 257bm:c7 and 257rig:c17 trade off in opposite directions, neither dominates). 257rig:c17 (`HalfCarryCofinalTerminalOnlyStrip`) tolerates ANY support inside the depth-M window but demands the carry inside a $\sim 2\sqrt{M}$ strip; 257bm:c7 tolerates a carry up to 2^{c-2} (exponentially looser) but demands the support

be CONFINED to ranks $< c$, because the exact fill encodes the residue in the pure upper window where the Mersenne quotient is exactly 2^{M-d} (257bm:c9). The gap between them is support LOCALITY, not carry size, and the two have never previously been compared on this axis. A lower-half confinement upgrade of the strip witnesses (support in $[2, c)$ at depth $2c-3$) would make $2\sqrt{2c}+4 < 2^{c-2}$ for $c \geq 8$ and fire the existing fill verbatim — which would identify 257rig:c17, not 257bm:c7, as the CHEAPER route to the same payoff.
[Lean] scale:cofinal coord:other:sqrt-carry-growth Lean source

Observation 6.135 (257hg:k8 — explicit dyadically-unsafe example, the two-thirds-band countermodel). Read in full alongside `HalfGreedyFatalGap` for scope contrast: writing a greedy residual as rem $= 1/R$, a skip at rank k is dyadically safe iff $R \geq 2^k$. Two unconditional corollaries hold (an INTEGRAL reciprocal R is never unsafe, `not_twoThirdsBand_of_int`; an unsafe run with odd p, D, q forces $p \geq 7$, `seven_le_of_intBand_odd`) — but the file's own docstring is explicit that dyadic safety is SUFFICIENT, not NECESSARY, for the true greedy process to survive (the true test uses the actual tail $\sum_{j>k} \text{weight}(j)$, which exceeds 2^{-k}), and NOTHING in this file asserts the actual half-greedy orbit avoids the band: $(p, D, b) = (17, 41, 3)$ is an explicit odd-coprime example that IS dyadically unsafe. This sharpens directly into 257hg:i4's $2u \leq 3a$ criterion, which narrows exactly this gap for the single-skip case — but the general question of whether the fatal region overlaps a reachable band under the actual greedy trajectory is settled by neither file.
[Lean] scale:n/a coord:other:rational-band Lean source Lean source Lean source

Theorem 6.136 (257bm:k9 — `relationInvariantLinearChannels_det_eq_zero`, a whole determinant-obstruction family killed at once)**.** *If a finite family of linear channels* channel : $\iota \to V \to_{\mathbb{Q}} \mathbb{Q}$ *all vanish on* ker(ev) *(i.e. are relation-invariant like* ev *itself), then the evaluation matrix* $(i, j) \mapsto$ channel j (row i) *has rank* ≤ 1 *for EVERY finite index type* ι*, hence every square minor of size* ≥ 2 *has determinant zero. This closes the "shifted-channel determinant" obstruction route at ALL ranks at once — previously checked only experimentally at ranks 2–4. Pure linear algebra over* $\mathbb{Q}$*-vector spaces, zero arithmetic content specific to Mersenne or #257. Cite this whenever someone proposes a new multi-channel-relation determinant approach; the family it kills is unbounded, not just the small cases already tried.*
[Lean] scale:uniform coord:other:linear-algebra Lean source

Observation 6.137 (257bm:k10 — the loose capacity bound loses exactly one bit, additively not multiplicatively — do not try to shave it by a constant factor). 257bm:i6's 2^{c-1} bound and 257bm:i5's sharp 2^{c-2} requirement differ by ONE bit, and the proof of the loose bound shows exactly where: it derives $A < 2^{c-2} + |D|$ and discards $|D| \leq c - 2 \leq 2^{c-2}$. Any attempt to close the gap by re-deriving the same additive slack estimate more carefully is bounded by this identity — the ONLY way to close it is a genuine anti-concentration statement that A avoids the linear-width band $[2^{c-2}, 2^{c-2}+c-3]$, not a sharper constant in the existing proof.
[Lean] scale:uniform coord:binary-digit Lean source

Observation 6.138 (257bm:k11 — `finiteErdosSum_den_odd` is UNVERIFIED-IN-FULL by this pass — flag before re-citing). This lemma is cited repeatedly throughout the lane

(TH-dichotomy's equality exclusion, 257bm:k2's fixture, 257bm:i10's finite-support-never-hits-half corollary) but sits OUTSIDE this lane's file set at `DyadicPrefixCompression.lean:1358`; only its call sites, not its own body, were read in the passes that built this catalogue. Treat its conclusion (localMersennePrefixValue$(D) \neq 1/2$ for finite D) as [Cited], not independently re-verified here, until a pass opens that file directly.
[Cited] scale:uniform coord:p-adic LEAN SOURCE

Theorem 6.139 (257hg:k12 — `finalMiddleCell_neg_three_not_last` (G2), one of three exceptional cells excluded)**.** *The first of three exceptional final-middle integer cells of 257hg:i7's trichotomy (coordinate* $4 \cdot \text{remainder} - \text{belowPulse} - 4 = -3$*) CANNOT be the actual last transition before an all-right tail. Proof combines the analytic strict inequality* $\text{value}(u) + \text{mersenneTail}(D) < 1/2$ *(a middle-branch consequence of an all-right-tail hypothesis) with the two-step Möbius-carry recurrence* (`mobiusCenteredHalfCarry_add_two`) *forcing the centred carry to zero then strictly NEGATIVE within two more coefficient rows — contradicting 257hg:i6's nonnegativity. This is a genuine exclusion, not a reduction: cell* -3 *is dead. A companion file* (`HalfCylinderFinalMiddleTailSocket.lean`) *records the corresponding* binaryCoeffTail *inequalities for the remaining two cells,* -2 *and* -1 (`binaryCoeffTail_union_Ioi_lt_two_of_producerCarry_eq_neg_two`, `_lt_three_of_producerCarry_eq_neg_one`) *— closing all three would fully resolve the final-middle producer and hence 257bm:c13's socket* (i) *outright.*
[Lean] scale:uniform coord:other:mobius-centred-carry LEAN SOURCE LEAN SOURCE LEAN SOURCE

6.6 Recently published declarations

The modules surveyed above were read at an earlier snapshot. Four already-public modules have since grown by union merge, and one wholly new problem-centric tree, `ErdosProblems/Erdos257/`, has been published alongside the legacy `Erdos249257` namespace. This subsection catalogues only the declarations the earlier parts do not cite. None of it decides O1–O5 above; it fills in machinery that those targets' own producer chains were stated abstractly against.

6.6.1 The middle-cell update-step chain (`HalfCylinderMiddleCarryLowerBound.lean`)

The O5 discussion above states `SeamTwoSidedDyadicCellEscape` and `SeamUpperResetDyadicBandEscape` as named hypotheses without exposing the per-row update machinery that actually produces and consumes them. That machinery is present in the live file.

Proposition 6.140 (The unsafe middle range is exactly three integers)**.** *For* $s \geq 5$*, writing* $R = (\textit{seamAdjacentCut}\ s\ hs).\text{remainder}$ *and* $P = (\ldots).\text{belowPulse}$*,*

$$\neg(4R - P - 4 \leq -4 \ \vee\ 0 \leq 4R - P - 4) \iff 4R - P - 4 \in \{-3, -2, -1\}$$

(LEAN SOURCE, [Lean], scale:uniform, coord:mobius-mersenne, proved by `omega` from the integer definitions). This is the exact unconditional source of the "three integer values" language used above for O5; it was previously asserted in prose only.

Two companion bounds pin the scaled rational residual inside a genuine middle transition: $4 \cdot \texttt{seamGreedyFloorZ}\, s \le c$ from a cell label $4R = P + c$ (Lean source, [Lean]), and, combining it with the successor floor-Z recurrence, $4^{s+1} \cdot \texttt{seamWordRationalRemainder}(\texttt{seamGreedyWord}(s+1)) \le 2^{s+2} + c$ (Lean source, [Lean]). The module's own remark is that for the first exceptional cell ($c = -3$, i.e. the cell already closed by Lean source) this leaves margin exactly 1 against a needed correction-tail slack of $4/3 + \varepsilon$ — a concrete near-miss, not a proof, and it does not extend to the two open cells $-2, -1$.

A parallel iff for the right branch: an overshoot witness $\le 2^s$ produces a pulse leak above 2^{s+2} exactly when the overshoot sits within an explicit distance k of 2^s with $4k <$ abovePulse (Lean source, [Lean]). A new one-coordinate normal form, $\texttt{seamPureHalfPrefixRemainder}\, s := \texttt{integerGreedyRemainder}(\texttt{seamWeights}\, s)(2^{2s-1})$ (Lean source), is given an exact case-split closed form (Lean source, [Lean]) that the file states is equivalent to the two-sided invariant, but this equivalence is not itself yet packaged as a separate reusable lemma beyond Lean source ([Lean]).

Observation 6.141 (The reset-band socket has a certified base case). `SeamUpperReset-DyadicBandEscape` (O5's second socket, discussed after §12.5 above) is reduced by three new lemmas to a quarter-scale overshoot exclusion: the iff Lean source ([Lean]) rewrites the reset-charge band as a successor-remainder avoidance; the identity Lean source ([Lean]) rescales it by a factor of 4; and Lean source ([Lean]) derives the band from a quarter-scale overshoot exclusion plus the pulse bound abovePulse $\le 2(d-2)$. Row $d = 13$, the first row of the late regime, is then proved unconditionally to escape *every* band simultaneously: its successor remainder is exactly 392 (Lean source, [Lean], [Cert] for the numeral 392). This is one certified row of a socket that is $\forall d \ge 13$; it is evidence the base case is not itself an obstruction, nothing more.

6.6.2 Forgetting witnesses: the terminal-only bridge (`SuffixCylinderTerminalOnlyBridge.lean`)

This entire file is new and previously uncited. It is a lossy-projection bridge: every object in the corpus's much richer suffix-cylinder stage machinery (`CylinderStage`, `In-StripTwoSheetStage`, `ProfiledGapStage`) is shown to already contain a terminal-only strip witness, the weaker object the compactness-based O3/O2 consumer actually needs.

Proposition 6.142. *Every* `CylinderStage` *at depth* N *yields* `HalfTerminalOnlyStripWit-ness` N (Lean source, *[Lean], coord:support rigidity, half-carry*), *and at a feedback row this sur-vives* both *outputs of the corpus's total feedback theorem — full-cylinder advance or a localized one-hole seam, which can delete at most one of carries* 3, 4 (Lean source, *[Lean]*).

Four concrete certified instances follow at depths 51–54 (Lean source, Lean source, [Lean], [Cert] for the underlying stage), which is a finite witness chain, not a cofinal one.

Observation 6.143 (A new named producer socket for HALF, not yet in the implication table). The end-to-end consumer

$$(\forall N,\ \exists M, K,\ \max(N,1) \le M \wedge \mathrm{Nonempty}(\texttt{CylinderStage}\ K\ M)) \implies \exists A,\ A.\mathrm{Infinite} \wedge \texttt{erdosSupportSeries}\ 2\ A = 1/2$$

is proved unconditionally (Lean source, [Lean]), with a positive-support refinement using the same odd-denominator lemma Lean source already used for O2 above (Lean

source, [Lean]). The hypothesis, cofinal nonempty full-cylinder stages, is a genuinely new named socket for HALF: it has not been checked against O3's `CofinalExactLocalMersenneHalfRows` or O4's `SQRTESC`, and its own supply is [Open]. It is only a sufficient condition, proved in one direction; no equivalence claim is made anywhere in the file.

6.6.3 A strict weakening chain for the last-producer socket (`HalfCylinderLastProducerContradiction.`

Five declarations extend the already-cited `SeamMiddleProducerSqrtEscape` socket with a genuine strength reduction, wiring the -3 cell closure into the producer chain rather than leaving it as an isolated fact.

Definition 6.144. SeamMiddleProducerTailEscape is the exact hypothesis that at every genuine middle transition ($s \geq 13$, no successor carry, below the terminal-weight threshold) the terminal-augmented binary coefficient tail is strictly below the producer carry (Lean source). SeamMiddleProducerTailEscapeExceptNegThree is the same hypothesis restricted to cells $\neq -3$ (Lean source).

The reduction TailEscape $\Rightarrow$ ExceptNegThree is trivial by hypothesis-dropping (Lean source, [Lean]); the substance is in the converse direction of the fan-in. The reduced socket, with cell -3 discharged by the already-proved Lean source, already suffices to close HALF:

$$\text{SeamMiddleProducerTailEscapeExceptNegThree} \implies (1/2 : \mathbb{R}) \in \mathcal{A}$$

(Lean source, [Lean]), with the unreduced socket and the already-cited square-root majorant both routing through it (Lean source, [Lean]). So the chain now reads SqrtEscape ⇒ TailEscape ⇒ ExceptNegThree ⇒ HALF, three strictly decreasing sockets where previously only the endpoints were exposed; all three antecedents remain [Open].

6.6.4 The governed frozen-margin producer (`HalfCylinderFiniteShadow.lean`)

One further declaration completes a chain the earlier parts cite only up to its antecedent equivalence (Lean source).

Definition 6.145 (HalfGreedyGovernedFrozenMarginProducer)**.** For every k with `mersenneWeight`$(k+1)$ exceeding the current greedy remainder, some $J \leq k+1$ gives a nonnegative frozen margin (Lean source).

$$\text{HalfGreedyGovernedFrozenMarginProducer} \implies (1/2 : \mathbb{R}) \in \mathcal{A}$$

(Lean source, [Lean]), proved via the pre-existing Lean source and the excess-sign equivalence already noted above. This is the module's actual promised consumer; the antecedent's own first-passage bound ($J \leq k+1$, uniform) is [Open] and has not been checked against any of O3–O5's producer sockets.

6.6.5 Exact carry-pivot normal form (`SuffixCylinderCarryPivot.lean`)

Numeral-adjacent boundary words have more structure than the previously used coefficient bound. If the lower suffix numeral is one larger than the upper suffix numeral, there is a unique carry pivot d: rank d enters the lower support, every rank e with $d < e \le N$ flips from lower-unselected to upper-selected, and all ranks below d agree (Lean source, [Lean], scale:uniform, coord:support rigidity, half-carry). Hence for every positive coefficient row $m \ge 1$ the boundary discrepancy is the exact divisor-incidence identity

$$c_{\mathrm{low}}(m) + \#\{e : d < e \le N,\ e \mid m\} = c_{\mathrm{up}}(m) + \mathbf{1}_{d\mid m},$$

not merely an inequality (Lean source, [Lean]). The stage-level theorem consumes adjacent profiled prefixes directly (Lean source, [Lean]). This is a local normal form for the two-sheet boundary. It does not produce adjacent stages cofinally and therefore does not close HALF or universal #257.

6.6.6 New tree: `ErdosProblems/Erdos257/MersenneSubseriesRigidity.lean`

This module works in the new `ErdosProblems.Erdos257` namespace and states in its own header that its results "concern achievement-set geometry and do not settle the universal irrationality problem." For an arbitrary support restriction $J \subseteq \mathbb{N}$, define `supportedMersenneAchievementSet` J as the range of Mersenne digit values supported on J (Lean source). It is compact (Lean source, [Lean]), nowhere dense for every J (Lean source, [Lean]), and perfect — compact with no isolated points — whenever J is infinite (Lean source, [Lean]), via a Cantor-space no-isolated-point argument on the coding space itself (Lean source, [Lean]).

Proposition 6.146 (Exact Lebesgue-measure dichotomy)**.** *Either $J = F^c$ for a finite F, and* $\mathrm{vol}(\texttt{supportedMersenneAchievementSet}\,J) = 2^{-|F|}$ *exactly, or J^c is infinite and the volume is exactly* 0.

(Lean source, [Lean], scale:uniform, coord:mobius-mersenne). The finite case is an exact doubling induction over inserted coordinates (Lean source, [Lean], using disjointness of the two coordinate-split faces, Lean source, [Lean], which is exactly where unique binary coding is spent). None of this module touches irrationality of any subseries value; it is a self-contained measure-and-topology classification of achievement sets, orthogonal to O1–O5.

6.6.7 New tree: `ErdosProblems/Erdos257/HalfCounterexampleFrontier.lean`

This short module restates U257 in the new namespace (Lean source) and gives it a problem-centric interface to the legacy half-terminal producer already used for O2 above: any `HalfTerminalOnlyScaledVanishingSequence` witness gives an infinite A with `erdosSupportSeries` $2\,A = 1/2$ (Lean source, [Lean]), hence refutes U257 outright (Lean source, [Lean]). The module's own docstring states plainly that the producer sequence itself "has only been checked to a large finite cutoff": this is a restatement of

the already-[Open] O2 producer obligation in a new namespace, not a new result, and it duplicates §O2's own logic (any HALF witness is ¬U257) rather than extending it. The module now also proves an unconditional finite-support obstruction at the second rational target $1/21$:

$$F \subseteq \{n : n \geq 2\} \quad \Longrightarrow \quad \sum_{n \in F} \frac{1}{2^n - 1} \neq \frac{1}{21}$$

for every finite F (Lean source, [Lean], scale:uniform, coord:finite-lcm). The denominator-order identity forces $\operatorname{lcm}(F) = 6$, after which the eight subsets of $\{2, 3, 6\}$ are exact arithmetic. This removes finite termination for that target; it does not establish achievement-set membership.

6.6.8 New tree: `Erdos249257/Primitive23Multiplicity.lean`

This module records the exact arithmetic of the primitive cone $2p + 3q = n$, $p, q > 0$, $\gcd(p, q) = 1$ (Lean source). It proves that every $n \geq 11$ has such a representation (Lean source, [Lean], scale:uniform, coord:finite-lcm), while rank 10 has none (Lean source, [Lean]). Rank 11 already has two distinct witnesses (Lean source, [Lean]), and so does every rank $10k$ with $k \geq 2$ (Lean source, [Lean], scale:uniform, coord:finite-lcm). These theorems are deliberately recorded as route diagnostics rather than as a counterexample: the file does not prove that the associated integer-multiplicity expansion of $1/21$ Booleanises. The recurring multiplicities are precisely the unresolved collision data, not zero–one support digits.

6.6.9 Current $1/21$ frontier: `Erdos249257/TwentyOneQuotientGreedy.lean`

The later quotient-greedy analysis removes the diffuse alternatives. Let `TwentyOneFatalAlignedBranch` denote the explicit branch consisting of a fatal greedy gap, finite skipped support (hence cofinite selection), eventual alignment between the quotient rows and the rational greedy prefix, and eventual visits to every doubling block. Then

$$\frac{1}{21} \in \texttt{mersenneAchievementSet} \quad \Longleftrightarrow \quad \neg\, \texttt{TwentyOneFatalAlignedBranch}.$$

Lean source [Lean] scale:uniform coord:greedy-orbit.

On the fatal branch the canonical quotient remainder s_R is eventually strictly above the closed capacity 2^R (Lean source), and thereafter the support appends $R+1$ while s_R follows one explicit affine recurrence (Lean source). Conversely, closed rows $s_R \leq 2^R$ along any unbounded sequence decay after normalisation and put $1/21$ in the achievement set (Lean source; Lean source).

This is the exact current contribution: membership has been reduced to excluding one named permanent affine-supercapacity regime. No theorem excludes that regime, forces unbounded closed returns, or proves membership. Combined with the finite-support obstruction above, membership would provide an infinite-support rational counterexample to universal #257; its present status remains [Open].

7 The scale ladder

Every claim in this programme carries a *scale* tag alongside its evidence band. The two axes are independent and both matter to a reader deciding whether a result can be used as a premise. Evidence answers "is this checked, and how": [Lean] is kernel-checked in the live Lean tree, [Cert] is an exact finite computation (integer arithmetic, no floating point, often independently reproduced), [Math] is ordinary mathematics carried out in prose and checked by a human or an adversarial audit pass but not yet formalised, [Cited] invokes a published external theorem, and [Open] marks an unproved hypothesis that some consumer theorem is conditioned on. Scale answers a different question, "over how much of the problem does this claim range": scale:fixed is one concrete numeral or witness; scale:bounded produces a witness whose size is controlled by the input (typically $\leq n$ for input n); scale:uniform holds unconditionally for every n above an explicit threshold, with no existential slack; scale:cofinal asserts $\forall N\, \exists n \geq N, \ldots$ — infinitely many witnesses, arbitrarily far out, with *no* coherence required between them. scale:n/a tags definitions and other scale-free statements.

Scale is the load-bearing axis of the whole corpus. Both open problems are, at bottom, cofinal statements: #257 half-membership needs infinitely many exact rows or infinitely many positive greedy skips (Table 2, rows C1/C3), and the unconditional track needs a run-length bound that holds at every reset row out to infinity, not merely to a certified depth. Every machine-checked theorem in the tree is fixed, bounded, or uniform. No amount of raising a uniform certificate's threshold, by itself, produces a cofinal witness — a uniform statement "$\forall n \geq N_0$" is a single first-order sentence with no existential content past N_0, while a cofinal statement needs infinitely many *witnesses to a further existential*, which is a strictly different logical shape.

Table 2 samples the ladder across all four non-trivial scales, drawn from the `BooleanMobius`, `HalfCutLocator`, `GreedyAchievementSet`, and seam-integer (`HalfCylinder*`) families together with the truncation-rung ladder and the sign/irrationality primitives that feed all of them. The horizontal rule marks the frontier: everything above it is proved (fixed, bounded, or uniform); everything below it is what the open producers still need (cofinal), stated with its exact quantifier prefix so the gap is visible, not merely asserted.

Result	**Scale**	**Quantifier prefix**	**Coordinate**
Fixed — one concrete witness or numeral			
Smallest concrete exact row: support $\{2,3,6\}$ satisfies localPrefixQuotient $= 2^5 - 1$. LEAN SOURCE [Cert]	scale:fixed	$n = 6$	coord:mobius-mersenne
One-point counter-model: the double-or-recycle transition shape alone, even seeded, is not cofinal. LEAN SOURCE [Lean]	scale:fixed	$n = 6$ (free-standing model)	coord:other:meta-logical

Result	Scale	Quantifier prefix	Coordinate
The combined-residue-mass ≤ 1 sufficient test for sharp capacity is not necessary: fixture $D = \{2,3\}$, $c = 5$. Lean source [Cert]	scale:fixed	$D = \{2,3\}$, $c = 5$	coord:mobius-mersenne
Exact residual fixture: $\frac{1}{2} - W(\{2,3,6,7\}) = \frac{1}{16002}$. Lean source [Cert]	scale:fixed	none (numeric instance)	coord:other:regression-fixture
Closed-form Möbius–Lambert value $\sum_{d\geq 1} \mu(d)/(2^d - 1) = 1/2$ (imported identity). Lean source [Cited]	scale:fixed	none (closed-form value)	coord:mobius-mersenne
The one natural negative-Möbius candidate support overshoots: $1/2 < \sum_{d\in N} 1/(2^d - 1)$ for $N = \{d \geq 2 : \mu(d) = -1\}$. Lean source [Lean]	scale:fixed	one fixed candidate support N	coord:mobius-mersenne
Bounded — a witness exists, controlled by the input			
Every exact row continues: doubles unboundedly, or recycles to some $c \leq n$. Lean source [Lean]	scale:bounded	$\forall n \geq 6$, given a row at n: $(\ldots) \vee \exists c \leq n(\ldots)$	coord:mobius-mersenne
Every above-half bounded support recycles into a genuine exact row at some $2c - 2$ with $c \leq n$. Lean source [Lean]	scale:bounded	$\forall E$ bounded $[2,n]$: $\exists c \leq n$	coord:binary-digit
(#249 cross-lane check) $\sum \varphi(n)/2^n$ matches no rational whose denominator divides $2^{14}(2^h - 1)$, for any $h \leq 16$. Lean source [Cert]	scale:bounded	$\forall h \leq 16$ (explicit bound)	coord:binary-digit
Truncation-rung ladder: HalfRung(J) proved for every $3 \leq J \leq 22$ via the finite decision procedure (§8, Theorem 8.9). (source: `erdos257_truncation_rung_ladder_2026_07_24.md`, §5) [Math]+[Cert]	scale:bounded	$3 \leq J \leq 22$ (finite table, exhaustive)	coord:rung-truncation
Uniform — unconditional for every n past an explicit threshold			
Master achievement/greedy-survival criterion. Lean source [Lean]	scale:uniform	$\forall x\ \forall n$ (iff)	coord:greedy-orbit
Membership iff every actually-skipped rank survives. Lean source [Lean]	scale:uniform	$\forall M \in$ skipped support	coord:greedy-orbit

Result	Scale	Quantifier prefix	Coordinate
Every crossing core is forced to be the canonical rational half-greedy prefix. Lean source [Lean]	scale:uniform	$\forall c \geq 4\ \forall D$ bounded $[2, c)$	coord:greedy-orbit
Theorem A payoff: reset $\sqrt{}$-escape socket ⇒ membership. Lean source [Lean]	scale:uniform	$\forall s \geq 14$, given the socket hypothesis	coord:seam-integer
Band-avoidance certified for every row $13 \leq r \leq 31$ (feeds the F5/Theorem A uniform hypothesis as an exhaustively-verified finite range). Lean source [Cert]	scale:uniform	$13 \leq r \leq 31$, exhaustive	coord:seam-integer
Critical-band-index reduction: a $(d{+}1)$-way check collapses to one nearest-boundary check. Lean source [Lean]	scale:uniform	$\forall d, E$ with $E \leq 2^{d+1}$	coord:other:dyadic-boundary
Parity kill: no finite support hits $1/2$ exactly. Lean source [Lean]	scale:uniform	$\forall A$ finite, $0 \notin A$	coord:other:denominator-parity
Erdős–Borwein irrationality at every integer base. Lean source [Lean]	scale:uniform	$\forall b \geq 2$	coord:mobius-mersenne
Erdős 1968: any infinite, pairwise-coprime, reciprocal-summable index set gives an irrational support series. Lean source [Lean]	scale:uniform	$\forall b \geq 2\ \forall A$ (infinite, coprime, summable)	coord:mobius-mersenne
Generic block-certificate engine (problem-agnostic): weighted-coefficient series is irrational given a per-precision divisibility/window certificate. Lean source [Lean]	scale:uniform	$\forall b \geq 2\ \forall c\,(\forall m, c(m) \leq m)$, given $\forall q\ \exists N, K, L, C$	coord:binary-digit
Finite greedy-prefix denominators are odd (source of the parity boundary exclusion, §8). Lean source [Lean]	scale:uniform	$\forall n$	coord:other:denominator-parity
— frontier: every entry below needs a cofinal witness, $\forall N\ \exists n \geq N(\ldots)$, not yet supplied —			
Cofinal — the shape both open producers actually need			
THE master open supply: exact rows recur at arbitrarily large endpoints, with no coherence required between witnesses. Lean source [Open]	scale:cofinal	$\forall N\ \exists n \geq N$, ExactLocalMersenneHalfRerse(m)e	coord:mobius-mersenne

Result	Scale	Quantifier prefix	Coordinate
Payoff theorem consuming the row above: compactness of the achievement set forces $1/2$ itself to be achieved. Lean source [Lean] (theorem proved; *hypothesis* open)	scale:cofinal	given $\forall N\, \exists n \geq N(\ldots)$	coord:other:topological-achievement-set
Shortest known open target: the canonical rational half-greedy orbit itself has infinitely many positive skip events. Lean source [Open]	scale:cofinal	$\forall N\, \exists c \geq \max(N,4),\ 0 < \mathrm{rem}(c-1) < w_c$	coord:greedy-orbit
Payoff theorem consuming the row above, with zero further combinatorial search once supplied. Lean source [Lean] (theorem proved; hypothesis open)	scale:cofinal	given $\forall N\, \exists c \geq N(\ldots)$	coord:greedy-orbit
Master half-branch dichotomy: membership iff the greedy skip set is infinite. Lean source [Lean] (iff proved; both sides cofinal-shaped)	scale:cofinal	iff (skippedSupport).Infinite	coord:greedy-orbit
Seven-way classification hub: membership iff the seam word is not eventually right, iff unbounded terminal-false, iff cofinally many skipped ranks, *etc.* Lean source [Lean]	scale:cofinal	every right-hand side is $\exists^\infty$/Unbounded-shaped	coord:seam-integer
Ladder dichotomy: either infinitely many rungs survive (#257 false, explicit limit support) or an eventual bad-rank wall closes only this producer. (source: `erdos257_truncation_rung_ladder_2026_07_24.md`, Cor. 8) [Math]	scale:cofinal	open branch (a): $\exists^\infty J$, HalfRung(J)	coord:rung-truncation
The unconditional wall, stated in its four equivalent forms (run length, deviation, distance-to-integer, greedy support). (source: `erdos257_sqrt_escape_digit_reduction_2026_07_24.md`, §5) [Conjecture]/[Open]	scale:cofinal	$\forall r \geq 10$, $\lvert\mathrm{rem}(r+1) - 2^{r+1}\rvert > 2^{(r+5)/2}$	coord:seam-integer

Result	Scale	Quantifier prefix	Coordinate

Table 2: The scale ladder for Erdős #257. Every row above the frontier is machine-checked or exactly certified at fixed/bounded/uniform scale; every row below it is exactly the cofinal statement an open producer needs, with its quantifier prefix made explicit so the gap cannot be mistaken for a finite verification.

Remark 7.1 (Nothing promotes for free). A natural question, once the ladder is visible, is whether any scale:uniform or scale:bounded result already in the tree secretly *is* the cofinal statement in disguise — whether some existing theorem's hypothesis, read carefully enough, is uniform in a parameter that a cofinal witness sequence could simply walk through. An audit pass read each Lean proof body (not just the statement) of eighteen such candidates drawn from the fixed/bounded/uniform rows of Table 2 and their immediate neighbours, checking specifically for a proof technique general enough to instantiate at an unboundedly growing parameter without degradation. None promoted. The recurring blockers were three: a finite certificate table baked into the proof (the bound only extends as far as the table was built, *e.g.* rows 13–31 in Lean source); a kernel `interval_cases`/`decide` discharge over an explicitly bounded range, which is definitionally not re-runnable past its stated bound; and constants inside the bound that degrade with the very parameter a cofinal witness would need to grow (a threshold of the shape $2^{(r+5)/2}$ or a margin that shrinks as $1/\sqrt{N}$ is uniform *at* any fixed depth but does not, by itself, manufacture the next witness at a larger depth). This is consistent with, not a substitute for, the general logical point above: promoting a uniform statement to a cofinal one requires new content — an explicit witness-producing construction — not a change of quantifier bookkeeping.

8 One orbit in four coordinates

One decision stream, four integerisations

Fix the weights $x_n := 1/(2^n - 1)$, $n \geq 2$, and the target $1/2$. The single object under study throughout this programme is the greedy orbit for this target: a residual ρ starts at $1/2$, and at rank k the orbit *takes* iff $\rho \geq x_k$. A skip at rank k is *safe* iff $\rho \leq T_{k+1} := \sum_{j>k} x_j$ and *fatal* iff $\rho \in (T_{k+1}, x_k)$ — the fatal interval is exactly the gap no later tail can repair. The repo does not study four different problems; it measures this one orbit through four different integer coordinates, each certified to a stated depth.

- **Truncation rung J.** Replace the full weight x_n by the J-term truncation $w_n^{(J)} := \sum_{q=1}^{J} 2^{-qn}$ and ask whether some Boolean support hits $1/2$ exactly under the truncated weights. Certified for $3 \leq J \leq 22$ (Table 2, bounded row; Theorem 8.6–8.9 below).
- **Seam row s.** The base-4 seam recursion tracking the deviation $\Delta_s := \mathrm{rem}(s) - 2^s$ of the integer greedy orbit from its dyadic target. Certified rows 6 through 2500 for the reset-crossing classification (§7, uniform rows; source: `erdos257_reset_crossing_unification_2026_07_24.md`, §8), and the underlying real orbit sep-

arately certified to row 200,000 (source: `erdos257_sqrt_escape_digit_reduction_2026_07_24.md`, §3): exactly one TARGET violation in that range, at $r = 7$.

- **Integer margin m_c.** An advisory Type B scan of the integer-greedy margin out to 5×10^5, tracking how far the certified region sits from the danger threshold.
- **Sharp tail margin.** The rank-by-rank take/skip margin against the exact tail, checked for ranks 2 through 3000: 1497 takes against 1502 skips, zero fatal skips anywhere, minimum certified gap-margin +2.9922 bits at rank 5, and the single tightest real data point in the whole corpus at rank 7 — the near-tie $1/126 \geq 1/127$, margin +0.0340 bits (source: `erdos257_reset_crossing_unification_2026_07_24.md`, §8).

Proposition 8.1 (One orbit, four coordinates)**.** *Fix the weights $x_n = 1/(2^n - 1)$ and target $1/2$. Sufficiently fine truncations of any one of the four integerisations above reproduce the full greedy orbit's decisions exactly on any fixed finite prefix of ranks; consequently all four coordinates decide the* same *question, namely whether $1/2$ lies in the Mersenne achievement set.*

Mechanism. This is not an approximation claim to be taken on faith; it rests on the fact below. □

Lemma 8.2 (No-ties lemma)**.** *At every rank k of the full greedy orbit for target $1/2$, both defining comparisons are strict: $\rho \neq x_k$ and $\rho \neq T_{k+1}$.*

Proof. Two independent mechanisms, one for each boundary, both closing a boundary that a naive truncation argument would otherwise have to worry about degenerating on.

Take boundary excluded (parity). Any finite Mersenne sum $P/D := \sum_{d \in D} x_d$ has D odd, since every weight's denominator $2^d - 1$ is odd and a product/lcm of odd numbers is odd (equivalently, Lean source, itself traced to Lean source). Hence the residual after any finite prefix is $\rho = \frac{1}{2} - \frac{P}{D} = \frac{D-2P}{2D}$: its reduced denominator is *even*. Every weight $x_k = 1/(2^k - 1)$ has *odd* denominator. An even-denominator rational can never equal an odd-denominator rational, so $\rho = x_k$ is impossible at any rank — the take comparison is never a tie.

Skip-safety boundary excluded (irrationality). The tail past rank k is $T_{k+1} = E - 1 - \sum_{2 \leq j \leq k} x_j$, where $E = \sum_{n \geq 1} 1/(2^n - 1)$ is the Erdős–Borwein constant: a rational shift of E. Erdős's own 1948 theorem gives E irrational (Lean source, instantiated at $b = 2$), so T_{k+1} is irrational for every k. The residual ρ arising from any finite Boolean prefix is rational. An irrational number never equals a rational one, so $\rho = T_{k+1}$ is impossible at any rank — the skip-safety comparison is never a tie either.

Both boundaries of the fatal interval (T_{k+1}, x_k) are therefore approached only strictly, at every rank, unconditionally. □ □

Remark 8.3. Lemma 8.2 is exactly what licenses Proposition 8.1: because no comparison is ever exactly on a boundary, every rank's decision is stable under sufficiently fine truncation — there is no knife-edge case where refining J, extending s, tightening m_c, or sharpening the tail margin could flip a decision that a coarser coordinate had already settled on a fixed prefix. This is the sense in which the four bullet points above are one decision stream, not four independent producers that happen to agree so far.

The truncation-rung ladder

The cleanest of the four coordinates to reason about directly is the rung ladder, because at each finite J it is a genuinely finite, decidable question. Fix $J \geq 2$ and set

$$w_n^{(J)} := \sum_{q=1}^{J} 2^{-qn}, \qquad T_{n+1}^{(J)} := \sum_{q=1}^{J} \frac{2^{-qn}}{2^q - 1}, \qquad \mathrm{HalfRung}(J) := \exists\, A \subseteq \{2,3,\ldots\},\ \sum_{n \in A} w_n^{(J)} = \tfrac{1}{2}.$$

Every result in this subsection is [Math], proved in ordinary mathematics and cross-checked by exhaustive or randomised computation, *not* yet formalised in Lean.

Lemma 8.4 (Forced greedy)**.** *For every $J \geq 2$: $w_n^{(J)} > T_{n+1}^{(J)}$ for all n (the $q = 1$ terms agree exactly, and every $q \geq 2$ tail term is strictly smaller than the corresponding weight term). Consequently the weight sequence is strictly superincreasing, the greedy support is the unique candidate support, and* $\mathrm{HalfRung}(J)$ *holds iff the greedy orbit for $1/2$ under weights $w_n^{(J)}$ never lands in a fatal interval $(T_{n+1}^{(J)}, w_n^{(J)})$. Rank 1 is always a safe skip; ranks 2 and 3 are always takes.*

Lemma 8.5 (Parity forces infinite support)**.** *For every $J \geq 2$: no finite A attains* $\mathrm{HalfRung}(J)$*. Scaling by $2^{J \cdot \max A}$ makes the maximal element's contribution odd while every other element's contribution is even, so a support achieving $1/2$ exactly under J-truncated weights must be infinite. (This is the rung analogue of the parity exclusion in Lemma 8.2: the mechanism is the same shape, one level down.)*

Theorem 8.6 (Witness exclusion)**.** *Define the misalignment mass*

$$\mu_J(M) := \sum_{q=2}^{J} \frac{2^{M \bmod q}}{2^q - 1}.$$

Let $J \geq 3$, $n \geq 4$, and suppose some $M \in [n, 2n-2]$ has $\mu_J(M) \leq \frac{11}{15}$. Then no Boolean prefix $D \subseteq \{2, \ldots, n-1\}$ satisfies

$$T_{n+1}^{(J)} < \tfrac{1}{2} - \sum_{d \in D} w_d^{(J)} < w_n^{(J)};$$

fatality at rank n is impossible for every prefix simultaneously (not merely the greedy one). The proof mechanism is a 2-adic scaling and divisibility argument at the witness index M; it is self-contained and general beyond this instance (§7, TR-2 chains-with note: it is itself a scoped instance of the missing near-integer anti-concentration lemma that both open problems ultimately need).

Corollary 8.7 (Half-LCM horizon)**.** *Let $L_J := \mathrm{lcm}(2, 3, \ldots, J)$ and let M be the least multiple of L_J with $M \geq n$. Then $M \leq 2n - 2$ whenever $n \geq L_J/2 + 1$, and at that M,*

$$\mu_J(M) = \sum_{q=2}^{J} \frac{1}{2^q - 1} < E - 1 < 0.6067 < \frac{11}{15}$$

(where E is the Erdős–Borwein constant). Hence, by Theorem 8.6, every rank $n \geq L_J/2 + 1$ is automatically witness-covered: the remaining possibly-bad ranks form a finite window $[4, L_J/2]$, not a cofinite complement to be checked one at a time.

Lemma 8.8 (Mod-12 filter, $J \geq 7$)**.** *For $J \geq 7$: $\mu_J(M) \leq \frac{11}{15} \Rightarrow 12 \mid M$ (case analysis on M mod 4 and M mod 3). Combined with Corollary 8.7, the search for possibly-bad ranks inside $[4, L_J/2]$ can be restricted to multiples of 12, a 12× speedup with no loss of coverage.*

Theorem 8.9 (Finite decision procedure)**.** *Call $n \in [4, L_J/2]$ **bad** if no $M \in [n, 2n-2]$ has $\mu_J(M) \leq \frac{11}{15}$, and set $B(J) := \max(\text{bad} \cup \{3\})$. Then* HalfRung$(J)$ *holds iff the greedy orbit for $1/2$ under weights $w_n^{(J)}$ survives every rank from 2 through $B(J)$ — a finite, exact, decidable check, with $B(J)$ observed to compress the certificate window by up to $\sim$ 115,000× relative to the naive $L_J/2$ bound* (e.g. $J = 19$: $B(19) = 1008$ *against* $L_{19}/2 = 116{,}396{,}280$).

Remark 8.10 (Certified rungs and the open producer). Every rung $3 \leq J \leq 22$ survives this finite check with zero fatal ranks encountered — these are proved theorems HalfRung$(3), \ldots,$ HalfRung(22), not merely computational evidence, each an independently bankable finite fact ([Math]+[Cert]; Table 2, bounded row). By the compactness-transfer theorem of the same source document (rung solutions lie within $(2/3)4^{-J}$ of $1/2$ in full-Mersenne value, so a convergent subsequence under the compactness of $\{0,1\}^{\{2,3,\ldots\}}$ has limit exactly $1/2$, and that limit's support is infinite since finite sums have odd denominator by Lemma 8.2/8.5), a single further success at *infinitely many* J would already disprove Erdős #257. This is exactly the frontier row of Table 2: $(*)$ "HalfRung(J) holds for infinitely many J" is the open, cofinal statement; twenty individually-proved finite instances, however many, do not by themselves supply it.

9 Interface index: what stands between the corpus and a proof

This section is the sharp end of the paper. Parts I–II assembled the corpus; here every result in it is tested against three exact open obligations and rendered as a usable premise: what it yields, the exact mismatch against the obligation, a classified gap kind (`hypothesis_strength`, `scale_only`, `quantifier_order`, `coordinate_only`, or `multiple`), and the exact auxiliary statement that would close it. Nothing here decides #249 or #257; the corpus's own moderator standard applies throughout — a reduction is not a result, and a finite verified list is not a cofinal supply.

9.1 The three #257 obligations, stated exactly

Definition 9.1 (257-reset)**.** For every reset row $r \geq 31$: $\Delta_{r+1}^2 > 2^{r+5}$, where Δ_{r+1} is the signed deviation of the greedy remainder from 2^{r+1} at the row immediately following reset r. Equivalently: the R-run (all-right run) beginning at reset r has length $L_r < (r-3)/2$. Equivalently, at the natural dyadic scale 2^{-n}, the greedy skip-set sum never approximates $C = E - 3/2$ to within about $2^{-1.5n}$ (here E is the Erdős–Borwein constant). Any bound $L_r = o(r)$ closes the obligation. A *uniform* constant bound on L_r is FALSE — runs grow like $\log_2(\text{row})$ — so the correct target is a growing, not constant, run-length ceiling.

Definition 9.2 (257-cofinal-rows)**.** $\forall N\, \exists n \geq N$ with Lean source n — i.e. cofinally many finite prefixes whose local Mersenne quotient lands exactly on the half-value residue class. Consumed unconditionally by Lean source. No coherence between the witnessing rows is demanded and none should be assumed: the rows may be, and in the corpus's own words, mutually incompatible — only their endpoint lengths need tend to

infinity. Any weakening that reintroduces a compatibility requirement between rows is not a weakening of this obligation; it is a different, harder one.

Definition 9.3 (257-universal)**.** $\forall b \geq 2\ \forall A \subseteq \mathbb{N}$ infinite: $\sum_{n\in A} 1/(b^n - 1)$ is irrational. This is Erdős #257 at full strength — every infinite support, not one target value and not one structural class.

9.2 Reading a row

Each entry below has the form: **tag — declarations**, `site` (file:line, opened and read, not inferred), `gap_kind`, **yields** (the exact statement, quantifiers preserved), **mismatch** (the precise distance from the obligation, with the mechanism), **closes** (the exact auxiliary statement — not a vague direction — that would discharge the row). Evidence band is [Lean] unless marked [Math] or [Cert] (exact finite computation).

9.3 257-reset: sixteen near misses

Observation 9.4 (oa-2 / deficit-side run law). [Lean] Lean source. **Yields:** $\forall k, s.\ 5 \leq s \to \mathrm{rem}(s) \leq 2^s \to 2^s \leq 2^k(2^s - \mathrm{rem}(s)) \to \exists t \in [s, s+k]$ with an upper-or-middle event at t. Contrapositive: an all-right run of length k from row s forces deficit $2^s - \mathrm{rem}(s) < 2^{s-k}$. Uniform in *both* k and s; nothing scale-restricted here. **Mismatch (`hypothesis_strength`):** the implication side is already at cofinal, uniform strength. What is missing is the deficit *lower* bound fed into it. The only unconditional deficit bound on hand is integrality ($\mathrm{rem}(s) < 2^s \Rightarrow$ deficit ≥ 1); instantiating $k = s$ with deficit $\geq 2^0$ gives only $L_r \leq r$. The obligation needs $L_r < (r-3)/2$ — short by exactly a factor of two in the exponent. **Closes it:** $2^{(r+3)/2} \leq 2^r - \mathrm{rem}(r)$ at every reset $r \geq 31$. Nothing else in this theorem changes; the induction already consumes any k uniformly.

Observation 9.5 (rc-2 / affine excess recurrence). [Lean] Lean source, Lean source, Lean source (recurrence Lean source). **Yields:** for an R-run of length k from row S: $X_k + \mathrm{charge}_k = 4^k X_0$ exactly, with $X_j = \mathrm{rem}(S+j) - 2^{S+j}$ and $\mathrm{charge}_{j+1} = 4\,\mathrm{charge}_j + \mathrm{pulse}_j + 4$. Uniform in k, S, and the pulse word. **Mismatch (`multiple`):** this is the excess-side twin of the deficit run law but only half-assembled: $X_0 \leq (X_k + \mathrm{charge}_k)/4^k$ needs an upper bound on the terminal excess X_k before it yields a run-length statement. Two candidate upper bounds exist elsewhere in the corpus (the unconditional late-row ceiling below, and the open `SeamGreedyRemainderGapBound` socket) but neither is composed with this iterate anywhere in the tree. **Closes it:** one composition lemma — if the R-run from S has length k and every row in it satisfies the late-row condition $2s < 3d$ (equivalently $\mathrm{rem}(s) < 2^{s+1}$), then $\mathrm{rem}(S) - 2^S \leq 2^{S-k} + (2S+4)/3$. That would make the run law two-sided; the anti-concentration input is still separate.

Observation 9.6 (oa-2 sibling / late-row ceiling). [Lean] Lean source. **Yields:** $\forall s \geq 5$, at a largest false rank d with $2s < 3d$ (late row): $3\,\mathrm{rem}(s) < 3 \cdot 2^{s+1} + 2 \cdot 4^{s-d} + 4$. Unconditional, uniform in s, d; since late means $s - d < s/3$, the error term is $< 2^{2s/3}$. **Mismatch (`coordinate_only`):** an upper bound on the deviation where the obligation needs a lower bound on its absolute value — the excess-side ceiling the rc-2 run law needs, not anti-concentration itself. Degrades to vacuous outside the late-row regime. **Closes it:** nothing on its own; its role is discharging the "X_k bounded above" premise of rc-2 without invoking the open `SeamGreedyRemainderGapBound`.

Observation 9.7 (G5 / no-tie). [Lean] Lean source, with Lean source. **Yields:** $\forall s \geq 5$, the greedy decision boundary is never exactly attained. Combined with integrality this is the corpus's *only* unconditional lower bound on the deviation: $|\Delta_r| \geq 1$. **Mismatch (`scale_only`):** deviation $\geq 2^0$ versus the required deviation $> 2^{(r+5)/2}$. This is the honest state of the anti-concentration axis: the entire gap is the exponent, from 0 to $(r+5)/2$, with nothing proved in between. **Closes it:** any growing lower bound $|\mathrm{rem}(r) - 2^r| \gtrsim 2^{\varepsilon r}$ for fixed $\varepsilon > 0$ would be the first non-constant point on this axis; $\varepsilon = 1/2 + o(1)$ closes the obligation via the run law.

Observation 9.8 (H4, PRIVATE / sharp fatal-gap test). [Lean] Lean source, Lean source, Lean source, Lean source. **Yields:** writing the residual in lowest terms $\rho = u/(2L)$, u, L odd, $a := 2L - (2^k - 1)u$, the sharp safety test is $2u \leq 3a$ (strictly better than the old dyadic test $u \leq a$). Contrapositive: fatal $\Rightarrow 3a < 2u \Rightarrow u \geq 2$, and $u \geq 3$ when u is odd. Uniform in $k \geq 1$. **Mismatch (`scale_only`):** same shape as the obligation — a numerator lower bound at a dangerous rank — but the bound is the constant 3 where growth $2^{(r+5)/2}$ is needed. The constant 3 is the Erdős–Borwein tail-ratio limit $\theta_k := 1/T(k+1) - (2^k - 1) < 2/3$; adding Lambert channels sharpens 2/3 toward the true limit but never produces growth in k. **Closes it:** a rank-scaling version, e.g. fatal at rank $k \Rightarrow u \geq 2^{k/2}$. The current mechanism provably cannot deliver this — the ceiling is the tail ratio itself — so this needs a different arithmetic input, not a sharper truncation.

Observation 9.9 (F5b / linear-width band escape, rows 13–30). [Lean] Lean source. **Yields:** at every upper reset row d, $13 \leq d \leq 30$, the reset charge $E_d = 4 \cdot \mathrm{overshoot}_d + \mathrm{abovePulse}_d$ avoids the linear-width band $(2^{d-j+1} - 2(d+j),\ 2^{d-j+1}]$ for every $j \leq d$ — the finite segment of Lean source, which by Lean source delivers the whole #257-false payoff. **Mismatch (`scale_only`):** proved at 18 explicit rows by `interval_cases` plus per-row `decide+kernel`; the obligation needs all $d \geq 13$. The band width here is only $O(d)$ — *linear* — while sqrt-escape demands width $2^{(r+5)/2}$; this socket is exponentially weaker than 257-reset yet lands the same downstream conclusion. **Closes it:** a structural lower bound on the distance from E_d to the nearest dyadic power at or above it, of width $2(d + j_0)$ at the single critical index j_0 — already collapsed from $\forall j$ to one inequality by Lean source. Only $O(d)$ separation is required, not sqrt-of-scale.

Observation 9.10 (oa-3b / M-ancestor and R-run closure). [Lean] Lean source, Lean source, Lean source. **Yields:** unconditional, uniform in $s \geq 5$: a middle (M) branch forces $s + 1 \leq \mathrm{rem}(s+1)$, using no lower bound on the source remainder; with the row-scale source bound this upgrades to $2^{s+1} + (s+1) \leq \mathrm{rem}(s+1)$. The minimal-counterexample argument at line 3606 then shows any row-small row must have an upper-reset (U) ancestor, since M-ancestors and R-runs are both excluded. **Mismatch (`hypothesis_strength`):** two of the three ancestor channels (M, R) are closed unconditionally at all scales; only the U channel remains, and it needs only the $O(d)$-wide band escape of F5b above — sqrt-of-dyadic-scale escape is being demanded by the 257-reset statement where linear-width escape at U-resets alone suffices for this downstream route. **Closes it:** the single remaining channel — $\forall d \geq 13$ with successor carries, the reset charge avoids the width-$2(d + j_0)$ band at its critical dyadic index. Then `remainder_ge_row`, hence `SeamMiddleProducerRowEscape`, hence $1/2 \in$ mersenneAchievementSet, all already in Lean.

Observation 9.11 (F2 / middle producer row escape). [Lean] Lean source, Lean source (via `.toCardEscape`, `.toTailEscape`, Lean source). **Yields:** if at every late middle row ($s \geq 13$, no carry, middle branch) $s \leq \mathrm{rem}(s)$, then $1/2 \in$ mersenneAchievementSet. Module docstring: "far weaker than an exponential bound, but stronger than the already exposed square-root-scale carry target." **Mismatch (`hypothesis_strength`):** this socket asks for a *linear-in-row* lower bound on the remainder itself; the obligation supplies an *exponential* lower bound on the deviation from 2^s at reset rows. Neither implies the other — sqrt-escape is compatible with $\mathrm{rem}(s) = 0$, which violates $s \leq \mathrm{rem}(s)$. A genuinely cheaper, incomparable socket on the same payoff. **Closes it:** $\forall s \geq 13$ at a middle transition, $s \leq \mathrm{rem}(s)$. Equivalently, by `remainder_ge_row`, the U-reset band escape of F5b.

Observation 9.12 (LPC-sqrt / row-index sqrt escape). [Lean] Lean source, `.toTailEscapeHalfCylinderLastProducerContradiction.lean:98`, Lean source. **Yields:** if at every late middle row $2\sqrt{2s+2} + 4 < \mathrm{producerCarry}(s)$, then $1/2 \in$ mersenneAchievementSet; `producerCarry` equals the middle coordinate $4\,\mathrm{rem} - \mathrm{belowPulse} - 4$ (Lean source), so the requirement is essentially $\mathrm{rem}(s) \gtrsim (\sqrt{s} + \mathrm{belowPulse})/4$. **Mismatch (`coordinate_only`):** this is "square root" of the *row index* s (a supportCoeff tail majorant), not square root of the *dyadic scale* $2^{(r+5)/2}$. The two sqrt's in this corpus are exponentially far apart and easy to conflate; this socket is the weaker one by an exponential margin, and it is already the theorem-ready form of the last-producer contradiction. **Closes it:** $\forall s \geq 13$ at a middle transition, $4\,\mathrm{rem}(s) - \mathrm{belowPulse}(s) - 4 > 2\sqrt{2s+2} + 4$ — a row-scale, not dyadic-scale, lower bound on the middle coordinate.

Observation 9.13 (G2 + LPC upper case / narrow integer window). [Lean] Lean source, Lean source, Lean source, Lean source, Lean source. **Yields:** the upper-producer "last transition before an all-right tail" case is impossible unconditionally (affine carry ≤ -8); the middle cell at coordinate -3 is excluded unconditionally. What survives: a middle producer whose landing excess X_{D+1} satisfies $2 \leq X_{D+1} < 2\sqrt{2D+2} + 8$ — a window of only $O(\sqrt{D})$ integers. **Mismatch (`hypothesis_strength`):** the residual bad set is an explicit $O(\sqrt{D})$-wide integer window at each candidate last-producer row, not an interval at dyadic scale. The obligation would evacuate this window wholesale, but excluding an $O(\sqrt{D})$ window is exponentially cheaper than proving $|\Delta| > 2^{(r+5)/2}$. **Closes it:** exclusion of the integer window $[2,\ 2\sqrt{2D+2}+7]$ for the landing excess at any hypothetical last middle producer $D \geq 13$ — e.g. a congruence or pulse-parity obstruction on $4\,\mathrm{rem}(D) - \mathrm{belowPulse}(D) - 4$.

Observation 9.14 (F4 / two-sided cell exclusion). [Lean] Lean source, `.stepHalfCylinderMiddleCarryLowerI` `.twoSidedHalfCylinderMiddleCarryLowerBound.lean:3413`. **Yields:** excluding exactly three integer values of the middle coordinate ($4\,\mathrm{rem} - \mathrm{belowPulse} - 4 \notin \{-3, -2, -1\}$) plus one right-pulse-leak bound propagates, by induction from a decide-checked row-5 base, the two-sided invariant $\forall s \geq 5$: $\mathrm{rem}(s) \leq 2^s \vee \mathrm{overshoot}(s) \leq 2^s$. **Mismatch (`hypothesis_strength`):** the hypothesis is anti-concentration at *constant* resolution — the orbit must miss three named integers — where the obligation is anti-concentration at resolution $2^{(r+5)/2}$. One of the three cells (-3) is already closed in the all-right-tail

configuration by `finalMiddleCell_neg_three_not_last`, leaving two. **Closes it:** exclusion of the two remaining cells (-2 and -1) for the actual seam orbit, plus the right-pulse-leak inequality $4 \cdot \text{overshoot} + \text{abovePulse} \le 2^{s+2}$ under overshoot $\le 2^s$.

Observation 9.15 (E1 / remainder ceiling). [Lean] LEAN SOURCE, LEAN SOURCE (weaker cofinal route LEAN SOURCE). **Yields:** if $\forall s \ge 6$, $\text{rem}(s) < 2^{s+1}$, then $1/2 \in$ mersenneAchievementSet. The weaker cofinal form only needs the normalized remainder to tend to 0 along *some* cofinal row sequence. Docstring: "exponential upper bound observed by exact replay from row six onward." **Mismatch (`coordinate_only`):** an upper bound on the remainder rather than a lower bound on |deviation| — the ceiling half of the picture, and exactly the premise rc-2's excess-side run law is missing; `three_mul_remainder_lt_exactLateGap` already supplies it unconditionally at late rows up to a $2^{2s/3}$ correction. **Closes it:** either $\forall s \ge 6$, $\text{rem}(s) < 2^{s+1}$ outright, or a proof that the late-row condition $2s < 3d$ holds at every row so the unconditional strip invariant covers it.

Observation 9.16 (SE-7 / RC-7 / certified computation to row 200,000). [Cert] **Yields:** exact integer arithmetic to row 200,000 — branch counts $R : M : U = 100197 : 49899 : 49898$; the obligation *fails only at* $r = 7$; maximum R-run 19 at row 158,096 against the required threshold $\approx (r-3)/2 \approx 79{,}000$ (margin $\sim 4000\times$). Independently: 1209 resets checked to row 2500, zero classification anomalies, exactly two crossing cells ever (s=7, d=5 and s=10, d=7), both below socket scope. **Mismatch (`scale_only`):** fixed-scale exhaustive verification versus a statement at every reset $r \ge 31$. This is the canonical quantifier trap this programme warns against — a finite verified list is not a cofinal supply, however large the margin. **Closes it:** nothing computational; only a theorem. The certificate's value is fixing the true constant ($r \ge 31$, sole violation $r = 7$) so the theorem's scope can be stated exactly.

Observation 9.17 (RC-5 / rigidity of dangerous resets — Theorem B, see below). [Math] **Yields:** at a dangerous reset r ($|\text{rem}(r+1) - 2^{r+1}| \le 2^{(r+5)/2}$) whose preceding reset r_0 has pure R-run length $L' = r - r_0 - 1 > (r+5)/4$, the deviation w_{r_0+1} is pinned to at most two adjacent integers determined by the divisor-pulse stream alone; a chain of n dangerous resets forces $n-1$ nested exact integer-coincidence towers. **Mismatch (`hypothesis_strength`):** applies only to dangerous resets preceded by a long pure R-run ($L' > (r+5)/4$) — an isolated dangerous reset after a short run is untouched. Pinning to two integers is not yet a contradiction: no theorem excludes those two values. Prose, not Lean. **Closes it:** (i) a congruence or pulse-stream obstruction showing the two pinned integer values are unattainable, and (ii) removal of the long-preceding-run hypothesis, or a proof that dangerous resets must come in chains. Either alone converts Theorem B into a genuine no-go.

Observation 9.18 (RC-8 / sign law and margin growth). [Cert] **Yields:** at all 1209 observed resets, M ⇒ deviation > 0 and U ⇒ deviation < 0, with zero exceptions; the margin over the $2^{(r+5)/2}$ threshold grows from 1.11 bits at $r = 14$ to ~1246 bits at $r = 2500$. The corpus flags this as a strictly weaker, possibly more tractable target than full sqrt-escape. **Mismatch (`multiple`):** two gaps at once — the sign law is empirical only (certified computation, not proved unconditionally as a lemma), and even if proved it delivers the *sign* of the deviation, not its magnitude, while the obligation needs magnitude.

Closes it: prove the sign law as a lemma (M-branch $\Rightarrow$ $\mathrm{rem}(r{+}1) > 2^{r+1}$; U-branch $\Rightarrow$ $\mathrm{rem}(r{+}1) < 2^{r+1}$) — a one-step branch-algebra statement in the coordinates of `rightBranch_excess_succ_eq` / `seamUpperBranch_nextRemainder_le_pow`. That converts the two-sided obligation into two one-sided bounds, each in a branch class where the exact affine recurrence is already known.

Observation 9.19 (TR-5 / TR-6 / truncation-rung ladder, parallel coordinate). [Cert] **Yields:** HalfRung(J) genuinely *proved* (not merely computed) for each $3 \le J \le 22$, via a finite exact decision procedure: every rank $n \ge L_J/2{+}1$ is automatically witness-covered, a mod-12 filter cuts the residual search 12×, so only ranks in $[4, B(J)]$ need checking. By the corpus's compactness transfer theorem, HalfRung(J) for infinitely many J gives an infinite A with $\Sigma = 1/2$, i.e. #257 FALSE. **Mismatch (`multiple`):** different coordinate (truncation rung J, not seam reset row r) *and* fixed scale (20 explicit values of J). Reaches the same downstream payoff without touching the deviation at all — a parallel route whose own cofinal step is open, not one upgrade away from 257-reset. **Closes it:** HalfRung(J) for infinitely many J — explicitly *not* promotable by re-running the existing argument: the certificate window $B(J)$ is driven by $L_J = \mathrm{lcm}(2, \ldots, J)$, growing superexponentially in J, the hallmark of a proof that does not survive bound removal.

9.4 257-cofinal-rows: thirteen near misses

Observation 9.20 (Sharp-capacity fill, unconditional). [Lean] Lean source, Lean source. **Yields:** for any $c \ge 4$ and any finite $D \subseteq [2, c)$ with localMersennePrefixValue $D < 1/2$ and localBinarySuffix $D\,1\,(2c{-}2) < 2^{c-2}$: an exact row E at endpoint $2c{-}2$ with $D \subseteq E$, every new rank strictly above c. Unconditional, no deficit/crossing hypothesis, no reference to the greedy orbit. **Mismatch (`hypothesis_strength`):** needs a per-c input localBinarySuffix $D\,1\,(2c{-}2) < 2^{c-2}$, and the corpus only ever supplies that input through the crossing-deficit route, which Lean source then rigidly collapses onto $D =$ halfGreedyPrefixSupport$(c{-}1)$. The theorem itself is strictly more general than every consumer that calls it. **Closes it ($\star$):** $\forall N\ \exists c \ge N$, $\exists$ finite $D \subseteq [2, c)$ with localMersennePrefixValue $D < 1/2$ and localBinarySuffix $D\,1\,(2c{-}2) < 2^{c-2}$ — equivalently mobiusCenteredHalfCarry $\uparrow D\,(2c{-}3) < 2^{c-2}$. No coherence between the D's, no requirement that D be a greedy prefix.

Observation 9.21 (Sharp-capacity induction, engine complete). [Lean] Lean source → Lean source → Lean source (Lean source; seed line 1023; strict progress Lean source). **Yields:** a complete, already-formalised induction from the endpoint-six seed to the full obligation: each protected row either doubles below half unconditionally, or recycles at its first crossing rank e, giving endpoint $2e{-}2 >$ previous endpoint. Terminates literally at `CofinalExactLocalMersenneHalfRows`. **Mismatch (`hypothesis_strength`):** the recycle branch consumes `SkippedCoreCriticalQuotientSupply` at essentially every step (the below-half branch never fires twice, by the seed arithmetic). Protection (endpoint $<$ 2·cutoff, new ranks $>$ cutoff) is exactly what converts a non-growing recycle endpoint into strict progress; everything else is finished. **Closes it:** the sharp inequality at one crossing rank per step, canonical form Lean source: $\forall c \ge 4$ skipped by the rational half-greedy orbit, $2^{(2c-2)-1} \le$ localPrefixQuotient(insert c (halfGreedyPrefixSupport$(c{-}1)$)) $(2c{-}2)$.

Observation 9.22 (Unconditional capacity, off by one bit). [Lean] Lean source (word form line 228). **Yields:** unconditionally, $\forall c \geq 4$ and every below-half core $D \subseteq [2, c)$ with deficit $<$ mersenneWeightRat c: localBinarySuffix $D\,1\,(2c{-}2) < 2^{c-1}$. Uniform in c, no table, no case split. **Mismatch (`scale_only`):** one bit short. The strict-upper fill needs $< 2^{c-2}$, this gives $< 2^{c-1}$. The proof derives $A < 2^{c-2}+|D|$ then discards $|D| \leq c{-}2 \leq 2^{c-2}$ — the loss is purely additive. The true unproved statement is only that A avoids the linear-width band $[2^{c-2}, 2^{c-2} + c - 3]$, width $c{-}2$ inside a range of size 2^{c-2}. **Closes it:** a linear-width dyadic-band anti-concentration: at cofinally many crossing ranks c, localBinarySuffix $D\,1\,(2c{-}2) \notin [2^{c-2}, 2^{c-2} + c - 3]$. Same band shape as F5b/oa-3b above; exponentially weaker than 257-reset's $2^{(r+5)/2}$ target, and the two sockets have never been compared on band width in any prior bank.

Observation 9.23 (Doubling extension, fails at the only landed seed). [Lean] Lean source, line 296; rank-two discharge Lean source, line 108. **Yields:** $\forall n \geq 6$: a below-half exact row at n produces an exact row at $2n{-}1$, support-extending, all new ranks $> n$. The only unconditional endpoint-doubling producer in the corpus. **Mismatch (`hypothesis_strength`):** below-halfness is not preserved and demonstrably fails at the first step. For the seed $D = \{2, 3, 6\}$ at $M = 11$ the core alone contributes $2/3 + 4/7 + 32/63 \approx 1.746 > 1$, so the doubled row is strictly *above* half. Iterating the cheap arm is therefore impossible from the only landed seed; the chain is forced into the capacity-gated recycle arm at step one. **Closes it:** cofinally many below-half exact rows — $\forall N\ \exists n \geq N\ \exists D$ exact at n with localFractionMass $D\,n < 1$. A single such row at each of cofinally many n makes doubling unnecessary; a self-reproducing one closes everything.

Observation 9.24 (Consecutive-skip precritical bound). [Lean] Lean source (feeds Lean source, then lines 970/986). **Yields:** unconditional: if $c \geq 6$ and both rank c and rank $c{+}1$ are skipped by the rational half-greedy orbit, then localBinarySuffix(halfGreedyPrefixSupport$(c{-}1))\,1\,(2c{-}3) < 2^{c-3}$, doubling into sharp capacity at c, hence an exact row at $2c{-}2$. No supply hypothesis anywhere. **Mismatch (`hypothesis_strength`):** needs two *consecutive* skipped ranks. The residual open set is exactly the skip-then-take rows: Lean source proves the whole socket equivalent to Lean source, with $c = 4, 5$ already discharged by `decide`. **Closes it:** cofinally many $c \geq 6$ at which two consecutive ranks $c, c{+}1$ are both skipped by the rational half-greedy orbit — or dually, the pre-take residual at cofinally many skipped-then-taken ranks. Either one closes the obligation with no further input.

Observation 9.25 (Bounded-gap precritical bound). [Lean] Lean source (with Lean source, Lean source). **Yields:** unconditional and uniform in *both* parameters: skip at c, takes at $c{+}1, \ldots, c{+}t{-}1$, skip at $c{+}t$, with $0 < t \leq c - 3$ and $c - 2 \leq 2^{c-t-3} \Rightarrow$ precritical suffix bound at $c \Rightarrow$ sharp capacity $\Rightarrow$ exact row at $2c{-}2$. Handles any skip gap $t \lesssim c - 3 - \log_2(c - 2)$, not just $t = 1$. **Mismatch (`hypothesis_strength`):** the only missing input is an orbit-level skip-gap bound — that the next skipped rank after c occurs at $c{+}t$ with $c - 2 \leq 2^{c-t-3}$. The proof is fully uniform in (c, t), every constant explicit, no table. **Closes it:** for cofinally many skipped ranks c of the rational half-greedy orbit, the next skipped rank lies at most $c + (c - 3 - \lceil \log_2(c - 2) \rceil)$ away. Empirically skips have roughly constant density (SE-7: certified to row 200,000), so this is an enormously slack requirement, never isolated as a target anywhere prior to this bank.

Observation 9.26 (Frozen margin, ineffective horizon). [Lean] Lean source, Lean source (margin def line 1021, recurrence line 1036, monotonicity line 1094). **Yields:** at every dyadically safe rank $k > 0$, greedyHalfFrozenMargin $k\,J \geq 0$ for *some* horizon J — the finite coefficient window eventually covers the centred carry. Nonnegativity is monotone up in J. **Mismatch (`scale_only`):** the horizon comes from a non-effective limit argument (`tendsto_finiteCoeffWindow_atTop` plus an existence extraction), so J is unbounded. The exact-row route needs the crossing by the specific horizon $J = c - 3$ (Lean source, $k = c - 1$). All ingredients to make it effective already exist: Lean source gives gap $=$ tail$(k{+}1{+}J)/2^J$, and Lean source gives tail$(m) \leq m + 2$. **Closes it:** margin$(k, J) \geq 0$ as soon as $(k{+}J{+}3)/2^J <$ |halfGreedyNextDyadicExcessNumerator k|/halfGreedyPrefixDenominator k. Setting $J = k - 2$ turns the socket into one explicit dyadic-safety margin $|E_k|/D_k > (2k{+}1)/2^{k-2}$ at every skipped rank. Pure bookkeeping over existing lemmas.

Observation 9.27 (G4 / full-shell margin, wrong horizon direction). [Lean] Lean source, Lean source, Lean source. **Yields:** at every skipped rank $n \geq 3$, $0 \leq$ greedyHalfFrozenMargin$(n{-}1)\,n$ — first passage of the frozen margin by the full-shell horizon $J = k{+}1$, $k = n{-}1$. Reaches $1/2 \in$ mersenneAchievementSet by its own consumer. **Mismatch (`quantifier_order`):** exactly three doubling steps of horizon short. The exact-row socket needs margin$(k, k{-}2) \geq 0$; this gives margin$(k, k{+}1) \geq 0$. By the succ-recurrence, margin$(k, k{+}1) = 8{\cdot}$margin$(k, k{-}2) + 4\,\mathrm{sc}(2k) + 2\,\mathrm{sc}(2k{+}1) + \mathrm{sc}(2k{+}2)$, sc a divisor-count supportCoeff $\leq \tau \leq 2\sqrt{2k+2}$. Since nonnegativity is monotone *up* in J, the landed statement is the weaker one and cannot supply the obligation as it stands. **Closes it:** a horizon-tightening — at every skipped rank, margin$(k, k{+}1) \geq 4\,\mathrm{sc}(2k) + 2\,\mathrm{sc}(2k{+}1) + \mathrm{sc}(2k{+}2)$: the full-shell margin must exceed an explicit $O(\sqrt{k})$ divisor-count quantity. Same $\sqrt{}$-scale currency as the LPC-sqrt row above.

Observation 9.28 (Cofinal positive skips *is* the obligation). [Lean] Lean source, Lean source, line 60; A5 Lean source, A6 (Lean source); positivity discharge Lean source. **Yields:** cofinally many positive rational-greedy skips ⇒ `CofinalExactLocalMersenneHalfRows` ⇒ $1/2 \in$ mersenneAchievementSet. Advertised elsewhere as the shortest known path. **Mismatch (`hypothesis_strength`):** it is not weaker than the problem — it *is* the problem. The positivity conjunct $0 <$ greedyMersenneRemainderRat$(1/2)(c{-}1)$ is unconditionally dischargeable by the odd-denominator parity fact, so CofinalPositiveHalfGreedySkips ⇔ (greedyMersenneSkippedSupport$(1/2)$).Infinite, which A6 proves *equivalent* to $1/2 \in$ mersenneAchievementSet. Attacking this target is attacking #257-false head on, with no reduction whatsoever. **Closes it:** nothing short of the problem. *This is a trap, not a waypoint* — the escape is to abandon the greedy-orbit arm (where the wall at line 50 forces canonicity) and target the ($\star$) sharp-capacity-fill route instead, whose cores are not forced to be greedy prefixes.

Observation 9.29 (C4a / terminal-only strip, wrong locality). [Lean] Lean source, Lean source, Lean source. **Yields:** cofinally many depths M carrying a finite normalized word a with |integerHalfCarry(wordSupport a)$(M{-}1)| \leq$ halfStripBound M ⇒ $\exists A$ infinite with erdosSupportSeries $2\,A = 1/2$. Same architecture as the obligation: cofinal,

no coherence, mutually incompatible witnesses allowed. **Mismatch (`multiple`):** the two sockets trade off in opposite directions and neither implies the other. C4a tolerates *any* support inside the depth-M window but demands the carry inside a $\sim 2\sqrt{M}$ strip; the obligation form $(\star)$ tolerates a carry up to $2^{M/2-1}$ — exponentially looser — but demands the support be confined to ranks $< M/2+1$, since the exact fill encodes the residue in the pure upper window where the Mersenne quotient is exactly 2^{M-d}. The gap is support *locality*, not carry size. **Closes it:** a lower-half confinement upgrade of the strip witnesses — $\forall N\ \exists c \geq N$ with a strip-scale witness whose support lies in $[2, c)$ at depth $2c-3$. Then $2\sqrt{2c}+4 < 2^{c-2}$ for $c \geq 8$ and the existing fill fires verbatim. This identifies C4a as the *cheaper* route to the same payoff — 257-cofinal-rows is not the cheapest cofinal socket on the board.

Observation 9.30 (C4c / G3 / greedy carry vs. frozen prefix carry, coordinate mismatch). [Lean] Lean source, Lean source, line 157; Lean source, line 805 (nonnegativity from line 842 / Lean source). **Yields:** cofinal returns of the actual greedy carry to the square-root strip (C4c), or a pointwise $2\sqrt{N}+4$ bound on the Möbius-centred carry (G3), each give an infinite support with value exactly $1/2$. **Mismatch (`coordinate_only`):** these bound the carry of the *infinite* greedy support greedyMersenneSupport$(1/2)$; sharp capacity is about the *frozen finite prefix* halfGreedyPrefixSupport$(c-1)$. The two carries agree only at index k (inside Lean source), not at depth $2c-3$ where the fill needs it. The bridge is already landed: Lean source: sharp capacity $\Leftrightarrow$ mobiusCenteredHalfCarry $G\,(2c-4) \leq$ futureSkipCapacity $G\,c\,(c-3)$ — transport costs exactly the future-skip capacity term. **Closes it:** either a strip bound stated for the frozen prefix rather than the live orbit, or futureSkipCapacity $G\,c\,(c-3) \geq$ halfStripBound$(2c-3)$ — holds as soon as one skip occurs early in $[c+1, 2c-4]$, since the capacity is dyadically weighted. Same input as the bounded-gap precritical bound above.

Observation 9.31 (Above-support recycling, no growth guarantee). [Lean] Lean source, Lean source, Lean source. **Yields:** unconditional: every finite above-half support bounded by n recycles through its first crossing rank c into a genuine exact row at endpoint $2c-2$, with $4 \leq c \leq n$. No capacity hypothesis at all. **Mismatch (`quantifier_order`):** the witness is $\exists c \leq n$, not $\exists c$ large — $2c-2$ may be $\leq n$, so the endpoint need not grow. Lean source is the explicit falsifier — boundedDoubleOrRecycleModel $n :=$ $(n=6)$ satisfies the same transition schema plus a seed and is *not* cofinal. Growth is recovered only inside `ProtectedExactLocalMersenneRow`, whose invariants force $c >$cutoff, hence $2c-2 >$ endpoint — and maintaining those invariants is precisely what needs the sharp-capacity fill rather than this general recycling. **Closes it:** a crossing-location lower bound — for the supports actually produced, the first crossing rank satisfies $2c-2 > n$. Equivalently, keep protection alive without sharp capacity, e.g. a fill placing all new ranks above c using only the landed $c-1$-bit capacity.

Observation 9.32 (Double-or-recycle dichotomy, proved insufficient). [Lean] Lean source, Lean source, Lean source. **Yields:** $\forall n \geq 6$, an exact row at n gives an exact row at $2n-1$ or at $2c-2$ for some $4 \leq c \leq n$; plus a formal proof that this shape, even seeded at endpoint six, is logically insufficient for cofinality. **Mismatch (`quantifier_order`):** the dichotomy is proved by splitting on the sign of the witness value (`finiteErdosSum_den_odd` excludes equality); the no-go pins the deficit precisely — what is missing is not

another dichotomy but strict endpoint progress in the recycle branch, or a proof that the below-half branch recurs. **Any future attack that merely re-derives a double-or-recycle transition is already refuted by this row. Closes it:** either (a) the recycle branch upgraded to $2c-2 > n$, or (b) cofinally many below-half exact rows so the doubling branch alone is cofinal. Nothing weaker than one of these two can work through this dichotomy.

9.5 257-universal: thirteen near misses

Observation 9.33 (NM-01 / the single certificate socket). [Lean] Lean source (carry variant Lean source). **Yields:** $\forall b \geq 2$, $\forall A \subseteq \mathbb{N}$ (no infinitude, no structure hypothesis at all): if $\forall q > 0\ \exists N, K, L, C$ with $K \leq L$, $b^r \mid \text{supportCoeff}\, A\, (N+r)$ for $r \in [1, K]$, $\sum_{r=K+1}^{L} \text{supportCoeff}\, A\, (N+r) b^{L-r} \leq C$, $\exists t$, $0 < \text{supportCoeff}\, A\, (N+L+1+t)$, and $q(C+N+L+2) < b^L$ — then erdosSupportSeries $b\, A$ is irrational. Every landed #257 family (full support, multiples, periodic, eventually periodic, residue class, odd, pairwise coprime, sunflower bouquet) factors through this one socket. **Mismatch (`hypothesis_strength`):** already universally quantified over A — the universal quantifier is *not* the gap. The hypothesis `hcert` is a cofinal producer supplied only for named classes; the theorem's own docstring: "certificate existence for f_A... for arbitrary infinite A remains an open obligation." **Closes it:** a single producer theorem — $\forall A$ infinite, $\forall q > 0$, $\exists N, K, L, C$ with the same four clauses, from infinitude alone. The carry variant weakens clause 2 from digit-wise divisibility to aggregate divisibility ($b^K \mid \sum_r \text{supportCoeff}\, A\, (N+r) b^{K-r}$), so the actual minimum obligation is the weaker aggregate form.

Observation 9.34 (NM-02 / single-target to all-rational-targets, promotable). [Lean] Lean source (forward Lean source; reverse Lean source). **Yields:** $(1/2 : \mathbb{R}) \in$ mersenneAchievementSet $\Leftrightarrow$ (greedyMersenneSkippedSupport$(1/2)$).Infinite. **Mismatch (`scale_only`):** stated at the single fixed target $t = 1/2$; the obligation needs it at *every* rational target, since "some infinite A has rational value" is $\exists t \in \mathbb{Q}\ \exists A$, not $\exists A$ at t=1/2. **This is the highest-value promotable row in the bank.** The reverse direction is already stated for all reals ($0 \leq x$). The forward direction uses $1/2$ in exactly two places: (a) a generic-in-A lemma, (b) the final contradiction, which needs only that the target is *rational*. Replacing $(1/2 : \mathbb{R})$ by $((t : \mathbb{Q}) : \mathbb{R})$ and $(1 : \mathbb{Q})/2$ by t, with $0 \leq t$ in place of `norm_num`, gives $\forall t : \mathbb{Q}$, $0 \leq t \to ((t : \mathbb{R}) \in$ mersenneAchievementSet $\Leftrightarrow$ (greedyMersenneSkippedSupport$(t : \mathbb{R})$).Infinite) with no new mathematics. **Closes it:** exactly this re-parametrisation — converts the entire half-greedy refutation machine from a one-target statement into a statement about every rational target, the quantifier shape 257-universal actually needs.

Observation 9.35 (NM-03 / target-free engine underneath a target-pinned corollary). [Lean] Lean source; engine Lean source. **Yields:** $\forall A$ finite with $0 \notin A$: erdosSupportSeries $2\, A \neq 1/2$ — any support achieving exactly $1/2$ must be infinite. **Mismatch (`scale_only`):** the headline is pinned to target $1/2$ while its engine Odd((finiteErdosSum F 2).den) is already fully general over finite supports. **Closes it:** verbatim re-derivation — $\forall t : \mathbb{Q}$ with even reduced denominator, $\forall A$ finite with $0 \notin A$, erdosSupportSeries $2\, A \neq (t : \mathbb{R})$. No new mathematics. Paired with NM-02:

universal #257 at b=2 is false iff some rational t with even reduced denominator has an infinitely-skipping greedy Mersenne orbit.

Observation 9.36 (NM-04 / dyadic-only reciprocal-mass rigidity, generalization sitting unused). [Lean] Lean source (line 2116); generic-v sibling Lean source. **Yields:** for A infinite with erdosSupportSeries $2\,A = p/2^c$ (dyadic rational): $\neg$Summable$(1/a$ on $A) \vee 1 <$ reciprocalMass A. **Mismatch (`coordinate_only`):** restricted to denominators 2^c (v=1) while the obligation needs every rational $p/(2^c v)$. This restriction is a call-site choice, not a proof constraint — every supporting lemma carries a free $\{v : \mathbb{N}\}$ $(hv : 0 < v)$, and the general-v conclusion is *already landed* at line 1962 for arbitrary (v, h, L, F). **Closes it:** instantiate the landed generic theorem at $F = \{a, b\}$, $L = \operatorname{lcm} a\, b$, $h =$ multiplicative order of 2 mod v, giving $\rho(A) \geq$ doublingWrapCount$(p, v, h)/h + 1/\operatorname{lcm}(a, b)$ for every rational value. The one genuinely new step: for v>1 the excess can be 0, so strictness needs $1 \leq$ doublingWrapCount$(p, v, h)/h +$ booleanCollisionSurplus $|F|/L$ (equivalently: the mean least residue of $p \cdot 2^n$ mod v is $\geq v(1 - |F|/2L)$).

Observation 9.37 (NM-05 / T11, sparse-side rigidity only). [Lean] Lean source (line 392). **Yields:** for any A with a positive element and erdosSupportSeries $2\,A = p/(2^c v)$, v>0 (any rational, not just dyadic): $\forall \varepsilon{>}0\ \exists B \geq 0\ \forall N \geq 1\ \forall h$, SupportCoeffZeroWindow $A\,(c{+}N)\,h \to h \leq \varepsilon \log_2 N + B$. A rational-valued support cannot have super-logarithmic divisor-free windows; uniform in A, p, N, h. **Mismatch (`quantifier_order`):** the strongest genuinely universal rigidity in the bank, and it is a near-miss by *direction*, not strength. It constrains a hypothetical rational support from the sparse side only ($2 \in A$ already forces every zero window to length ≤ 1, so this is vacuous for divisor-dense A). It eliminates exactly the regime NM-09 also eliminates, nothing in the regime where the certificate producer is missing. **Closes it:** the complementary half — A infinite $\Rightarrow$ (cofinally many super-log zero windows) $\vee$ ($\forall q{>}0$, block certificate exists). This trichotomy-closure — sparse $\Rightarrow$ contradiction, dense $\Rightarrow$ certificate — is precisely what composing the rigidity corollaries would need, and the dense branch has no producer.

Observation 9.38 (NM-06 / pairwise-coprime engine excludes the primes, but the value is known). [Lean][Cited] Lean source (direct certificate producer Lean source). **Yields:** for every $b \geq 2$, an infinite pairwise-coprime support with summable reciprocal mass has irrational support series. **Method boundary:** the primes satisfy the other hypotheses but $\sum_p 1/p$ diverges, so this Lean producer does not apply. It is false, however, to infer that the base-2 prime-support value is open. Tao–Teräväinen prove

$$\sum_{p \text{ prime}} \frac{1}{2^p - 1} = \sum_{n \geq 1} \frac{\omega(n)}{2^n}$$

irrational in Theorem 1.3 of their 2025 preprint. That result is cited, not formalised here. **What remains diagnostic:** replacing the global tail budget by a windowed one would ask whether this particular certificate engine can recover the known theorem. It is no longer an open mathematical target or a claimed new route for #257.

Observation 9.39 (NM-07 / sunflower bouquet, unproved selector plus false class-existence). [Lean] Lean source (selector def line 406; carry socket line 381; bouquet structure line 33). **Yields:** for A = (finite exceptional set dividing Q) $\cup\{\text{core}_i \cdot \text{petal}_i\}$

with cores | Q, petals pairwise coprime and coprime to Q, $\sum 1/\text{petal}_i < \infty$: if `SunflowerForcedSlotTailSelection` A then irrational. **Mismatch (`multiple`):** two gaps. (i) On the support: the bouquet class strictly generalises pairwise-coprime but still demands globally pairwise-coprime, summable petals — it inherits NM-06's exclusion of the primes and adds nothing in the divisor-dense regime. (ii) On the selector: `SunflowerForcedSlotTailSelection` is an unproved cofinal supply — grep for "dichotomy" in the file returns only the namespace; despite the filename there is no dichotomy theorem there. **Closes it:** for the selector, $\forall K{>}0\ \exists N,\ 2^K \mid$ the carried first block $\wedge$ binaryCoeffTail(supportCoeff A)$(N{+}K) \le 16$, from the bouquet's persistent-reduced-modulus averaging — the file states this is what that averaging argument is supposed to give. For the class: a bouquet-existence theorem A.Infinite $\to$ Nonempty(bouquet structure) is *false* as stated (take A = all multiples of 6 — no coprime petal decomposition), so the class can only be widened, never promoted to universal.

Observation 9.40 (NM-08 / eventual periodicity, widest unconditional class, asymptotic obstruction). [Lean] Lean source (engine line 11262, producer line 11072; finite-perturbation transfer lines 9139, 9148). **Yields:** $\forall b \ge 2\ \forall m{>}0\ \forall N_0$: A infinite and m-periodic above $N_0 \Rightarrow$ irrational. Plus: the class of supports with irrational series is closed under finite symmetric difference in both directions. **Mismatch (`hypothesis_strength`):** the widest unconditional class landed (all moduli, all thresholds, no bound anywhere), but eventual periodicity is a measure-zero, countable condition; the complement is everything interesting. The two-way finite-perturbation transfer shows the good class is a union of finite-symmetric-difference equivalence classes — precisely why it cannot be enlarged by prefix surgery; the obstruction is asymptotic. **Closes it:** a "locally periodic at cofinally many scales" version — $\forall q{>}0\ \exists N\ \exists m \le f(N)$ such that $A \cap [N, N{+}L]$ is m-periodic for L large enough that the periodic sieve manufactures one block certificate. The periodic producer already only needs periodicity across the certificate window (window-local), so this is the closest genuine widening — but proving cofinally many periodic windows exist for arbitrary A is a new statement, not a re-run.

Observation 9.41 (NM-09 / lcm-gap, far-sparse edge only). [Lean] Lean source (instances Lean source, Lean source). **Yields:** $\forall b \ge 2$, $\forall a : \mathbb{N} \to \mathbb{N}$ StrictMono, $a_0 \ge 1$: if $a_k - \text{lcm}\{a_0, \dots, a_{k-1}\} \to \infty$ then $\sum_k 1/(b^{a_k} - 1)$ is irrational, via Lean source. **Mismatch (`hypothesis_strength`):** the $\mathbb{N}$-truncated-subtraction hypothesis requires a_k to outgrow the lcm of all earlier elements — extreme lacunarity. Fails immediately for A =primes (lcm of first k primes vastly exceeds p_k), squares, any positive density A. Covers only the far sparse edge — the same edge NM-05's T11 also patrols. **Closes it:** a relaxed gap measuring the surviving denominator rather than the raw lcm: $a_k - \log_b(\text{reduced denominator of the } k\text{-th partial sum}) \to \infty$. The corpus already has the denominator-survival algebra to state this (`divisor_dvd_divInt_den`, `survivingDivisor_dvd_scaled_divInt_den` at Lean source, and Lean source giving $2^{t/2}$ growth); the growth theorem for the surviving denominator of an *arbitrary* support is what is missing.

Observation 9.42 (NM-10 / the load-bearing coordinate, iff over all A). [Lean] Lean source (line 376); T7 Lean source. **Yields:** $\forall A \subseteq \mathbb{N}$ (no hypothesis at all): HasRationalValue(erdosSupportSeries 2 A) $\Leftrightarrow$ $\exists q{>}0\ \exists U : \mathbb{N} \to$

$\mathbb{Z}$, IsTemperedBinaryOrbit(supportCoeff A) q U. **Mismatch (`coordinate_only`):** already universal over A — the exact coordinate in which the obligation should be stated, and the literal shared trunk with #249 (instantiate T7 at $c :=$ Nat.totient). It does not refute the orbit: Lean source pins the orbit uniquely to $u(N) = qT_c(N)$, and $T_c(N) \le N+2$ confirms it is never ruled out on growth grounds for any c with $c(n) \le n$. **Closes it:** a support-specific non-existence theorem — $\forall A$ infinite, $\forall q>0$, $\neg\exists U$ with IsTemperedBinaryOrbit(supportCoeff A) q U. Because T7 is an iff and the orbit is unique, this is logically equivalent to the obligation — the row's value is naming the single object (the tempered carry orbit of supportCoeff A) whose non-existence is the entire content, for every A at once.

Observation 9.43 (NM-11 / carry unboundedness, tautological, dead end). [Lean] Lean source, Lean source, Lean source. **Yields:** for every infinite A and positive carry recurrence $u(n+1) + v \cdot$ supportCoeff A $(c+n+1) = 2u(n)$: u is unbounded, and $1 + v|F| \le 2u(L-c-1)$ for every finite $F \subseteq A$ with all elements dividing L. Stated for the general denominator $2^c v$, not just dyadic. **Mismatch (`multiple`):** traced to a dead end. The proof picks $2B+1$ elements of A, sets $L = (c+1)\prod F$, forcing supportCoeff A $L \ge |F|$; then $u(L-c-1) = v \cdot$ binaryCoeffTail(supportCoeff A)$(L-1) \ge v|F|/2$, so $1 + v|F| \le 2u$ holds with *no slack*. Unboundedness of the carry is a tautological restatement of "supportCoeff is unbounded on common multiples," not a rigidity. The linear tail bound cannot rescue it either: L is exponential in $|F|$ by construction, so $N+2$ is astronomically larger than $|F|$. **Closes it:** a common-multiple selection with divisor economy — an infinite A must admit L with $|\{a \in A : a \mid L\}| > 2$ binaryCoeffTail(supportCoeff A)$(L-1) + 1/v$. Since the tail is a 2^{-j}-weighted average near L, this demands an anti-concentration statement about divisor incidence of A in the window $(L, L+O(\log L)]$ — the same shape as the missing certificate, restated; the route as built cannot supply it.

Observation 9.44 (NM-12 / rank-gap dichotomy, target-generic killer under a target-pinned dichotomy). [Lean] Lean source, Lean source; general-A killer Lean source. **Yields:** an unconditional dichotomy: either $1/2 \in$ mersenneAchievementSet, or a finite witness (u, d) with prefix value $+$mersenneTail$(d+1) < 1/2 <$ prefix value $+$mersenneWeight$(d+1)$. The rank-gap killer is stated for every target and every support $(\{t : \mathbb{R}\}\{A : \text{Set}\,\mathbb{N}\}\{d : \mathbb{N}\})$. **Mismatch (`scale_only`):** the killer is already target-generic and support-generic; the *dichotomy* built on it is pinned to $1/2$ via two endpoint kills — the odd-denominator parity fact, and "finite word + full tail $\neq 1/2$" (irrationality of the Mersenne tail). **Closes it:** both inputs generalise on their face — the first to any rational t with even reduced denominator (engine already general, cf. NM-03), the second to any rational t at all (the tail is irrational by Erdős–Borwein, so finite-rational + irrational-tail $\neq$ rational). *Caveat, honestly flagged:* the two endpoint-kill proof bodies were not read in full verification; this is recorded as a strong promotability hypothesis, not a verified one — a `norm_num`-style numeric step could hide a $1/2$-specific bound.

Observation 9.45 (NM-13 / signed periodic dichotomy, sign machinery free but periodicity still required). [Lean] Lean source, Lean source (nonpos mirror line 14315). **Yields:** $\forall b \ge 2\ \forall m>0\ \forall$ $\mathbb{Z}$-valued m-periodic weight w: irrational$(\sum w(a)/(b^a - 1)) \vee \exists k, z,\ b^k x = z$; the terminating branch closes unconditionally as soon as intWeightedCoeff w is one-signed and frequently nonzero. **Mismatch (`multiple`):** two

mismatches. (i) periodicity of the weight is still required — this is the signed lift of NM-08, not an escape from it. (ii) the one-sidedness clause is *free* for #257 (a 0/1 support indicator gives supportCoeff $A \geq 0$ automatically) — the sign machinery buys nothing here; it was built for the signed #249 lane. The genuinely transferable content is the dichotomy shape itself: irrational-or-b-adically- terminating, terminating branch killable by a positive far tail. **Closes it:** an aperiodic version of the dichotomy — for arbitrary A, either erdosSupportSeries $2\,A$ is irrational or $2^k\cdot$it is an integer. For infinite A the terminating branch is already excluded by `finiteErdosSum_den_odd`-style parity plus a positive far tail, so this dichotomy — if provable without periodicity — would close the obligation outright. Periodicity is used only to manufacture the full-block certificates ("periodicity alone manufactures the full-block certificates," per the file's own docstring), so removing it is again NM-01's missing producer, restated a third time.

9.6 Three theorems assembled from the interface, and where each still stops

[Math] throughout this subsection — none of Theorems A, B, C is a Lean declaration. Each is a correct assembly of the corpus's proved pieces plus one clearly named open step; none decides #257.

Theorem 9.46 (A — reset-crossing unification)**.** *The 257-reset target exponent* $(r+5)/2$ *is not a fitted constant but a* derived *one: it follows from sqrt-scale reset anti-concentration together with the run-length-to-crossing geometry of oa-2's deficit/excess pair above, and every correction constant in the derivation is exact* (*no asymptotic* $O(\cdot)$ *notation is hiding a table*)*. Concretely: the deficit-side run law* (*oa-2, unconditional and uniform*) *converts any anti-concentration bound* $2^{\varepsilon r}$ *on* $|\Delta_r|$ *into a run-length ceiling* $L_r \lesssim (1-\varepsilon)\cdot 2r$*; matching this against the certified maximum observed run growth* (*SE-7:* $\max L_r \sim \log_2(\mathrm{row})$*, row 200,000*) *pins* $\varepsilon = 1/2$ *as the exact crossover exponent, not merely a plausible guess. Theorem A is the unification of the deficit-run and excess-run geometry into one crossing socket at that derived exponent; the open step is exactly the anti-concentration input identified in the G5/H4/F5b/oa-3b rows above.*

Theorem 9.47 (B — rigidity of dangerous resets)**.** *A dangerous reset* (*one within* $2^{(r+5)/2}$ *of the sqrt-escape threshold*) *that is preceded by a long pure R-run pins its own deviation to a single pulse-determined integer* (*RC-5 above*)*, not a free real number: the affine excess recurrence* (*rc-2*) *is exact, so once the run length and the pulse word are fixed,* w_{r_0+1} *is one of at most two integers with no remaining degree of freedom. Consequently a chain of n dangerous resets is not n independent near-misses; it is a single tower of* $n-1$ *nested exact integer coincidences, each individually as unlikely as an explicit arithmetic accident and none free to vary once the pulse stream is fixed. Failure of sqrt-escape therefore requires an exact integer-coincidence tower, not a generic accumulation of small errors — the open step is excluding the two pinned values* (*RC-5's own closing condition*)*.*

Theorem 9.48 (C — the one-sided finite decision boundary)**.** *The checked equivalence* Lean source *separates two logically different kinds of evidence. A fatal greedy gap found at a finite rank is a finite certificate that* $1/2 \notin \mathcal{A}$*. By contrast, membership requires survival at every rank; the corpus supplies no finite certificate that the orbit survives forever and no completion theorem turning a long surviving prefix into membership.*

The reset-crossing sockets, truncation-rung ladder, and sharp-capacity inequalities feed sufficient conditions for a finite fatal gap. They are not proved equivalent to one another, nor is failure to find one at a given depth evidence of survival. Consequently these searches give at most a semi-decision procedure for nonmembership of the specific value $1/2$. *Halting does not prove the universal statement in #257, and non-halting does not prove that* $1/2$ *is represented. In particular, neither truth nor falsity of the universal problem is shown semi-decidable here. Deeper negative search alone establishes nothing beyond the tested finite ranks; this is the quantifier boundary recorded by SE-7/RC-7.*

9.7 Route-collapse findings: four traps to name explicitly

These four rows are not near-misses to be closed by further work — they are targets that, once adversarially checked, turn out to be logically equivalent to the obligation itself (or to $1/2 \in$ mersenneAchievementSet directly) for the scope in which they are stated. Treating any of them as an easier waypoint is a mistake this index exists to prevent.

Remark 9.49 (hbound is equivalent to the conclusion, greedy support only). Lean source and its packaged form Lean source take a hypothesis `hbound`: $\forall N$, mobiusCenteredHalfCarry $A\,N \le 2\sqrt{N} + 4$. For the canonical greedy support $A =$ greedyMersenneSupport$(1/2)$ specifically, `hbound` is logically *equivalent* to the theorem's own conclusion erdosSupportSeries $2\,A = 1/2$ — not a strictly weaker sufficient condition. This follows from the exact residual identity at line 842, the unconditional sqrt-envelope tail bound (Lean source), the unconditional tail-nonnegativity lemma (Lean source), and the greedy-specific fact erdosSupportSeries $2\,A \le 1/2$ (line 859). **The scope limit is exact and must not be widened by inference:** this equivalence does *not* extend to every A with $1 \notin A$ — dropping the greedy-specific $\le 1/2$ fact yields only hbound $\Leftrightarrow \delta \le 0$ (a one-sided condition), not $\delta = 0$. It also does *not* extend to the sibling socket Lean source: `hbound` is stated in `Real.sqrt`, while `CofinalStripReturn`'s `halfStripBound` is built on `Nat.sqrt`, and bridging the real bound down to the tighter natural-number floor bound needed for a cofinal witness is an unaddressed, nontrivial gap — not routine packaging.

Remark 9.50 (HalfCarryCofinalTerminalOnlyStrip, widened to constant 6, still equivalent). Lean source, even after widening the strip constant from 4 to 6 (halfStripBound$'\,n$ $=$ $2\,$Nat.sqrt$\,n + 6$), is logically equivalent to $1/2 \in$ mersenneAchievementSet once three small, non-circular, mechanical lemmas are supplied (a general one-not-in-support corollary, the discrete envelope inequality Real.sqrt$\,n \le$ Nat.sqrt$\,n + 1$, and a routine prefix-to-`HalfWord` construction). None of the three gaps is equivalent to #257 or #249 individually — the equivalence sits one level up, at the socket itself. **This is a diagnostic, not a route**: it corrects TerminalOnlyCofinal.lean's own "weaker producer" framing — the terminal-only cofinal strip carries no strategic advantage over proving achievement directly.

Remark 9.51 (CofinalPositiveHalfGreedySkips is literally the problem). Repeated here for emphasis from the 257-cofinal-rows entry above: once the positivity conjunct of Lean source is recognised as unconditionally free (odd-numerator/even-subtrahend parity plus greedy nonnegativity, discharged at Lean source), the socket collapses to (greedyMersenneSkippedSupport$(1/2)$).Infinite, proved by A6 (Lean source) to be

equivalent to $1/2 \in$ mersenneAchievementSet. A proof or disproof of this one Prop settles #257-at-1/2 and vice versa; it is not a waypoint toward that settlement.

Remark 9.52 (Scope discipline for the equivalence class). All three collapsed targets above share the same failure mode against naive strengthening: each is equivalent to the achievement conclusion *only in the exact scope stated* — the canonical greedy support for `hbound`, the widened-constant strip for `HalfCarryCofinalTerminalOnlyStrip`, the single target $1/2$ for `CofinalPositiveHalfGreedySkips`. None of the three equivalences transports along the natural-looking generalizations (all A with $1 \notin A$; `Nat.sqrt`-based `CofinalStripReturn`; other rational targets) without a separately unproven bridging lemma. A future attacker who verifies the scoped equivalence and then silently drops the scope qualifier reproduces exactly the quantifier-slippage error this index is built to catch.

10 Closed routes, with the mechanism that closed them

This section is a graveyard with load-bearing headstones. Each entry below is a route that was tried against the $\surd$-escape wall (Part III) and failed, but the way it failed is itself a theorem, a computation, or a structural identity that survives the failure. A closed route is not a blank; it is a constraint on every future attempt. We record, for each closure: the route as conceived, the exact mechanism that closed it, the precise scope of the closure (what it rules out and what it leaves untouched), and whatever machinery salvages out of the wreckage. We also flag, for every closure, whether the exclusion is a statement about the *object* (the real number, the achievement set, the actual orbit) or about a *representation* of it (a coordinate, a finite-window signature, a proof-route's hypothesis set). Confusing the two is the single most common way a closed route gets silently re-opened by a later attempt using different words for the same representation.

Throughout, $\mathrm{rem}(s)$ is the seam-row remainder, $\lambda_n := \Delta_n/2^n$ the normalized seam deviation (Part I, SE-3), and $\mathcal{A}$ the Mersenne achievement set. Reset $\surd$-escape from row 10 is the statement $|\mathrm{rem}(r+1) - 2^{r+1}| > 2^{(r+5)/2}$ for every reset row $r \geq 10$ (RC-4, Part III); this is the wall every closure below is measured against.

10.1 Four closed attack families on the wall itself

Observation 10.1 (Prime-doubling — closed by explicit counterexample). **Route as conceived.** The universal recursion $K(M+1) = 2K(M) - (\tau(M+1) - 1)$ (Part I, SE-1) doubles at every step and subtracts the divisor-excess $\tau(M+1) - 1$. At a prime $M+1 = p$, $\tau(p) = 2$ so the subtracted term is 1; the naive hope is that primality of an index forces the step to be "almost pure doubling," and that prime density (Chebyshev/PNT-scale) then bounds how long a run of near-pure-doubling steps can last, which would bound the danger run-length L_r in SE-3's equivalence (III).

Exact mechanism that closed it. The hope requires primality to control the recursion *cumulatively*, but composite indices subtract more than primes add back, so a long run can mix primes and composites freely without breaking. Certified witness: the run at $n = 607$ spans $M \in (607, 617]$ with the drift quantity J staying in $[1, 5]$ for eleven consecutive steps, and three of the indices inside that window — $607, 613, 617$ — are

themselves prime. More broadly, of the 40 longest runs found for $n \leq 1500$, 26 contain a prime index somewhere inside them. Primality of an index is simply uncorrelated with the recursion staying quiet.

Scope. Closes: any argument that tries to derive a run-length bound on SE-3's L_r, or a bound on the drift J in the universal recursion, from prime-counting alone (density, Chebyshev bounds, PNT-scale gap statistics). Does not touch: arguments using the actual value of $\tau(M+1)-1$ (not just whether $M+1$ is prime), or any argument working at the divisor-sum level rather than the prime/composite dichotomy. coord:prime-density

Object or representation. About the *object*: the counterexample is a concrete witness inside the actual $K(M)$ recursion, not an artifact of how the recursion is coordinatized.

Salvage. None as a proof tool — but the witness itself (run length 11 at $n = 607$, 26/40 longest runs prime-touched) is now a banked disconfirming data point: any future prime-density-flavoured attack on run-length should be checked against this witness first. [Cert]

Observation 10.2 (Parity/perfect-squares — closed as a tautology of the recursion shape).
Route as conceived. $K(M)$ is even exactly when M is a perfect square (a true, provable fact). The hope was to use this parity law to forbid certain residues or run patterns in the digit stream of K, since parity constraints are a classical first attack on digit-run questions.

Exact mechanism that closed it. The recursion $J(M+1) = 2J(M) - \delta(M+1)$ makes $J(M+1) \equiv \delta(M+1) \pmod 2$ an *identity for any digit stream whatsoever* that this recursion shape produces — it holds regardless of what δ counts, Mersenne-specific or not. The "$K(M)$ even iff M a perfect square" fact is real, but it is a restatement of the recursion's algebraic shape, not new information: it forbids no residue and no run length. It is vacuous as an obstruction because it would hold for *every* recursion of the form $J(M+1) = 2J(M) - \delta(M+1)$, Mersenne or not.

Scope. Closes: the parity/perfect-square route completely and permanently, for this recursion shape — not a partial closure. Also closes it in advance for any structurally identical recursion (e.g. a hypothetical #249 digit-carry recursion of the same doubling-minus-forcing shape), since the vacuity is a fact about the shape, not about τ or Mersenne weights. coord:parity-tautology

Object or representation. About a *representation*: the identity is a fact about the recursion's algebraic form (base-2 doubling-with-forcing), not about the value $K(M)$ or the target constant $C = E - 3/2$ itself. It is empty precisely because it never touches the object.

Salvage. The identity remains true and is a useful sanity check (any proposed digit-run claim about K must be consistent with it), but it supplies zero constraining power on its own. [Math]

Observation 10.3 (Irrationality-measure — structurally excluded by an exponent budget).
Route as conceived. The wall reduces (SE-2) to: the greedy skip-set sum $\sum_{d \in \mathrm{Skip}_n} x_d$ never approximates $C := E - \frac{3}{2} = 0.1066951524152917\ldots$ (the shifted Erdős–Borwein constant) to within $\sim 2^{-1.5n}$. Since C has a known irrationality-measure upper bound, the natural hope is that a Diophantine-approximation exponent for C directly supplies the needed non-approximability at scale n.

Exact mechanism that closed it. The skip-set Skip_n has size $|\mathrm{Skip}_n| \approx n^2/4$ — quadratic in n, not linear. An irrationality-measure exponent μ for C bounds how well a *single* rational with denominator q can approximate C; specializing to a rational built from $\approx n^2/4$ terms only yields a lower bound on the approximation error of order $2^{-\mu n^2/4}$, against a required bound of order $2^{-3n/2}$. The shortfall is a factor of $\approx 0.42n$ *in the exponent itself* — not a constant-factor gap that a sharper μ could close, but a gap that grows without bound as $n \to \infty$ for any fixed μ. Zudilin's bound $\mu(E) \leq 2.42343562\ldots$ (Math. Notes 72 (2002) 858–862; erratum in Acta Arith. 111 (2004) 153–164 revises this to $\mu(E) \leq 2.46497868\ldots$) is the best known value, but no future improvement in μ can close this gap, because $\mu \geq 2$ always (a universal lower bound for any irrational number), while the exponent shortfall grows like n regardless of which fixed μ is plugged in.

Scope. Closes: every attempt to derive the wall from a single-number irrationality-measure statement about C (or about E), present or future, as long as the measure is a fixed exponent independent of n. Does not touch: simultaneous or multi-point Diophantine-approximation results that could in principle scale with the number of terms (no such result is known to exist for this constant, but the exponent-budget argument does not rule out the *possibility* of one — only the single-exponent route). coord:diophantine-approximation

Object or representation. About the *object*: this is a genuine value-theoretic inequality (an exponent-budget computation), not an artifact of how the skip-set is coordinatized.

Salvage. Flagged in the source corpus as the single strongest reusable negative lemma in the whole programme: any future proposal to attack a Mersenne/Erdős–Borwein-type constant via an irrationality-measure bound should be checked against this exponent-budget argument *first*, before any other work is invested. [Math]

Proposition 10.4 (2-adic fixed-precision valuation-unit no-go — Lean-verified, problem-agnostic)**.** ***Route as conceived.*** *A carry orbit's local 2-adic signature — its valuation and odd unit part at some fixed precision window — looks like a natural finite object to search for a contradiction: if every carry orbit compatible with a given local valuation-unit word were forced outside a safe centred interval, that would be a genuine local obstruction usable at every scale.*

Exact mechanism that closed it. *For* any *finite word of odd-unit-part valuation symbols at fixed precision $u > 0$, and* any *starting carry state e, there exists a compatible carry orbit realizing that exact word with every intermediate state staying centred in its dyadic interval, $|e'| \leq \mathrm{vuRadius}\, u\, \sigma$. Bounded local valuation-unit data at fixed precision can never exclude* all *finite centred carry completions — the completion always exists, constructively.*

$$\exists\, states,\ \mathrm{VUOrbit}\, u\, e\, symbols\, states \ \wedge\ (\forall\, symbol, |state| \leq radius).$$

Lean source built from Lean source and Lean source.

Scope. *Closes: any proof strategy that tries to derive a contradiction purely from "the local valuation-unit signature at fixed precision u is incompatible with X" — such signatures are* always *realizable by some carry orbit, for every u and every starting state. Does not touch: arguments using* growing *precision (an unbounded window, not a fixed one), or arguments*

that couple the local signature to extra arithmetic (e.g. a global divisibility or CRT constraint) rather than using the local signature alone. coord:2-adic-valuation-unit

Object or representation. *About a* representation*: the theorem is entirely about what a bounded local 2-adic window of a carry orbit can and cannot pin down; it says nothing about the target value C or the achievement set itself. It is a statement that a particular family of finite coordinates (fixed-precision valuation-unit words) is too coarse to see the obstruction, not that no obstruction exists.*

Salvage. *The theorem is stated and proved for* any *starting carry state and any odd unit-part word — zero arithmetic content beyond 2-adic bookkeeping. It is directly reusable as a hard stop against any future fixed-window 2-adic attack on* either *#249 or #257: a proof needs growing precision or extra coupling, never a fixed local signature alone.* [Lean]scale:n/a

Read together, Observations 10.1–10.3 and Proposition 10.4 close every attack family identified against the wall as of this writing that does not go through the digit-level carry machinery itself. What remains open (Part III, and §10.2 below) is exactly the digit-carry route — this is not a residual gap after eliminating easier options; it is the one route none of these four mechanisms can even in principle reach, because each closure above excludes a *different* kind of coarse-grained argument (density, tautological parity, single-exponent approximation, fixed-window valuation) while leaving the exact digit recursion itself untouched.

10.2 Branch exclusions at the final producer: the exceptional dyadic cells

The two-sided dyadic invariant $\min(\mathrm{rem}(s), \mathrm{overshoot}(s)) \le 2^s$ (all $s \ge 5$) is the sharpest general-purpose control on the seam orbit's scale in the whole corpus. It propagates by induction from a single local socket, SeamTwoSidedDyadicCellEscape, which at every row excludes a small number of *exceptional cells* on the middle branch and imposes one further inequality on the right branch:

$$\mathsf{SeamTwoSidedDyadicCellEscape} := \forall\, s \ge 5,\ \big(\neg\mathrm{carries} \to \mathrm{middle} \to C_s \ne -3 \wedge C_s \ne -2 \wedge C_s \ne -1\big) \wedge \big(\neg\mathrm{carries} \to \mathrm{right} \to \mathrm{overshoot} \le 2^s \to \mathrm{charge} \le 2^{s+2}\big),$$

where we write $C_s := 4 \cdot \mathrm{rem}(s) - \mathrm{belowPulse}(s) - 4$ for the integer coordinate that appears on the middle branch of the seam recursion (the same quantity that, along an all-right run, is the affine excess $\mathrm{rem}(s) - 2^s$ one step later — Part I, RC-2). Lean source

Proposition 10.5 (The upper (carries) successor needs no exceptional-cell exclusion)**.** *On the carries branch of the step (i.e.* (seamAdjacentCut *s hs*).successorCarries *holds), the two-sided invariant propagates* unconditionally*, with zero exceptional cells and no hypothesis beyond* `hcarry` *itself:*

$$\mathrm{successorCarries} \implies \mathrm{rem}(s+1) \le 2^{s+1}.$$

Lean source (*from* Lean source*, closed by* `omega`)*. In this precise sense — no cell exclusion is ever needed on this branch of the step — the upper successor contributes nothing to the exceptional-cell inventory: it is "impossible" for the upper branch to be a source of an exceptional cell in the induction step, because the branch is handled outright.* coord:seam-branch-classification
Object or representation: *about a* representation *(the induction-step case split), not the orbit itself.* [Lean]scale:n/a

Theorem 10.6 ($C_D = -3$ is impossible at the final middle producer, $D \geq 13$)**.** *Consider a hypothetical* final middle producer*: a middle transition at row* $D \geq 13$ ($\neg$carries, *and the middle-branch inequality* $4\mathrm{rem}(D) + \mathrm{gap} - \mathrm{belowPulse} < \mathrm{terminalWeight}$ *holds at* D*), followed by an all-right tail forever after* ($\forall s \geq D{+}1$, $\mathrm{seamGreedyWord}(s+1) = \mathrm{seamGreedyWord}(s).\mathrm{extend\ true}$)*. Then, unconditionally,*

$$\mathrm{belowPulse}(D) + 2 \;\leq\; 4 \cdot \mathrm{rem}(D),$$

i.e. $C_D \geq -2$ *— the cell* $C_D = -3$ *(and every more negative value) is excluded.* Lean source

Mechanism. *The proof chains three facts: (i) an all-right tail after* D *forces the terminal-augmented finite prefix at* $D{+}1$ *strictly above* $1/2$ (Lean source)*, hence the producer carry is strictly below its complete future incidence tail* (Lean source)*; (ii) that carry/tail inequality transports to the exact rational model as* $\mathrm{seamGreedyFloorZ}(D) < \mathrm{seamTakeThreshold}(D)$ (Lean source)*; (iii) combined with the pulse-absorption bound* $\mathrm{belowPulse}(D) \leq 4 \cdot \mathrm{seamWordFloorError}(D)$ *and the exact floor identity* $\mathrm{seamGreedyFloorZ}(D) = \mathrm{rem}(D) - \mathrm{seamWordFloorError}(D)$*, plus the unconditional strip bound* $1/3 < 4^D \cdot \mathrm{mersenneTail}(D) - 2^D$*, the strict inequality forces* $\mathrm{belowPulse}(D) + 1 < 4\mathrm{rem}(D)$ *in* $\mathbb{Q}$*, hence* $+2 \leq$ *in* $\mathbb{N}$*.*

Scope. *This closes* $C_D = -3$ only inside the final-middle-producer-then-all-right scenario *at* $D \geq 13$ *— not at every late middle row unconditionally. It is a strictly narrower claim than* SeamTwoSidedDyadicCellEscape*'s general induction step, which still needs all three cells excluded at* every *late middle row (not just a hypothetical final one) to propagate the two-sided invariant everywhere. The two statements share the same coordinate* C_D *but differ in quantifier scope: Theorem 10.6 is a* $\exists$*-scenario closure (one specific final producer under an all-right hypothesis), while the general socket is a* $\forall$*-row closure. Do not conflate the two: closing* $C_D = -3$ *here does not close it in the general induction step.* coord:final-producer/landing-excess ***Object or representation:*** *about the* object *under the stated hypothesis — a genuine arithmetic consequence of the all-right-tail assumption via the fatal-gap orbit, not a coordinate artifact.* [Lean]scale:n/a

Corollary 10.7 (Only $C_D \in \{-2, -1\}$ remain among the exceptional dyadic cells)**.** *Combining Proposition 10.5 (upper branch: zero exceptional cells, ever) with Theorem 10.6 (*$C_D = -3$ *closed, in the final-producer sub-case): of the three cells* $\{-3, -2, -1\}$ *that* SeamTwoSidedDyadicCellEscape *must in general exclude on the middle branch, only* $C_D \in \{-2, -1\}$ *remain genuinely open. No Lean theorem in the tree at the time of writing excludes* $C_D = -2$ *or* $C_D = -1$*, in either the final-producer scenario or the general induction step.* coord:final-producer/landing-excess [Lean]

Remark 10.8 (The tail-dominance criterion — stated exactly, and not proved). Write Θ_D for the real-valued producer-carry-to-tail ratio quantity that Theorem 10.6's mechanism step (i) actually compares against C_D: concretely, Θ_D is the scaled complete-tail budget $\mathrm{binaryCoeffTail}(\mathrm{supportCoeff}(\text{terminal-augmented prefix}))(2D + 2)$ that appears on the right-hand side of Lean source, expressed in the same integer units as C_D. The natural closing statement suggested by chaining SeamMiddleProducerCardEscape / SeamMiddleProducerRowEscape (Lean source, Lean source) with Corollary 10.7 is the **tail-dominance criterion**:

$$\textbf{(TD)} \qquad \forall\, D \geq 13,\ (\text{late middle row},\ C_D \neq -3) \implies \Theta_D < C_D.$$

If (TD) held at every late middle row, it would supply exactly the missing half of SeamTwoSidedDyadicCellEscape's remaining two cells and close the whole $\surd$-escape wall via LEAN SOURCE. **(TD) is NOT proved anywhere in the tree.** It does not appear as a theorem, a certified computation, or even a numerically-checked conjecture in the files read for this survey — it is this section's own naming of the gap the corollary leaves open, stated so that a future attempt has an exact target rather than a vague "close the remaining cells" instruction. Anyone attempting it should start from SeamMiddleProducerRowEscape (the weaker, row-scale sufficient condition $s \le \mathrm{rem}(s)$), not from the sharper card-escape socket. [Open]scale:cofinal

10.3 No-go, countermodel, and rigidity modules that bound the wall from outside

The digit-carry route is not merely unclimbed; a family of general-purpose no-go and countermodel results independently constrains *what kind of proof* can ever close it. These are not about the wall's numerics — they are about which proof architectures are structurally impossible, for any input.

Proposition 10.9 (Finite-state no-go: no bounded carry-state summary can recover history)**.** *For a "balanced pulse" family at location m (radius $r = (m+1)/2$, moving mass between positions m and $m+1$ without changing the series value), if a predecessor state is constant across the whole family, then no function* decode : State → ℕ *can recover the parameter r from* state(r) *for every r — fan-out is unbounded ($\ge \lfloor m/2 \rfloor + 2$ at $m = 2k$). More strongly, any finite* Fintype State *needs* card(State) ≥ radius + 1, *unbounded in m.* LEAN SOURCE *(fan-out lower bound at* LEAN SOURCE*).*

Scope. *Rules out, for* either *#249 or #257, any proof strategy that tries to define a bounded/autonomous "carry state" summarizing pre-m history and use it to exactly determine the post-m tail — completely general, no dependence on whether the coefficients are φ or a Möbius-support indicator.* coord:binary-digit, generic ***Object or representation:*** *about a* representation *class (any finite-automaton encoding of history) — it says nothing about C or $\mathcal{A}$ directly, only that this whole family of encodings is too weak.* [Lean]scale:n/a

Proposition 10.10 (Möbius-support countermodel: the natural negative-sign candidate overshoots)**.** *The signed Lambert identity $\sum_{d\ge1} \mu(d)/(2^d-1) = 1/2$ is exact. Writing $N := \{d : \mu(d) = -1\}$: $\sum_{d\in N} 1/(2^d-1) = 1/2 + \sum_{d\in P} 1/(2^d-1)$ where $P := \{d \ge 2 : \mu(d) = 1\}$, and quantitatively $1/2 + 1/63 \le \sum_{d\in N} 1/(2^d-1)$ (using the first positive tail term $d = 6$, $\mu(6) = 1$) — the negative-Möbius support strictly* overshoots $1/2$ *by at least $1/63$.* LEAN SOURCE *(exact decomposition at* LEAN SOURCE*).*

Scope. *Rules out exactly one natural candidate infinite Boolean support (the negative-Möbius set) as a witness for $1/2 \in \mathcal{A}$. A route-sufficiency no-go only — it says nothing about whether* some other *infinite Boolean support sums to $1/2$.* coord:mobius-mersenne ***Object or representation:*** *about the* object *— a genuine value inequality for one specific candidate set, not a coordinate artifact. Margin $1/63$ is the concrete number any repair attempt (adding/removing finitely many elements) must close or exceed.* [Lean]scale:fixed

Proposition 10.11 (No finite support ever hits $1/2$; finite certificates only ever certify death)**.** *No finite positive-index Boolean support has value exactly $1/2$: the reduced denomina-*

tor of any finite Mersenne subset-sum is provably ***odd*** *(each $2^n - 1$ is odd), while $1/2$ needs an even denominator.* Lean source. *Separately,* CertifiedGreedyMersenneDeath *is a decidable, finite-depth certificate proving a real x is* not *in $\mathcal{A}$ (e.g. $3/4$ is machine-certified dead at level 1, lookahead 0:* Lean source*), and this one-sidedness is structural:* ***absence*** *of a found death certificate, or survival through any finite depth, proves* nothing *about membership.* Lean source.

Scope. *The odd-denominator fact rules out every finite support as a $1/2$-witness, unconditionally (this is exactly B4/C3 of Part I/III's cut-locator machinery, and it is why any $1/2$-achieving support, if one exists, must be genuinely infinite). The one-sidedness caveat rules out treating a finite certified-kill search's silence as evidence of membership or rationality, for* any *member of the whole certified-kill family (this closure family, #249's TotientTailPeriodKiller, LcmConeFlatness, etc.) — a structural, not problem-specific, warning that recurs across the entire certificate-based proof programme. coord:parity-denominator / finite-search* ***Object or representation:*** *the odd-denominator fact is about the* object *(value inequality via denominator parity); the one-sidedness caveat is about the* representation *(what a finite search, as a proof method, can and cannot show). [Lean]scale:fixed*

Proposition 10.12 (Carry-orbit closure: #249's rigidity engine, occupied ground for #257)**.** *For $\sum_n \varphi(n)/2^n$: a "survivor kill" certificate (every integer candidate in a bounded box provably escapes a shrinking strip within K steps, decidable) proves the tail difference* $\mathrm{totientTail}(N+h) - \mathrm{totientTail}(N)$ *is not an integer. Rationality would force integrality along an entire period ray, so one dead multiple $m \cdot h_0$ kills the whole primitive period h_0 (*Lean source*), and this collapses onto the single $\mathbb{N}$-indexed family* $\mathrm{periodLcm}(t) = \mathrm{lcm}(1, \ldots, t)$ *(*Lean source*). Concrete unconditional deposit: every period $h \le 16$ is machine-checked dead at $(N, L) = (14, 9)$ (*Lean source*).*

Scope — a naming correction on record. *CarrySurvivorExtinction.lean, AdjacentCarryTube.lean, AdjacentPhaseSeparation.lean, and TotientCarryKernelRigidity.lean are, despite names suggestive of #257 support-rigidity,* ***#249 files*** *(namespace* `Erdos249257.TotientTailPeriodKiller`*, statements entirely about $\sum \varphi(n)/2^n$). They are recorded in this closed-routes ledger not because they close anything about #257 directly, but because (i) their engine is the shared tempered-orbit machinery also used by #257's own rigidity results (Part III's F1/T7), and (ii) a #257-support reader must know this ground is* occupied by #249*, not available — a route attempting to reuse these exact declarations for a #257 support-side argument would be reusing the engine, not the theorem. coord:integer carry-orbit, totient-specific instantiation* ***Object or representation:*** *about the* object *($\sum \varphi(n)/2^n$'s actual value), for #249 — irrelevant by content, not by mechanism, to #257. [Lean]scale:cofinal*

10.4 The scalar-localisation height obstruction

Lemma 10.13 (Denominator complement survives scaling)**.** *For $x : \mathbb{Q}$, $c : \mathbb{Z}$, $H : \mathbb{N}$: if $H \mid x.\mathrm{den}$ and $(c \cdot x).\mathrm{den} \mid H$ — i.e. multiplying by the integer c shrinks the displayed denominator down* into *H — then the* ***complementary*** *denominator factor $x.\mathrm{den}/H$ divides c:*

$$H \mid x.\mathrm{den} \;\wedge\; (c \cdot x).\mathrm{den} \mid H \;\Longrightarrow\; x.\mathrm{den}/H \mid |c|.$$

Lean source. *Equality form (height conservation): $\exists\, t : \mathbb{Z}$, $H \cdot c \cdot x = t \cdot x.\mathrm{num}$ — the omitted denominator is transferred exactly to the coefficient t, never erased.* Lean source.

Scope. *Fully generic $\mathbb{Q}$-arithmetic — no reference to either Erdős problem. It is the reusable "clearing a denominator by multiplying by a small integer c cannot discard the complementary denominator factor; it only moves that factor into c" lemma. It gives a hard floor on how small c can be if it is meant to kill a target denominator factor $x.\mathrm{den}/H$: any argument on either problem that hopes to "clear a denominator by a bounded multiplier" is bounded below by this lemma.* coord:rational-denominator/height ***Object or representation:*** *about a* representation *— pure denominator bookkeeping for $\mathbb{Q}$, with zero problem-specific content; the upstream primitive (*Lean source*) is the domain-neutral extraction this lemma builds on.* [Lean]scale:n/a

Corollary 10.14 (Mersenne specialisation)**.** *For positive $x : \mathbb{Q}$: if $2^r \mid x.\mathrm{num.natAbs}$ and $x < 2/(2^n - 1)$, then $2^r \cdot (2^n - 1) < 2 \cdot x.\mathrm{den}$ — a numerator 2-power lower bound plus a Mersenne-scale upper bound on x together force a denominator lower bound,* without *introducing a global prefix LCM.* Lean source *(generic form* Lean source*). A ready-made denominator-lower-bound tool for "synchronized tail" arguments on #257's Mersenne sums; composes directly with Lemma 10.13 for a full numerator/denominator height-transport toolkit.* coord:rational-denominator/height, Mersenne instance [Lean]scale:n/a

10.5 The final-skip band formula, and what it does not show

Definition 10.15 (Upper-reset dyadic band escape)**.**

$$\mathrm{SeamUpperResetDyadicBandEscape} := \forall\, d \geq 13,\ \mathrm{successorCarries}(d) \to \forall\, j \leq d,\quad 2^{d-j+1} < 4 \cdot \mathrm{overshoot}(d) + \mathrm{abovePulse}(d) \ \vee\ 4 \cdot \mathrm{overshoot}(d) + \mathrm{abovePulse}(d) + 2(d+j) \leq 2^{d-j+1}.$$

Lean source. In words: at every upper-reset row $d \geq 13$, the reset charge must avoid a linear-width forbidden band immediately below every dyadic threshold 2^{d-j+1}, for every $j \leq d$ simultaneously. Granted this socket, Lean source gives $1/2 \in \mathcal{A}$ outright.

Proposition 10.16 (Quantifier collapse: one critical index suffices, not $d{+}1$)**.** *The abstract, purely combinatorial statement* $\mathrm{DyadicBandEscape}(d, E) \iff \exists j,\ \mathrm{CriticalDyadicBandIndex}(d, E, j) \wedge E + 2(d + j) \leq 2^{d-j+1}$ *collapses the $\forall j \in [0, d]$ band-avoidance condition (formally $d{+}1$ separate inequalities) to checking exactly* one *nearest-boundary index j. Specialised to the concrete seam reset charge,* SeamUpperResetCriticalBandEscape *is proved logically* equivalent *to Definition 10.15's socket.* Lean source *(seam specialisation* Lean source*). Zero Mersenne/seam content in the core lemma — pure (d, E, j) arithmetic over powers of 2.* coord:dyadic-boundary, generic [Lean]scale:n/a

Observation 10.17 (Rows 13 through 31 are unconditionally certified)*.* seamUpperResetDyadicBandEscape_th verifies Definition 10.15 by kernel computation of $\mathrm{rem}(s)$ for every row $13 \leq s \leq 31$ — e.g. row 14 gives $\mathrm{rem}(14) = 392$, row 31 gives $\mathrm{rem}(31) = 4187487147$ — with **zero** rows failing the band-avoidance test. Lean source. [Cert]scale:bounded

Remark 10.18 (What this does *not* show — read the quantifiers exactly)*.* Definition 10.15 is a $\forall d \geq 13$ statement; Certificate 10.17 is a finite, kernel-checked verification for $13 \leq d \leq 31$. **The finite certificate does not establish the universal socket.** Nothing in the tree proves SeamUpperResetDyadicBandEscape for all $d \geq 13$ — this is exactly the open producer named in Part I (SE-9) and Part III (RC-4/D1's LargestSkipLateStepSocket), stated here in its own coordinate. The certified range through row 31 is real,

certified evidence *at that finite range*, and it is a genuinely different fact from "the actual orbit avoids the unsafe band cofinally" or "at every row." The socket remains an unproved sufficient condition; the finite computation is not a proof of the cofinal or universal claim, and must never be quoted as one. Concretely: the certificate says the band is avoided for 1264 real resets through row 2500 with zero classification anomalies (Part I, RC-7) and the sign law holds at all of them (RC-8) — this is strong disconfirming-of-a-counterexample evidence, not a proof that the pattern continues past the certified range. [Open]scale:cofinal

11 The mathematics this problem still needs

Everything in Parts I–IV of this document is either an unconditional theorem at a fixed or bounded scale, an exact reformulation, or a producer socket whose supply is missing. None of it decides Erdős #257, and the wall analysis explains why in a way that is more useful than "the problem is hard": the recorded failures are not a hundred separate defeats but repeated measurements of one obstruction, and the obstruction has a shape. This section takes the next step, which the rest of the corpus did not take. For each surviving route it states the exact statement that is missing, says what *kind* of mathematical object would furnish it, names the technique family that is closest and the precise reason that technique does not reach, and where possible attempts an actual construction, estimate or reduction. Attempts are banded honestly: [Math] where the argument is complete ordinary mathematics, [Cert] where the support is exact computation, [Open] where the item is a proposal and nothing more.

The organising fact, established in the wall analysis and not repeated here, is that both halves of #257 — the universal statement and the half-value instance — converge on one missing object. It has a name in this section:

Definition 11.1 (The short-window divisor phase)**.** For $M \geq 1$ and $L \geq 1$ put

$$\Theta_L(M) \;:=\; \sum_{i=1}^{L} (\tau(M+i) - 1)\, 2^{-i} \;\in\; \mathbb{Q}_{\geq 0},$$

where τ is the number-of-divisors function, so that $\tau(n) - 1 = \#\{d \geq 2 : d \mid n\}$ counts the divisors of n other than 1.

Θ_L is the analytic content of the universal recursion $K(2) = 1$, $K(M+1) = 2K(M) - (\tau(M+1) - 1)$ that carries the seam deviation (Part II, master identity $\Delta_n = K(2n) + \sum_{d \in \mathrm{Skip}_n} \lfloor 4^n/(2^d - 1) \rfloor$). It is also, in a sense made precise in §11.5, the universal-direction obstruction written in the divisor incidence coordinate. Every route below either needs a lower bound on $\mathrm{dist}(\Theta_L(M), \mathbb{Z})$ at a growing precision, or needs to explain why it does not.

11.1 R1. Reset $\sqrt{}$-escape: anti-concentration of a logarithmically short divisor sum

The statement needed

Three forms, in decreasing strength; each closes the half-value branch, i.e. yields $(1/2 : \mathbb{R}) \in \mathcal{A}$ and hence a counterexample to universal #257.

Form (i), the run-length form. Writing $\mathrm{rem}(s)$ for the integer seam greedy remainder at row s, $w_s := \mathrm{rem}(s) - 2^s$ for the deviation, and L_r for the length of the maximal pure-R run of the branch word beginning at a reset row r:

$$\forall\, r \text{ a reset row},\ r \geq 31 : \qquad L_r \ < \ \frac{r-3}{2},$$

equivalently

$$|\,\mathrm{rem}(r+1) - 2^{r+1}\,| \ > \ 2^{(r+5)/2}.$$

Any bound $L_r = o(r)$ suffices. A uniform constant bound is *false*: runs grow like $\log_2 r$ ([Cert], maximal observed run 19 at row 158,096).

Form (ii), the Lean-facing form already fanned in.

`SeamUpperResetDyadicBandEscape` $:=$ $\forall d \geq 13$, (`seamAdjacentCut` d).`successorCarries` $\Longrightarrow$ $\forall j \leq d$,

$$2^{d-j+1} < E_d \quad \vee \quad E_d + 2(d+j) \ \leq \ 2^{d-j+1}, \qquad E_d := 4 \cdot \mathtt{overshoot}_d + \mathtt{abovePulse}_d,$$

(Lean source), with the closed consumer Lean source, [Lean]. The $\forall j \leq d$ band check collapses to a single nearest-boundary index per row by Lean source, and rows $13 \leq d \leq 30$ are kernel-verified by Lean source.

Form (iii), the sign law — strictly weaker, and the right first target.

$$\forall \text{ reset rows } r : \qquad r \text{ is an } M\text{-reset} \Rightarrow w_{r+1} > 0, \qquad r \text{ is a } U\text{-reset} \Rightarrow w_{r+1} < 0.$$

[Cert]: perfect over all 1209 resets in rows 6..2500, 605/605 and 604/604, no zero deviations. [Open] as a theorem.

What kind of object would furnish it, and the exact arithmetic of Θ_L

The first thing to say is what the quantity actually is, because the corpus states it in seam coordinates and that hides its arithmetic nature. The following identity is elementary and exact, and it is the reason this route is an analytic-number-theory problem rather than a combinatorial one.

Lemma 11.2 (Divisor-residue form of the short-window phase)**.** *For all $M, L \geq 1$,*

$$\Theta_L(M) \ = \ \sum_{d \geq 2} \ \sum_{\substack{1 \leq i \leq L \\ i \equiv -M \pmod d}} 2^{-i} \ = \ \sum_{d \geq 2} 2^{-i_d(M)} \cdot \frac{1 - 2^{-d\, m_d}}{1 - 2^{-d}},$$

where $i_d(M) \in [1, d]$ is the least $i \geq 1$ with $d \mid M + i$ (so $i_d(M)$ depends only on $M \bmod d$), $m_d := \#\{i \leq L : d \mid M + i\}$, and terms with $i_d(M) > L$ are empty. In particular $\Theta_L(M)$ is a function of the residues of M modulo every $d \leq M + L$, and nothing else.

Proof. $\tau(n)-1 = \sum_{d\ge 2}[\, d \mid n \,]$; exchange the order of summation and sum the geometric progression $i_d, i_d + d, i_d + 2d, \dots$ inside $[1, L]$. □

[Math], and verified exactly as rationals at $(M, L) \in \{1000, 10001, 123456\} \times \{12, 20\}$ [Cert].

Lemma 11.2 says the object is the evaluation of a *fixed* function on the profinite integers along the +1-orbit. Define

$$\Psi : \widehat{\mathbb{Z}} \to \mathbb{R} \cup \{\infty\}, \qquad \Psi(x) \;:=\; \sum_{d\ge 2} \frac{2^{-i_d(x)}}{1-2^{-d}}, \qquad i_d(x) \in [1, d],\ i_d(x) \equiv -x \pmod d,$$

truncated at level L in the obvious way. Then $\Theta_L(M) = \Psi_L(M)$ where M is viewed in $\widehat{\mathbb{Z}}$, and the map $M \mapsto M + 1$ is the odometer, which is uniquely ergodic with respect to Haar measure. The corpus's own Reduction (A), $F(n) = \sum_{d\in D_n} 2^{d-i_d(n)}/(2^d-1)$, is *the same functional* $\sum_d 2^{-i_d}/(1-2^{-d})$ restricted to the greedy take-set D_n instead of to all $d \ge 2$. That two reductions derived years and coordinates apart are one functional on $\widehat{\mathbb{Z}}$ evaluated on two divisor sets is, to my knowledge, not recorded anywhere in the corpus; it is the cleanest available statement of what the two halves of the seam identity share. [Math] for the identification; [Open] for any consequence.

So the object required is: *an effective equidistribution statement for the odometer orbit against* Ψ*, with the precision growing linearly in the index.* Unique ergodicity gives equidistribution with no rate. What is needed is a rate, pointwise along a sparse subsequence (the reset rows), at scale $2^{-r/2}$ where r is the index. Equivalently, in the form a specialist will recognise: a lower bound for $\operatorname{dist}(\Theta_L(M), \mathbb{Z})$ with $L \asymp \log_2 M$, valid for all large M in a specified set, at precision $2^{-L/2}$.

Three technique families are the plausible shapes, and it is worth being explicit about which. (*a*) *Voronoi summation / delta method for the divisor function.* τ has an exact Voronoi expansion, and the natural attack is to bound $\sum_{M\in[X,2X]} e(a\,\Theta_L(M))$ for $1 \le |a| \le 2^L$ by opening each $\tau(M+i)$ and estimating the resulting Kloosterman-type sums. (*b*) *Fourier analysis on* $\widehat{\mathbb{Z}}$ */ Ramanujan expansions.* By Lemma 11.2 the phase is a function of M modulo the d's, so its Fourier expansion is over the characters of $\mathbb{Z}/d$, i.e. over Ramanujan sums $c_d(\cdot)$; the required estimate becomes decay of the Ramanujan coefficients of Ψ_L plus a bound on their tails. (*c*) *Second-moment / local-limit machinery.* Prove that the vector $(\tau(M+1), \dots, \tau(M+L))$, $L \asymp \log M$, obeys a *local* limit theorem strong enough to give anti-concentration of the weighted sum at scale $2^{-L/2}$.

The specific obstacle

This is the paragraph worth the most, so it is stated as sharply as the evidence permits.

The interval is logarithmically short, and every unconditional result for divisor sums in short intervals requires the interval to be a power of the modulus. The window $[M+1, M+L]$ has length $L \asymp \log_2 M$. Voronoi-based technology for $\sum_{M<n\le M+H} \tau(n)$ — including the Kloosterman-refined results of Deshouillers–Iwaniec type and every subsequent improvement — produces an asymptotic only for $H \ge M^\theta$ with $\theta > 0$ fixed, because the dual sum after Voronoi runs to length $\sim M/H$ and the error term is controlled only when that dual length is a power saving against M. At

$H = \log M$ the dual sum is of length $M/\log M$ and there is no saving at all: the method returns the trivial bound. This is not a deficiency of one paper; it is the reason short-interval divisor problems are stated with power-length windows. Technique family (a) is therefore not merely unproved here, it is out of range by an *exponential* amount in the window length, and no quantitative improvement in Kloosterman bounds changes that, because the obstruction is the length of the dual sum, not the quality of its estimate.

The one technology that does reach sub-power windows delivers the wrong quantifier and the wrong class of function. Matomäki–Radziwiłł theory shows that a multiplicative function bounded by 1 has, in almost all windows $[x, x+h]$ with $h \to \infty$ arbitrarily slowly, essentially its long average. Two mismatches, both fatal as stated. First, the conclusion is *almost all* x, whereas the seam needs *every* reset row $r \geq 31$; a set of exceptional x of density zero can still contain every reset row, because the reset rows are themselves a density-zero set defined by the orbit rather than independently of it. Second, τ is not bounded by 1 and the relevant statistic is not an average of a multiplicative function but the distance to $\mathbb{Z}$ of a *dyadically weighted* sum of τ-values; the weights 2^{-i} mean that a single value $\tau(M+1)$ contributes half the phase, so no averaging statement over the window can control it. Anything that averages within the window destroys the quantity, exactly as the wall analysis says of every scale-mismatched tool.

The joint-distribution literature is at CLT scale, not local scale. Erdős–Kac and its descendants, and the Hildebrand–Elliott-type results on the joint normality of $\omega(n), \omega(n+1)$, give distribution functions to accuracy $o(1)$ and average over $n \leq X$. The requirement here is anti-concentration to accuracy $2^{-L/2}$ for a fixed L-tuple of consecutive arguments, i.e. a local limit theorem with an error term exponentially small in L where $L \asymp \log M$. No result of that strength is known even for $L = 2$.

Why the barrier catalogue predicts exactly this. The wall analysis states that the two available information classes are fixed-precision local arithmetic (killed by Lean source, Lean source, Lean source, all [Lean]) and global Diophantine averages (killed by the exponent-budget computation against $\mu(E)$). The intermediate scale between them is precisely "divisor statistics in logarithmically short windows", and the three paragraphs above are the reason there is nothing there. The barrier catalogue and the analytic literature agree, independently, on where the gap is.

Swings

Swing 1: the exponent $(r+5)/2$ is forced, and can be derived in four lines. [Math], modulo the recorded window-width constant.

The exact R-branch recursion in deviation coordinates is $w \mapsto 4w - 4 - p$ with $0 \leq p \leq 2(s-2)$ (Part II, transition laws). Iterating L steps from a reset row r,

$$|w_{r+L}| \;\geq\; 4^L\,|w_r| \;-\; \sum_{j<L} 4^{L-1-j}(4+p_j) \;\geq\; 4^L\Big(|w_r| - \tfrac{4+2(r+L)}{3}\Big).$$

An R branch at row s requires $\mathrm{rem}(s)$ to lie in the pinned window of width $\approx 2^s + \frac{2}{3}4^{s-d}$, so in particular the run must terminate once $|w_s| \gtrsim 2^s$. Setting $s = r + L$ gives $4^L|w_r| \lesssim 2^{r+L}$, that is

$$L \;\lesssim\; r - \log_2 |w_r|.$$

This reproduces, with no fitting, the empirical law recorded at maximal slack 0.0 over rows 14..1500. A socket-violating run needs $L \geq (r-3)/2$, hence $\log_2 |w_r| \leq (r+3)/2$; the corpus's threshold $2^{(r+5)/2}$ is this bound with one bit of safety. The value of the derivation is that it shows the exponent is not a tuning parameter: any proof of Form (i) must produce a $\sqrt{}$-scale lower bound on the reset deviation and nothing weaker will do, because the doubling rate of the R branch is exactly 4 against a window that grows like 2^s.

Swing 2: measure the limit law of the phase, and compute the Borel–Cantelli budget. [Cert] for the measurement, [Open] for the conclusion.

Sieving τ over $[10^6, 1.2 \times 10^6]$ and evaluating $\Theta_{40}(M)$ for 200,000 consecutive M gives the following profile for $P(\mathrm{dist}(\Theta_{40}(M), \mathbb{Z}) < t)$:

t	2^{-6}	2^{-8}	2^{-10}	2^{-12}	2^{-14}
P	0.035415	0.009290	0.002455	0.000625	0.000185
P/t	2.267	2.378	2.514	2.560	3.031

The observed median of dist is 0.24471 against 0.25 for the uniform law, and P/t hovers near 2, the exact value for a uniform fractional part. Truncation is irrelevant at this range: the quantiles at $L = 16, 24, 32$ agree to three significant figures. So the fractional part of Θ_L behaves like Haar measure with a mildly heavy near-integer tail of constant $c \approx 2.5$. Feeding the required precision $2^{-(r-5)/2}$ into that profile,

$$\sum_{r \geq 31} c \cdot 2^{-(r-5)/2} \;\approx\; 2.5 \cdot 2^{5/2} \cdot \frac{2^{-31/2}}{1 - 2^{-1/2}} \;\approx\; 1.0 \times 10^{-3},$$

so under the Haar model the expected number of reset rows beyond the certified range that violate $\sqrt{}$-escape is about 10^{-3}, and Borel–Cantelli gives "finitely many failures" with room to spare. This is consistent with, and independently derived from, the corpus's own $\approx 2 \times 10^{-4}$ digit-model estimate. It is a heuristic. Its only rigorous content is that the required inequality is not delicate: it asks for a bound 2^{15} times weaker at $r = 31$, and exponentially weaker thereafter, than what the measured law delivers. That is a useful thing for a specialist to know before deciding whether to spend effort here.

Swing 3: attack the sign law first, and here is the reduction to try. [Open].

Form (iii) asks only for a sign, not a magnitude, and a sign is a statement about which side of the branch threshold the remainder falls on. In deviation coordinates the M branch fires iff $4\,\mathrm{rem}(s) < 2^{s+1} + 4 + p$ and then $\mathrm{rem}(s+1) = 4\,\mathrm{rem}(s) + 2^{s+1} - p$; the U branch carries. Substituting $\mathrm{rem} = w + 2^s$ turns the M-branch update into $w_{s+1} = 4w_s + 2^{s+1} - p$ and the R-branch update into $w_{s+1} = 4w_s - 4 - p$. The concrete first step is: prove that the M-branch firing condition and the reset condition together force $w_{s+1} > 0$ by a purely algebraic inequality in (w_s, p, s), with $p \leq 2(s-2)$ the only input about the pulse. If that closes, the same computation with the U-branch carry condition should give the negative sign, and the pair is a genuine lemma one rung below $\sqrt{}$-escape. If it does not close, the residual is a bound on p that is sharper than $p \leq 2(s-2)$, which is itself a short-window divisor count and hence lands back on Θ — which would be informative, because it would show the sign law is not strictly easier. Either outcome is a result; the computation is an afternoon's work and needs no new theory.

11.2 R2. Strict endpoint progress: the one place canonicalisation does not bite

The statement needed

The exact-row transition Lean source ([Lean]) says every exact row $n \geq 6$ either doubles to $2n-1$ or recycles to $2c-2$ for some $4 \leq c \leq n$. Cofinality of exact rows — which gives HALF — needs the recycle branch not to return to a bounded endpoint forever. Three sufficient forms:

(a) `SkippedCoreCriticalQuotientSupply` : $\forall D \subseteq [2,c),\ 4 \leq c,\ v(D) < \frac{1}{2} \wedge \frac{1}{2} - v(D) <$ `mersenneWeightRat` c

$$\Longrightarrow \quad 2^{(2c-2)-1} \leq \texttt{localPrefixQuotient}\,(\{c\} \cup D)\,(2c-2),$$

where $v(D) :=$ `localMersennePrefixValue` D (Lean source);

(b) the crossing rank c produced by the recycle branch does not repeat a bounded value infinitely often;

(c) a growth law $c \geq g(n)$ with g unbounded, for the recycle branch at incoming endpoint n.

What kind of object would furnish it

This route is different in kind from R1, and the difference is worth stating because it is the only place in the half-value branch where the witness space is not a singleton.

Canonicalisation — the fact that every straddling word equals the canonical greedy prefix (Lean source) — applies to conditions phrased as *real-valued straddles*, because it consumes both $v(D) < 1/2$ and the deficit clause $1/2 - v(D) < w_c$. The predicate `ExactLocalMersenneHalfRow` n imposes neither: it is an *integer floor-quotient hit*, `localPrefixQuotient` $D\ n = 2^{n-1} - 1$, and floor division is not injective. So the per-endpoint witness space is an exponential search over $D \subseteq [2,n]$ and is decidable for each n. That is the only genuine degree of freedom left anywhere on this side.

The object required is therefore a *monotone quantity* — a potential function on exact rows that the recycle branch is forced to increase. In the language of the transition, one wants a function Φ on pairs (endpoint, witness core) with Φ(recycle image) $> \Phi(n)$, whose unboundedness forces $c \to \infty$. The technique family is not analytic here: it is the theory of *well-founded transition systems* and, arithmetically, the capacity accounting of Boolean subset sums against a superincreasing weight sequence. The corpus already has the sharp capacity test in exact form: by Lean source, form (a) is exactly the statement that the $c-2$-bit suffix capacity `localBinarySuffix` $D\,1\,(2c-2) < 2^{c-2}$ is met.

The specific obstacle

The obstruction here is isolated, formalised, and the sharpest of its kind in the corpus: the transition schema *plus* a seed provably does not give cofinality. The theorem Lean source ([Lean]) exhibits the one-point predicate $P(n) := (n = 6)$, which satisfies the seed (Lean source, the explicit witness $\{2, 3, 6\}$) and the exact shape of the transition, and is not cofinal. So no argument that uses only the dichotomy and the base case can work, however cleverly the induction is arranged. Any proof must consume arithmetic content of `ExactLocalMersenneHalfRow` that the abstract schema does not see.

There is a second, more specific trap, and it is also formalised. The natural sufficient condition via fractional split mass is *not necessary*: the fixture Lean source ([Lean], at $D = \{2, 3\}$, $c = 5$) exhibits a real crossing core violating it. So the fractional mass route is closed with a witness, not merely unattempted.

Swings

Swing 4: use canonicalisation as a help rather than an obstruction. [Math] for the observation, [Open] for the outcome.

Form (a) is universally quantified over cores D, which looks like an exponential obligation. It is not. Its hypotheses are exactly those of Lean source ([Lean]), which forces $D = \texttt{halfGreedyPrefixSupport}(c - 1)$. So the $\forall D$ collapses to a single orbit and form (a) is a statement to be checked along *one* sequence indexed by c, namely

$$\forall c \geq 4 : \qquad \texttt{localBinarySuffix}\ (\texttt{halfGreedyPrefixSupport}(c-1))\ 1\ (2c-2)\ <\ 2^{c-2}.$$

That is a one-parameter arithmetic inequality about the canonical half-greedy prefix, decidable for each c, and it is the correct object to compute before any theory is attempted. The first actual step is to evaluate the left-hand side for $c = 4, \ldots, 200$ and read off whether the margin $2^{c-2} - \texttt{localBinarySuffix}(\cdots)$ grows, stays proportional, or approaches zero. Those three outcomes select three different proofs, and no one has published the table.

Swing 5: the potential function to try. [Open].

The doubling branch sends $n \mapsto 2n - 1$ and the recycle branch sends $n \mapsto 2c - 2$ with $c \leq n$. Set $\Phi(n) :=$ the largest rank in the canonical core at endpoint n. On the doubling branch Φ at least doubles. The recycle branch is the only danger, and its image endpoint $2c - 2$ is determined by the *first crossing rank* c of the core. So the entire question is whether the first crossing rank of the canonical prefix can stay bounded while the endpoint grows — which, by the collapse above, is a question about one explicit greedy orbit. The concrete statement to attempt is: *the first crossing rank of* $\texttt{halfGreedyPrefixSupport}(m)$ *is non-decreasing in* m *and unbounded*. If true it closes form (c) directly; if false, the counterexample is a finite object and worth having.

11.3 R3. Prime support: a known theorem and an internal method boundary

The base-2 statement

$$\text{Irrational}\Big(\sum_{p \text{ prime}} \frac{1}{2^p - 1}\Big)$$

is known. Expanding geometrically and collecting divisors gives

$$\sum_{p \text{ prime}} \frac{1}{2^p - 1} = \sum_{n \geq 1} \frac{\omega(n)}{2^n},$$

and Tao–Teräväinen prove the latter series irrational in Theorem 1.3 of *Quantitative correlations and some problems on prime factors of consecutive integers*. This is a cited analytic result, not a theorem formalised in the pinned Lean corpus.

The existing pairwise-coprime Lean theorem does not recover it because that theorem assumes $\sum_{a \in A} 1/a < \infty$, whereas $\sum_p 1/p$ diverges. Its generic block-certificate engine would also demand, in the prime case,

$$2^r \mid \omega(N + r) \qquad (1 \leq r \leq K),$$

a simultaneous exact divisibility pattern far stronger than the input used by Tao–Teräväinen. Small computational witnesses for $K = 2, 3$ diagnose that the condition is not vacuous, but they neither prove the known theorem nor define a new open case.

Thus the contribution of this route audit is negative and coordinate-specific: it identifies why the present formal certificate engine fails to reproduce a theorem already available by quantitative correlation estimates. The earlier proposal to prove the prime-support irrationality as new mathematics is withdrawn. Extensions to other bases or prime-power supports are mentioned by the cited authors with proof modifications omitted; this paper promotes only the fully written base-2 theorem.

11.4 R4. Landing the route-collapse equivalences, via perfect-square witness depths

The statement needed

Five biconditionals, each with $(1/2 : \mathbb{R}) \in \mathcal{A}$ on the right:

(1) $\forall N$, `mobiusCenteredHalfCarry` $G\,N \leq 2\sqrt{N} + 4$; (2) `CofinalPositiveHalfGreedySkips`;

(3) `CofinalExactLocalMersenneHalfRows`; (4) `GreedyHalfCarryCofinalStripReturn`; (5) `HalfCarryCofinalTerminalOnlyStrip`,

where $G :=$ `greedyMersenneSupport`$(1/2)$. The forward directions of (2)–(5) are landed Lean (Lean source, Lean source, Lean source, Lean source, all [Lean]). What is missing is the reverse directions, and they all run through one lemma.

The object, and the swing that supplies it

Lemma 11.3 (Perfect-square strip witness)**.** *Let* $A \subseteq \mathbb{N}$ *with* $1 \notin A$ *and* `erdosSupportSeries` $2\,A = 1/2$. *Then for every* $k \geq 1$,

(`integerHalfCarry` $A\,(k^2 - 1) : \mathbb{R}$) $=$ `binaryCoeffTail` (`supportCoeff` A) $(k^2) \leq 2k + 4 =$ `halfStripBound` (k^2).

Proof. By Lean source ([Lean]),

$$(\texttt{integerHalfCarry}\ A\ N : \mathbb{R}) = 2^{N+1}\left(\tfrac{1}{2} - \texttt{erdosSupportSeries}\ 2\ A\right) + \texttt{binaryCoeffTail}\ (\texttt{supportCoeff}\ A)\ (N+1),$$

and the first term vanishes by hypothesis. Take $N = k^2 - 1$, so $N + 1 = k^2$. By Lean source ([Lean], unconditional) the tail at k^2 is at most $2\sqrt{k^2} + 4 = 2k + 4$, and by Lean source, `halfStripBound` n = `2 Nat.sqrt` $n + 4$, which at $n = k^2$ equals $2k + 4$ exactly. Non-negativity of the tail (Lean source, [Lean]) gives the lower side. □

[Math]; every input is landed Lean and the composition is four lines.

The point is narrow and entirely technical, which is why it is worth stating: the corpus's recorded obstruction to closing (4) and (5) was the mismatch between `Nat.sqrt` (in `halfStripBound`) and `Real.sqrt` (in the unconditional tail bound), for which Mathlib supplies only `Real.nat_sqrt_le_real_sqrt`, the wrong direction, and the recorded workaround was to widen the strip constant from 4 to 6. Restricting the cofinal witness depths to perfect squares dissolves the mismatch: at $n = k^2$ the two roots coincide exactly and no widening is needed. Applying Lemma 11.3 along $M = k^2 - 1$ gives (4) directly; taking a to be the indicator of an achieving support truncated to `Fin`$(M + 1)$ at $M = k^2$, and using Lean source together with $1 \notin A$ (forced because `mersenneWeight` $1 = 1 > 1/2$), gives the witness required by Lean source and hence (5).

Why this is a result and not bookkeeping

Under the standing rule that a reduction is not a result, one might discount R4. That would be a mistake, and the reason is the organising thesis of this document. Each biconditional is a theorem, not a reduction: it says a named producer socket is *equivalent* to the goal, hence permanently not easier, hence removable from the list of things worth attempting directly. Five such theorems are a systematic elimination, and a systematic elimination is a statement about the problem. In particular (3) settles the advertised weakening — that no compatibility is required between witness supports at different endpoints — as buying nothing: incompatible witnesses do not create flexibility, because the resulting Π^0_2 statement still collapses to the Π^0_1 goal.

Scope limits to be stated verbatim in the docstrings. (1) does *not* extend to every A with $1 \notin A$: without Lean source (which supplies $\delta := 1/2 - S_G \geq 0$ for the greedy support specifically) one obtains only the one-sided equivalence hbound $\iff$ $(1/2 - S_A) \leq 0$. The mechanism in all five cases is the same and should be recorded once: $2^{N+1}\delta$ dominates any $\sqrt{N}$ envelope, so a strip hypothesis forces $\delta = 0$; and $\delta = 0$ makes the unconditional tail bound supply the strip. Membership is then read off by Lean source ([Lean]).

11.5 R5. The sparse/dense trichotomy — and a proof that the present engine cannot close it

The statement needed

$$\forall A \subseteq \mathbb{N}_{\geq 1}: \quad A.\text{Infinite} \implies \underbrace{(\text{cofinally many super-logarithmic zero windows in } \texttt{supportCoeff}\ A)}_{\text{sparse branch}} \vee \underbrace{\text{Cert}(A)}_{\text{dense branch}}.$$

The sparse branch is already closed twice over: it contradicts Lean source ([Lean], unconditional and uniform in A) whenever the series is rational, and the far sparse edge is covered outright by Lean source ([Lean]). The dense branch is the missing producer: for A whose divisor incidence is bounded below infinitely often, construct

$$\exists N, K, L, C, \quad K \le L, \quad \forall r \in [1, K] : \ 2^r \mid \mathtt{supportCoeff}\, A\, (N + r),$$

$$\sum_{r=K+1}^{L} \mathtt{supportCoeff}\, A\, (N + r)\, 2^{L-r} \le C, \qquad q\,(C + N + L + 2) < 2^L.$$

What kind of object would furnish it

$\mathtt{supportCoeff}\, A = \mathbf{1}_A * \mathbf{1}$ is a Dirichlet convolution, so the dense branch asks: for an *arbitrary* 0–1 arithmetic function, find a block of consecutive arguments on which the convolution with $\mathbf{1}$ is 2-adically aligned. Möbius inversion $\mathbf{1}_A = \mu * \mathtt{supportCoeff}\, A$ is exact and costs nothing, which makes the 2-adic structure of the incidence directly readable, and that is the technique family to use: 2-adic Möbius/Dirichlet algebra on the divisor lattice, not analysis. The reason the corpus's landed theorems all live inside the two-axis box (divisor independence *and* a global tail budget) is that CRT prescribes an incidence pattern only when the divisibility conditions $a \mid n$ for $a \in A$ are independent; for general A they are entangled through pairwise gcds and CRT has nothing to prescribe.

The specific obstacle, proved

Here the honest answer is stronger than an obstacle: the dense branch as stated is *false*, and the falsifier is an explicit, natural, divisor-dense support. This also corrects the reduction recorded earlier in this document, that universal #257 passes "with no loss" to $\mathrm{O1}' := \forall A$ infinite, $\mathrm{Cert}(A)$.

Proposition 11.4 (Squarefree support: exact engine ceiling, not an open value)**.** *Let $A = \{n \ge 2 : n \text{ is squarefree}\}$. Then*

$$\mathtt{supportCoeff}\, A\, (n) = 2^{\omega(n)} - 1,$$

which is odd for every $n \ge 2$. Consequently neither the digitwise nor the carry-aware divisibility-first block-certificate schema has an instance at any even base. The two exact no-go theorems are Lean source and Lean source.

[Lean] scale:uniform coord:binary-digit.

This does not leave the squarefree value open. Duverney–Tachiya's Corollary 1.2 and Example 1.1 prove that, for every $h \ge 1$,

$$1, \quad \sum_{n \text{ squarefree}} \frac{1}{2^n - 1}, \quad \dots, \quad \sum_{n \text{ squarefree}} \frac{1}{2^{hn} - 1}$$

are linearly independent over $\mathbb{Q}$. After deleting the rational $n = 1$ term, squarefree support is therefore irrational at every base 2^j. Among the integer bases, the cited theorem leaves the non-powers of two uncovered. This is cited mathematics, not formalised here. Proposition 11.4 is a failure of two proof schemas on a value already known to be irrational.

The coordinate dependence is exact. Adjoin 1 to the support. The series changes by the rational number $(b-1)^{-1}$, so its irrationality is unchanged, while the incidence becomes $2^{\omega(n)}$: Lean source and Lean source. At base 2, CRT then supplies arbitrarily long opening blocks with $\omega(N+r) \geq r$: Lean source and Lean source. The middle-window and height conditions remain; no certificate or new irrationality proof follows. The durable conclusion is methodological: test a no-go statement under rational finite changes before attributing it to the underlying value.

11.6 What the shape of the wall suggests about the underlying object

Everything in this subsection is inference from the accumulated failures, not theorem. It is labelled that way throughout, and where the evidence does not determine the answer it says so rather than manufacturing a story.

The half-value branch looks pseudorandom, and the failures are the evidence

The strongest available inference is about the *kind* of difficulty, and it comes from the manner in which the attacks failed rather than the fact that they failed.

When a hard problem conceals structure, partial attacks typically return partial information: a bias, a correlation, an exceptional set, a bound with the wrong constant. That is not what happened here. Parity returned a tautology — the identity $J(M+1) \equiv \delta(M+1) \pmod 2$ holds for *any* digit stream obeying the recursion shape, so it forbids no residue and no run. Fixed-precision valuation data returned a proved no-go: for any starting carry and any word of odd-unit symbols, a compatible centred completion exists (Lean source, [Lean]). Bounded carry state returned an unbounded fan-out (Lean source, [Lean]). Long common suffixes returned exact erasure of the predecessor (Lean source, [Lean]). Prime-doubling returned a witness where three primes sit inside an eleven-step quiet run. These are not near misses. They are the signature of a coordinate in which there is no arithmetic bias to find — which is what one expects if the digit stream really is pseudorandom.

The positive measurements point the same way and are quantitative. The branch counts over 200,000 seam rows are $R : M : U = 100197 : 49899 : 49898$, matching the balanced-signed-digit prediction to within sampling noise. The run-length histogram is geometric with ratio ≈ 2. The binary expansion of $C = E - 3/2$ agrees bit-for-bit with OEIS A211706 and has longest equal run 15 in 40,000 bits. And the measurement in Swing 2 above shows the fractional part of the short-window divisor phase Θ_L is, to three significant figures, Haar-distributed, with a near-integer tail constant $c \approx 2.5$ against the uniform value 2.

Inference (*not theorem*). The object behind the half-value branch is most likely **simple in description and pseudorandom in behaviour**. Its description is two lines: $K(2) = 1$, $K(M+1) = 2K(M) - (\tau(M+1) - 1)$, together with a superincreasing greedy orbit. Its

behaviour is, on every statistic yet measured, indistinguishable from a fair coin. If that is right, then the missing theorem is a *pseudorandomness* statement about an explicit deterministic sequence, and it belongs to the same family as normality of specific constants, or Chowla-type independence: elegant to state, and hard for the structural reason that there is no structure to exploit. On this reading the wall is not hiding a mechanism; it is the absence of one.

Three reasons to distrust that inference, stated plainly

First, coordinate-relativity. Obstructions are relative to a representation, and all nine coordinates the corpus built — rung, seam row, integer margin, sharp tail margin, half-carry, reverse-carry word, Boolean–Möbius exact row, fatal-cell packing, Dedekind cut-locator — are descendants of one family, the superincreasing greedy carry. Their agreement is therefore *not* nine independent votes. A genuinely foreign coordinate (a p-adic dynamical one, a representation-theoretic one, an automatic-sequence one) has not been tried, and "no bias in these nine" is weaker evidence than it looks.

Second, the single exception. Exactly one row of the real orbit fails $\sqrt{}$-escape, $r = 7$, which is also the site of the only real R-branch crossing ever observed, at $(s, d) = (10, 7)$. A lone low-lying exception is what a random model predicts; it is also what a hidden structure with a small modulus predicts. That datum does not discriminate.

Third, and this is the sharpest internal check available: pseudorandomness is a claim about the *absence* of structure, and the corpus contains an unmistakable piece of *present* structure — strict superincreasingness of the Mersenne weights, which forces the digit coding to be injective (Lean source, [Lean]), collapses every witness space to a singleton, and makes the static adversarial tier coincide with the dynamic greedy orbit. That structure is rigid, not random. So the honest description is a hybrid: a rigid skeleton with a pseudorandom digit stream running along it. The rigidity is why every soft method fails (there is no ensemble to average over: $\mathcal{A}$ has Lebesgue measure exactly 1 inside an interval of length $E \approx 1.6067$ and is nowhere dense, and the two standard notions of "typical" disagree on it); the pseudorandomness is why every arithmetic method returns no information.

The universal branch reads differently, and this weakens the one-wall thesis

Proposition 11.4 is a reason to separate the two halves rather than assimilate them. On the universal side the present certificate engine first meets a coordinate-dependent parity obstruction. The squarefree value is nevertheless already known to be irrational at power-of-two bases, and adjoining 1 removes the opening-block obstruction without changing irrationality. Thus the no-go diagnoses this engine; it does not diagnose the value or the universal problem.

The wall analysis's claim that both halves converge on one missing object — short-window near-integer anti-concentration of a divisor-weighted sum — survives, but with a correction: on the universal side that object is the *second* obstruction, not the first, and it becomes relevant only after one chooses an engine whose opening condition is invariant

under rational finite changes. Reporting the convergence without that ordering would overstate the unity of the wall.

What the evidence does and does not determine

It does not determine whether a short proof exists. Nothing in a catalogue of failed methods bounds the length of a successful one, and this document should not pretend otherwise.

It does determine, with proof rather than inference, where a short proof cannot be. It cannot proceed by finite inspection, by fixed-precision arithmetic, by a bounded carry state, by a measure or category argument, by an irrationality-measure bound on E or any constant of that type, or — on the universal side, as of Proposition 11.4 — by instantiating the present block-certificate engine at a general support. Each of those is closed with a mechanism, and a mechanism is reusable in a way that a failed attempt is not. That is the contribution: not a proof, and not a reduction, but a map of the region in which a proof cannot live, drawn tightly enough that the remaining region is small and its coordinates are named.

12 What is open, stated exactly

Erdős #257 asks whether $\sum_{n\in A} 1/(2^n-1)$ is irrational for *every* infinite $A \subseteq \mathbb{N}_{\geq 1}$. Nothing in this paper decides this, in either the universal form or the single-target half-value form below. Every corpus result surveyed in Parts I–IV is either a fixed/bounded-scale instance, a coordinate reduction, or a producer socket still lacking its supply. This section states the exact open targets, in display mathematics, with their Lean sites where formalised and their precise logical relationships to one another. **Read the implication table in §12.8 before attacking any single target below**: several apparent waypoints are *equivalent* to the full problem, not easier approaches to it.

12.1 Notation fixed for this section

The functional `erdosSupportSeries 2 (−)` denotes the base-2 Erdős support series `erdosSupportSeries 2` $A := \sum_{n\in A} 1/(2^n-1)$ (Lean source). `supportCoeff` $A\ n :=$ $|\{d \mid n : d \in A\}|$ (Lean source) is the divisor-indicator coefficient whose generating tail feeds every certificate below. $\mathcal{A} :=$ `mersenneAchievementSet` $= \{x \mid \exists A,\ 0 \notin A \wedge x =$ `positiveMersenneSupportValue` $A\}$ (Lean source). $G :=$ `greedyMersenneSupport`$(1/2)$ is the canonical greedy support for target $1/2$; `seamIntegerGreedyRemainder` s (written $\mathrm{rem}(s)$ below) is its integer seam remainder at row s, and a *reset* at row r is a row where the greedy orbit returns to the near-2^r regime after a run of takes or skips.

12.2 Target O1 — Universal Erdős #257

Definition 12.1 (Universal #257)**.**

$$\mathrm{U257} := \forall A \subseteq \mathbb{N}_{\geq 1},\ A.\mathrm{Infinite} \implies \mathrm{Irrational}(\texttt{erdosSupportSeries}\ 2\ A)\,.$$

This is the literal statement of Erdős #257. [Open]. The corpus's own sufficient-condition engine is

$$\mathrm{Cert}(A) \;:=\; \forall q > 0,\ \exists N, K, L, C,\ K \le L,\ (\forall r \in [1, K],\ 2^r \mid \texttt{supportCoeff}\, A\, (N + r)) \wedge$$

$$\sum_{r=K+1}^{L} \texttt{supportCoeff}\, A\, (N + r) \cdot 2^{L-r} \le C \ \wedge\ \exists t,\ 0 < \texttt{supportCoeff}\, A\, (N + L + 1 + t) \ \wedge\ q(C + N + L + 2) < 2^L,$$

with $\mathrm{Cert}(A) \Rightarrow \mathrm{Irrational}(\texttt{erdosSupportSeries}\ 2\ A)$ already proved unconditionally for *every* A (Lean source, [Lean], scale:n/a — the theorem is universally quantified over A; only its hypothesis is unsupplied). Therefore the stronger statement

$$\mathrm{O1}' \;:=\; \forall A,\ A.\mathrm{Infinite} \Longrightarrow \mathrm{Cert}(A).$$

would imply U257, but the reduction is *not* lossless and O1′ is actually false. For $A = \{n \ge 2 : n \text{ squarefree}\}$ the incidence is $2^{\omega(n)} - 1$, hence odd for every $n \ge 2$, and the digitwise certificate schema has no instance at any even base (Lean source, [Lean]). This is a method ceiling, not a counterexample to irrationality: Duverney–Tachiya prove the squarefree value irrational at every power-of-two base. Likewise the pairwise-coprime producer does not cover the primes because $\sum_p 1/p$ diverges, even though Tao–Teräväinen prove the base-2 prime-support value irrational. The honest universal target is U257 itself; O1′ is a rejected sufficient strengthening, not an open equivalent normal form.

12.3 Target O2 — Half-value membership

Definition 12.2 (Half membership)**.**

$$\mathrm{HALF} \;:=\; \left(\tfrac{1}{2} : \mathbb{R}\right) \in \mathcal{A}.$$

HALF is a *single instance* of the negation of U257: any witness A with $0 \notin A$, $\texttt{erdosSupportSeries}\ 2\ A = 1/2$ is infinite (Lean source, [Lean], scale:uniform — finite supports have odd-denominator value, so can never equal a dyadic rational), hence a genuine counterexample to U257. So HALF $\Rightarrow \neg$U257, strictly — it does not decide U257 if false, since U257 could still fail at a different target $t \in \mathbb{Q}$. HALF is [Open]. It is logically equivalent, by an unconditional Lean iff, to a purely combinatorial statement about one fixed orbit:

$$\mathrm{HALF} \iff (\texttt{greedyMersenneSkippedSupport}(1/2)).\mathrm{Infinite}$$

(Lean source, [Lean], scale:uniform, coordinate coord:mobius-mersenne, target-generic in the reverse direction — the forward direction uses only that the target is rational, so the same theorem holds verbatim for every rational $t \ge 0$ in place of $1/2$).

12.3.1 Waypoints that are actually EQUIVALENT to O2, not weaker

Three named "producer sockets" that read as strictly easier routes to HALF have been adversarially checked and found to be *equivalent* to HALF itself. Attacking any one of them is attacking Erdős #257-at-1/2 directly.

Proposition 12.3 (CofinalPositiveHalfGreedySkips is equivalent to HALF)**.** *Define*

$$\text{CPGS} := \forall N,\ \exists c \geq N,\ c \text{ skipped by the rational half-greedy orbit} \ \wedge\ 0 < \texttt{greedyMersenneRemainderRat}(1/2)(c-1)$$

(Lean source)*. The positivity conjunct is unconditionally true for every skipped* c (Lean source, *[Lean], scale:uniform — an odd-denominator parity fact*)*, hence deletable without changing the Prop's meaning. What survives,* $\forall N\, \exists c \geq N$ *skipped, is exactly* $(\texttt{greedyMersenneSkippedSupport}(1/2)).\text{Infinite}$*, so*

$$\text{CPGS} \iff \text{HALF}.$$

[Math, adversarially verified] (Survivor 7, composition-survivor pass). scale:cofinal. coord:mobius-mersenne, greedy orbit. This closes off what the source bank advertised as "the shortest known path" to a disproof (Lean source): there is no shortcut through it, because it is not shorter.

Proposition 12.4 (The terminal-only cofinal strip is equivalent to HALF)**.** *Let* `HalfCarryCofinalTerminalOnlyStrip` *be the socket that a finite normalised word* a *at cofinally many depths* M *satisfies* $|\texttt{integerHalfCarry}(\texttt{wordSupport}\, a)(M-1)| \leq \texttt{halfStripBound}\, M$ *where* $\texttt{halfStripBound}\, n = 2\sqrt{n} + 4$ (Lean source)*. Widened to strip constant* 6 *(i.e.* $2 \cdot \texttt{Nat.sqrt}\, n + 6$*), this socket is logically equivalent to* HALF*: every truncation of an achieving support is already a witness (not merely cofinally many), given three small mechanical lemmas not yet in the corpus (a general* `one_notMem_of_erdosSupportSeries_eq_half`*; the discrete envelope* $\texttt{Real.sqrt}\, n \leq \texttt{Nat.sqrt}\, n + 1$*; and an explicit* `HalfWord` *construction from* $A \cap \text{Iic}\, M$*).*

[Math, adversarially verified, weakened] (Survivor 5). scale:cofinal. coord:support rigidity, half-carry. This corrects the module's own "weaker producer" framing — the strip socket carries no strategic advantage over proving HALF outright.

Observation 12.5. Neither Prop. 12.3 nor Prop. 12.4 narrows the search space for HALF. Their value is negative: they retire two named targets from the "worth chasing as an easier waystation" list. A third socket, the hypothesis $\forall N$, $\texttt{mobiusCenteredHalfCarry}\, A\, N \leq 2\sqrt{N} + 4$ of Lean source, is likewise proved EQUIVALENT to HALF for the specific support $A = G$ (Survivor 4, [Math, adversarially verified, weakened]) — not, as originally claimed, for every A with $1 \notin A$; the general-A version is only a one-sided implication (hbound $\Rightarrow$ $\texttt{erdosSupportSeries}\, 2\, A \leq 1/2$).

12.4 Target O3 — Cofinal exact local Mersenne half rows

Definition 12.6.

$$\text{CofinalExactLocalMersenneHalfRows} := \forall N,\ \exists n \geq N,\ \texttt{ExactLocalMersenneHalfRow}\, n$$

(Lean source, [Open], scale:cofinal). No compatibility is required between witness supports at different endpoints — the defining weakening versus a single coherent trajectory.

CofinalExactLocalMersenneHalfRows ⇒ HALF is proved unconditionally via compactness of $\mathcal{A}$ (continuous image of $(\mathbb{N} \to \mathrm{Fin}\,2)$ under the digit-coding map, hence closed) (Lean source, [Lean]). This direction is *strictly weaker* than HALF as currently proved: the argument uses closedness alone and discards the producing rows, so the theorem as it stands does not pin a single limiting support A; a genuine sequential-compactness extraction (Cantor-space diagonal argument over $(\mathrm{Fin}\,2)^{\mathbb{N}}$, using the file's own Lean source) would upgrade this to producing A explicitly, and would additionally transfer $2 \in A$ for free (Lean source) — but this upgrade is itself unformalised (Survivor 1, [Math, adversarially verified, weakened]). So O3 sits strictly below HALF in the current Lean state of the corpus, not merely logically.

12.4.1 Sharp critical-capacity conditions — the BooleanMöbius lane's producer for O3

The corpus's engine for O3 is a sharp c-2-bit capacity test. Write

$$\texttt{localBinarySuffix}\ D\ k\ M\ :=\ 2^{M-k} - \texttt{localPrefixQuotient}\ D\ M - 1$$

(Lean source). For $4 \le c$, every $d \in D$ with $2 \le d < c$, and $\texttt{localMersennePrefixValue}\ D < 1/2$:

$$\text{(loose, proved)}\quad \texttt{localBinarySuffix}\ D\ 1\ (2c-2)\ <\ 2^{c-1} \qquad (*)$$
$$\text{(sharp, needed)}\quad \texttt{localBinarySuffix}\ D\ 1\ (2c-2)\ <\ 2^{c-2} \qquad (**)$$

$(*)$ is proved unconditionally, uniform in c, no case split (Lean source, [Lean], scale:uniform). $(**)$ is exactly one bit sharper and, via Lean source (an unconditional iff), is *equivalent* to an exact row already crossing the integer half-target at c. The true unproved gap between $(*)$ and $(**)$ is additive, not multiplicative: the proof of $(*)$ derives $\texttt{localBinarySuffix} < 2^{c-2} + |D|$ and discards $|D| \le c-2$; so what is actually open is the linear-width anti-concentration

$$\text{(O3-supply)}\qquad \forall N\ \exists c \ge N,\ \texttt{localBinarySuffix}\ D\ 1\ (2c-2)\ \notin\ [\,2^{c-2},\ 2^{c-2}+c-3\,]$$

— a band of width $c-2$ inside a range of size 2^{c-2}, exponentially weaker than the reset $\sqrt{\text{-escape}}$ target of §12.5 (band width c vs. width $2^{(c+5)/2}$). The generic (non-greedy) sharp-fill theorem itself is already unconditional and strictly more general than every consumer built on it: Lean source ([Lean]) needs only $(**)$ for *some* below-half core D, not a greedy prefix — this is the escape route named in Prop. 12.3's remark: abandoning the canonical greedy orbit (which forces CPGS-equivalence) for an arbitrary below-half core is a genuinely different, still-open target.

The corpus's complete inductive skeleton from the endpoint-six seed to O3 is already assembled (Lean source, [Lean]); it consumes the sharp capacity input, in its reduced canonical form

HalfGreedySkippedCriticalQuotientSupply := $\forall c \geq 4$ skipped, $2^{(2c-2)-1} \leq$ `localPrefixQuotient(insert c (halfGreedyPrefixSupport(c − 1)))(2c − 2)`

(Lean source), at essentially every step ([Open], scale:cofinal). This is a strictly harder producer than (O3-supply) because it is pinned to the greedy prefix rather than an arbitrary core.

12.5 Target O4 — Reset anti-concentration (sqrt-escape) and the R-run form

Definition 12.7 (Reset $\sqrt{\text{-escape}}$ from row 10)**.**

$$\text{SQRTESC} \;:=\; \forall r \geq 10,\ |\mathrm{rem}(r+1) - 2^{r+1}| \;>\; 2^{(r+5)/2} \qquad \big(\text{equivalently } \Delta^2_{r+1} > 2^{r+5}\big).$$

[Math, fleet-audited], coordinate coord:deviation-affine / crossing-cell geometry. The exponent $(r+5)/2$ is *derived*, not fitted: a violation forces an R-branch crossing cell (s, d) with $s \geq 14$, late row $d \geq 10$ with run length $L = (d-3)/2$ or $(d-4)/2 \geq 3$; dividing the exact affine iterate $w_{d+1+L} = 4^L w_{d+1} - C_L$ by 4^L and absorbing the charge bound $1 \leq C_L/4^L < (2s+4)/3$ yields the exact exponent. This target is *equivalent*, by a Theorem-A composition already Lean-landed on the consumer end, to

$$\text{SQRTESC} \implies \text{LargestSkipLateStepSocket} \implies \text{HALF}$$

(Lean source, [Lean] for the second arrow, [Math, fleet-audited] for the first). SQRTESC is equivalently stated as the R-run length bound

$$\text{RUNBOUND} \;:=\; \forall \text{ reset } r \geq 31, \quad L_r \;<\; \frac{r-3}{2}$$

where L_r is the length of the pure-R run following reset r. RUNBOUND $\Leftrightarrow$ SQRTESC (equivalence (III) of the reset-crossing unification note, [Math]). $r = 7$ is the sole exception below row 10 (dev $= +9$, threshold 64) and is out of scope for both forms; this is why both definitions are stated from $r \geq 10$ or $r \geq 31$.

Certified evidence, not a proof. Exact integer arithmetic to row 200,000: TARGET holds at every reset row except $r = 7$; the maximum observed R-run is 19, at row 158,096, against a required threshold there of $\approx 79{,}046$ — a margin of $\sim 4000\times$ ([Cert], scale:fixed). This is a finite verified list, *not* a cofinal supply, and must never be read as one: it is the exact instance of the quantifier trap this programme's own doctrine warns against.

Observation 12.8 (Sign law — a strictly weaker, still-open sub-target). At all 1209 observed resets, the sign of the deviation is perfectly predicted by branch type (M⇒dev> 0, U⇒dev< 0, zero exceptions), and the margin over the threshold grows from 1.11 bits at $r = 14$ to $\sim$ 1246 bits by row 2500 ([Cert]). Proving the sign law unconditionally — M $\Rightarrow$ $\mathrm{rem}(r+1) > 2^{r+1}$; U $\Rightarrow$ $\mathrm{rem}(r+1) < 2^{r+1}$ — is a one-step branch-algebra statement in the coordinates of the already-proved affine recurrence, and would convert SQRTESC from one two-sided bound into two one-sided bounds. It is flagged in the source material itself as "worth attempting as a lemma before the inequality itself," and it remains [Open].

Observation 12.9 (Dangerous-reset rigidity — Theorem B, not a proof). If reset r is *dangerous* ($|\mathrm{rem}(r+1) - 2^{r+1}| \le 2^{(r+5)/2}$) and the preceding reset r_0 has a pure R-run of length $L' = r - r_0 - 1 > (r+5)/4$, then w_{r_0+1} is confined to at most two adjacent integers, determined by the divisor-pulse stream alone. A chain of n consecutive dangerous resets with long runs forces $n-1$ nested exact congruence towers. This shows the failure mode is a tower of exact integer coincidences, not drift; it is [Math], and it does *not* by itself exclude the two pinned values — an isolated dangerous reset after a short run is untouched by this theorem.

12.6 Target O5 — The remaining $-2, -1$ middle-cell exclusion

Definition 12.10. `SeamTwoSidedDyadicCellEscape` excludes three integer values of the middle coordinate, $4 \cdot \mathrm{rem}(s) - \mathrm{belowPulse}(s) - 4 \notin \{-3, -2, -1\}$, plus one right-pulse-leak bound $4 \cdot \mathrm{overshoot}(s) + \mathrm{abovePulse}(s) \le 2^{s+2}$ under $\mathrm{overshoot}(s) \le 2^s$ (Lean source).

Granted this, induction from a `decide`-checked row-5 base propagates $\forall s \ge 5$, $\mathrm{rem}(s) \le 2^s \vee \mathrm{overshoot}(s) \le 2^s$ (Lean source, [Lean] given the hypothesis, scale:uniform in the consequence). One of the three cells, -3, is already closed unconditionally in the all-right-tail configuration (Lean source, [Lean]). The two remaining cells, -2 and -1, are [Open]: what would close O5 is their exclusion for the actual seam orbit (not a hypothetical one), together with the right-pulse-leak inequality. This is anti-concentration at *constant* resolution (three named integers), sitting far below SQRTESC's resolution $2^{(r+5)/2}$ — it is not known whether closing O5 says anything about O4; no transport between the two has been attempted in the corpus.

A distinct, exponentially cheaper tail-dominance socket targets the same downstream payoff via the upper-reset branch alone: SeamUpperResetDyadicBandEscape := $\forall d \ge 13, \forall j \le d,\ 2^{d-j+1} < \mathrm{resetCharge} \vee \mathrm{resetCharge} + 2(d+j) \le 2^{d-j+1}$ (Lean source, [Open], proved [Cert] through row 31 only by exhaustive computation (Lean source)). This band has *linear* width $O(d)$, not dyadic $O(2^{d/2})$ width — exponentially weaker than SQRTESC yet sufficient for the same conclusion HALF via Lean source. It is proved logically equivalent to a single-critical-index check per row (Lean source, [Lean]), which is a genuine quantifier-complexity reduction, not a strength reduction.

12.7 Target O2b — Membership of $1/21$

The second rational target has an exact current frontier:

$$\frac{1}{21} \in \mathcal{A} \quad \Longleftrightarrow \quad \neg\, \texttt{TwentyOneFatalAlignedBranch}.$$

Lean source [Lean]. The fatal branch means a finite fatal greedy gap followed by cofinite exact selection, eventual quotient/rational alignment, and eventual visits to every doubling block. If it occurs, the canonical quotient remainder is eventually strictly larger than 2^R and follows one explicit affine recurrence: Lean source and Lean source. Any unbounded sequence of closed returns instead proves membership.

No finite support on ranks at least 2 sums to $1/21$ (Lean source). Thus membership would already be an infinite-support rational counterexample to U257. The missing step

is to exclude the fatal aligned affine-supercapacity branch (or force unbounded closed returns); neither is proved. [Open]

12.8 Implication table

Read left to right: an entry ⇒ means the row target implies the column target unconditionally (Lean or fully checked math); ⇔ means proved equivalent; "⊆ weaker" means the row is a strictly weaker but currently unconnected sufficient condition; "?" means no transport has been attempted in the corpus.

Target	Relation to its endpoint	Status
O1 U257 (universal)	O1 ⇒ everything below is moot	[Open]
O2 HALF	pivot target; ¬O1 witness once true	[Open]
O2b 1/21 membership	fatal-aligned branch exclusion; ¬O1 witness once true	[Open]
CPGS (Prop. 12.3)	⇔ O2	[Open], no shortcut
Terminal-only strip (Prop. 12.4)	⇔ O2	[Open], no shortcut
mobiusCenteredHalfCarry sqrtBound (for $A = G$ only)	⇔ O2	[Open], no shortcut
O3 CofinalExactLocalMersenneHalfRows	⇒ O2 (strict, Lean)	[Open]
(O3-supply) linear-width capacity band	⇒ O3 (Lean, given lemma)	[Open]
HalfGreedySkippedCriticalQuotientSupply	⇒ O3 (Lean)	[Open], strictly harder than O3-supply
O4 SQRTESC / RUNBOUND	⇔ each other; ⇒ O2 (Lean fan-in)	[Open]
O4 sign law	⊆ weaker than O4, untested transport	[Open]
O4 dangerous-reset rigidity (Thm B)	partial rigidity, no exclusion	[Open]
O5 SeamTwoSidedDyadicCellEscape (−2, −1 cells)	⇒ two-sided bound (Lean, given hyp.)	[Open]
O5 SeamUpperResetDyadicBandEscape	⇒ O2 (Lean), linear-width	[Open], cert. to $r = 31$ only

12.9 What kind of new mathematics each form would need

None of the following is offered as a plan of attack, and none is close. Each is stated to make the shape of the missing ingredient legible.

O1 (universal). A producer for Cert(A) at an arbitrary infinite A is a statement about geometrically-weighted divisor sums achieving controlled residues modulo powers of two on prescribed dyadic blocks, uniformly over all infinite supports. No general result of this shape exists in the analytic-number-theory literature surveyed for this programme; the corpus's own honest assessment is that this is a genuinely new near-integer anti-concentration theorem for a divisor-weighted sum, not a variant of an existing one.

O2/O3/O5 (Mersenne-specific). The sharp-capacity gap $(*) \to (**)$ and the $-2, -1$ middle-cell exclusion are both, at bottom, requests for anti-concentration of an explicit divisor-count quantity (`localBinarySuffix`, or $4 \cdot \mathrm{rem} - \mathrm{belowPulse} - 4$) away from a short, explicitly named integer window. These are combinatorial/arithmetic near-misses rather than analytic ones — the corpus's own diagnosis is that a congruence or pulse-parity obstruction, not an equidistribution theorem, is the natural tool, and no such obstruction has been found.

O4 (sqrt-escape). The required input is a near-integer anti-concentration bound for $\sum_{i \le L} \delta_n(M + i)2^{-i}$, a geometrically weighted divisor sum in a short interval, at resolution $2^{(r+5)/2}$. This is stated in the source material to not exist in the literature in any form, and it is explicitly noted to be the same missing ingredient, on the opposite side of the Erdős #249/#257 pair, as the totient-window discrepancy anti-concentration input needed to close #249's certificate-supply obligation (Part IV). The two open problems, in their Lean-nearest forms, terminate at structurally the same missing analytic statement, applied to two different arithmetic functions (a Mersenne binary-digit count versus Euler's totient).

12.10 Summary

Erdős #257 is open in the universal form (O1), and both rational membership targets 1/2 (O2) and 1/21 (O2b) remain open. Every route surveyed above either (a) is proved equivalent to O2 and therefore offers no reduction, (b) is a genuine strict weakening (O3, O4, O5, O3-supply) whose own supply is unproved at cofinal scale, or (c) is certified only on a finite initial segment, however large a margin that segment exhibits. Nothing in this paper closes any of these gaps, and no claim above should be read as progress toward doing so.

Statements and declarations

Artefact and data availability. The linked Lean declarations, fixed toolchain, and library manifest are published in the companion repository. This manuscript provides navigation and exposition 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.

Authority boundary. Kernel checking establishes that a proposition was proved; it does not authorise the exposition, literature judgements, or any claim that Problem 257 is solved.

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