

EIGHT PRIME DIVISORS IN ODD DISTINCT COVERING SYSTEMS

MICHAEL SCHROEDER

ABSTRACT. We prove that the least common multiple of the moduli of a finite covering system with distinct odd moduli greater than one has at least eight distinct prime divisors, uniformly in the primes and their exponents. A capped-gain lift converts a normalized one-fibre bound into an increasing supermodular functional on individual congruence labels. Injective reindexing of complete exponent types gives a scalar upper envelope. A certificate of 54 rational inequalities and six affine identities yields $124847/125000 < 1$, with exact affine tails. We also prove a method-specific ceiling: on the first eight odd reference primes, the exact scalar value is at least $1001/1000$ for every admissible real threshold vector and block order, and adding admissible blocks cannot restore a value below one. Separate Lean 4 companions verify the covering theorem and this obstruction to the specified scalar architecture, not to other proofs of higher rank bounds.

1. INTRODUCTION

A *covering system* is a finite family

$$\mathcal{C} = \{a_i \pmod{m_i} : 1 \leq i \leq s\}$$

whose union is $\mathbb{Z}$. Here $a_i \in \mathbb{Z}$ and $m_i > 1$ are integers. It is *distinct* if its moduli are pairwise different, and *odd* if every modulus is odd. Write

$$N = \text{lcm}(m_1, \dots, m_s), \quad \omega(N) = |\{p : p \text{ is prime and } p \mid N\}|.$$

Covering systems were introduced by Erdős in 1950 [9]. The still-open Erdős–Selfridge problem (Erdős Problem 7 [21]) asks whether an odd distinct covering system exists. We establish the following necessary condition.

Theorem 1.1. *Every finite odd distinct covering system with moduli greater than one satisfies $\omega(N) \geq 8$.*

Date: 6 September 2026; revised 14 September 2026.

2020 Mathematics Subject Classification. Primary 11B25; Secondary 11A07, 05D40.

Key words and phrases. Erdős–Selfridge odd covering problem, Erdős Problem 7, covering systems, prime divisors, supermodularity, distortion method, formal verification.

Zenodo edition 6.1. DOI: 10.5281/zenodo.22747841.

Author ORCID: 0009-0004-3249-0195. Paper and prose documentation: CC BY 4.0; original verification code: MIT..

Corollary 1.2. *Every such covering system satisfies*

$$N \geq 3 \cdot 5 \cdot 7 \cdot 11 \cdot 13 \cdot 17 \cdot 19 \cdot 23 = 111,546,435.$$

Proof. The product of any eight distinct odd prime divisors of N divides N and is at least the product of the first eight odd primes. $\square$

There is no assumption that the moduli are square-free, that they form a complete divisor list, or that the cover is minimal or disjoint. The exponents and residue classes are unrestricted. The theorem is not a resolution of the odd covering problem: it excludes support on at most seven primes.

The odd-covering question appears explicitly in Erdős's 1965 paper [20] and has remained open for more than sixty years. The historical account in [4] identifies that paper as the apparent first statement of the question; the 1950 date concerns the introduction of covering systems.

Theorem 1.1 strengthens the author's earlier unrestricted bound $\omega(N) \geq 7$ [19], whose preprint is archived on Zenodo. That argument uses digitwise distortion and a reverse causal gate comparison to bound five charges. Here the capped-gain lift permits a different, whole-block induction on fixed full-label functionals. The proof below is self-contained: the earlier rank-seven theorem is credited as antecedent work, not assumed as a lemma.

Our second result identifies a limitation of the resulting scalar relaxation, rather than of the covering problem itself.

Theorem 1.3 (Scalar-method ceiling). *For the architecture with one fixed threshold per prime, unit ending charge, whole prime blocks, and the saturated single-index envelope $K \mapsto K(d+1)$, set $V(\mathbf{p}, \mathbf{t}) = (T_{p_1, t_1} \cdots T_{p_r, t_r} 0)(1)$, using the exact operators (26), terminal function zero, and the leftmost operator outermost. For every ordering $\mathbf{p}$ of the first eight odd primes and every real vector satisfying $0 \leq t_i \leq p_i - 3$,*

$$V(\mathbf{p}, \mathbf{t}) \geq \frac{1001}{1000} > 1.$$

The same lower bound holds after inserting any finite number of additional admissible prime blocks, in arbitrary positions.

Thus this architecture reaches the uniform rank bound eight but not nine, and cannot recover a subunit bound at a later reference rank. The theorem does not assert the existence of an eight-prime covering system, or that $\omega(N) \geq 9$ is false. It does not rule out sharper residue-sensitive, multi-statistic, joint-block, or history-dependent arguments, nor larger reference primes justified by a case split on the small actual primes.

1.1. Context and scope. Balister's survey [2] gives a broad account of the classical questions and modern probabilistic methods for covering systems. Hough resolved the minimum modulus problem by proving a universal upper bound for the smallest modulus of a distinct covering system [14]. Hough and Nielsen subsequently proved that every distinct covering system has a modulus divisible by 2 or 3 [15]. For an odd system this forces $3 \mid N$; it

does not require 3 itself to be a modulus. The distortion method of Balister, Bollobás, Morris, Sahasrabudhe and Tiba controls uncovered density and provides further divisibility restrictions [3]. Our normalized local reweighting belongs to this line of methods.

The square-free setting must be distinguished from the general problem. Guo and Sun proved a lower bound of 22 prime divisors under the assumption that every modulus is square-free [11]. Balister, Bollobás, Morris, Sahasrabudhe and Tiba later excluded odd distinct square-free covering systems altogether [4]. Neither result assumes arbitrary prime powers as Theorem 1.1 does. Work on odd coverings in which one modulus may repeat [13, 6] studies a different hypothesis and does not contradict this theorem.

Other recent directions include the structure and enumeration of minimal covers [5], reciprocal-sum gaps [10], overlap restrictions on complete divisor lists [1], and systems with prime divisors restricted to 2, 3, 5 [12]. In formal verification, Mian and Siddique give a Lean proof of the finite exclusion $N > 10000$ [17]; this is a bound on the size of N , not a uniform exclusion by the number of its prime factors. We cite that work as a reported formalization, not as a dependency of the present proof.

The rearrangement inequality used below is a finite, bounded form of the classical extremal property of comonotone dependence; see Puccetti and Wang [18]. We give a short proof. All random objects in the covering argument have finite support, so the distinctions concerning infinite means discussed by Côté and Wang [7] do not arise. The contribution here is the capped-gain lift and its full-label use in a seven-block obstruction, together with a small rational certificate, an end-to-end formalization, and a separately verified ceiling for the associated scalar architecture. No priority claim about all unpublished work is needed for the argument. The literature discussion reflects the original manuscript's 6 September 2026 cutoff.

1.2. Proof architecture. The Chinese remainder theorem converts a hypothetical cover into a cover of a product of prime-power prefix spaces. Distinct moduli become distinct complete depth vectors. We first avoid every cylinder supported on only one coordinate. Six normalized fibre bounds are then composed using fixed increasing supermodular functionals on the remaining labels. A seventh coordinate supplies a terminal covering obstruction.

Two distinctions are essential. First, the continuation functional is fixed before any fibre outcome is selected; it never depends on an unresolved comparison run. Second, labels with different complete depth vectors are never merged because they share a residue, a depth, or an earlier projection. A scalar bound is applied only after the full-label functional has been justified. The final contradiction is

$$(1) \quad 1 \leq F_1(A) \leq \frac{124847}{125000} = 1 - \frac{153}{125000} < 1.$$

Sections 2–8 give a self-contained finite proof of Theorem 1.1. Sections 9–11 prove Theorem 1.3, explain its scope, and give a retuned scalar witness. Section 12 describes the two formal counterparts.

2. PREFIX SPACES AND REMOVAL OF PURE CYLINDERS

For integers $P \geq 2$ and $a \geq 0$, a word of length a over $\{0, \dots, P-1\}$ represents its base- P digits, least significant digit first. A depth- d prefix, $0 \leq d \leq a$, fixes its first d digits. Depth zero imposes no restriction. A *cylinder* in a product of word spaces fixes one prefix in each coordinate. Its *type* is the vector of prefix depths. A cylinder is *pure* if precisely one depth is positive, and *mixed* if at least two are positive.

Lemma 2.1 (Arithmetic reduction and padding). *If an odd distinct covering system has $\omega(N) \leq 7$, there is a finite covering of a product of seven word spaces by cylinders with distinct nonzero types, with alphabet sizes*

$$(2) \quad P_j \geq p_j, \quad (p_1, \dots, p_7) = (3, 5, 7, 11, 13, 17, 19).$$

The heights of these spaces are arbitrary nonnegative integers.

Proof. List the distinct prime divisors of N as $P_1 < \dots < P_r$ and write $N = \prod_{j=1}^r P_j^{a_j}$. Every modulus divides N and therefore has a unique exponent vector $(e_1, \dots, e_r)$ with $0 \leq e_j \leq a_j$. The Chinese remainder theorem identifies residues modulo N with the product of the r word spaces. A congruence modulo $\prod P_j^{e_j}$ is exactly a cylinder of type $(e_1, \dots, e_r)$. Nontriviality excludes the zero vector, and distinctness gives injectivity of the type map. Coverage of $\mathbb{Z}$ gives coverage of this finite product.

The j th odd prime is at least the j th entry of (2). For $r < 7$, append coordinates $j = r+1, \dots, 7$ with alphabet size p_j and height zero. Every cylinder has zero depth in them. Such coordinates have a singleton word space, so this preserves coverage, nontriviality, and type injectivity. The case $r = 0$ is included: no nonzero type could then occur in a cover. The appended alphabet sizes need not be additional prime divisors of N . $\square$

Fix a hypothetical cylinder family as in Lemma 2.1. For each coordinate and each positive depth choose one forbidden prefix: use the prefix of the pure cylinder of that depth if it exists, and choose an arbitrary prefix otherwise. There is at most one such pure cylinder, by type injectivity.

Lemma 2.2 (Safe fibres). *In an alphabet of size $P \geq 3$ and height a , let Y be the words avoiding the chosen prefix at every depth. Put*

$$\theta = (P-1)|Y|/P^a, \quad \beta_P(d) = (P-1)P^{-d} \quad (d \geq 1).$$

Then Y is nonempty, $\theta \geq P-2$, and the uniform law μ on Y assigns a depth- d prefix probability at most $\beta_P(d)/\theta$.

Proof. A depth- d prefix excludes P^{a-d} words. The union bound gives

$$|Y| \geq P^a - \sum_{d=1}^a P^{a-d} = \frac{(P-2)P^a + 1}{P-1} > 0.$$

Every prefix has at most P^{a-d} representatives in Y . Dividing by $|Y|$ proves the probability bound. For $a = 0$ the sum is empty and the word space is a singleton. $\square$

Use these safe fibres in coordinates $1, \dots, 6$. In coordinate 7, use a slightly different safe set. At digit position d , forbid the last digit of the chosen depth- d prefix, regardless of the preceding digits. Every word in the resulting product avoids every forbidden prefix. The uniform law λ_7 on this product assigns any depth- d prefix probability either zero or $(P_7 - 1)^{-d}$, hence at most 18^{-d} . This remains valid at height zero. The product of the seven safe sets is nonempty, avoids all pure cylinders, and would therefore be covered by the original mixed cylinders.

3. A FINITE REARRANGEMENT BOUND

For a finite label universe U , a function $F : 2^U \rightarrow \mathbb{Q}$ is *increasing* if $A \subseteq B$ implies $F(A) \leq F(B)$, and *supermodular* if

$$F(A) + F(B) \leq F(A \cap B) + F(A \cup B).$$

Equivalently, its single-label increments are increasing in the base set. We always keep the individual labels, including when their events coincide.

Lemma 3.1 (Nested marginal majorization). *Let X be a random subset of $\{1, \dots, n\}$ on a finite probability space, and suppose $\Pr(i \in X) \leq v_i$, where $0 \leq v_i \leq 1$ for each i . If F is increasing and supermodular, then, with $C_i = \{1, \dots, i\}$ and $C_0 = \emptyset$,*

$$\mathbb{E}F(X) \leq F(\emptyset) + \sum_{i=1}^n v_i (F(C_i) - F(C_{i-1})).$$

Proof. For each deterministic A , add its elements in increasing order. The base set before adding i is a subset of C_{i-1} , so increasing increments give

$$F(A) \leq F(\emptyset) + \sum_{i \in A} (F(C_i) - F(C_{i-1})).$$

Take expectations. Each coefficient is nonnegative by increasingness, so the marginal upper bounds can be substituted. No independence is used. $\square$

No ordering of the marginal bounds is needed. In the depth application below, labels are ordered by depth to align the chain with its level sets.

Fix $p \geq 3$ and an integer $D \geq 1$. Give each label a depth in $\{0, \dots, D\}$ and put $S_d = \{x \in S : e(x) \leq d\}$. Write $\beta_d = (p-1)p^{-d}$ for $d \geq 1$ and

define

$$\begin{aligned} F_0(S) &= F(S_0), \\ (3) \quad G(S) &= \sum_{d=1}^D \beta_d (F(S_d) - F(S_{d-1})), \\ (4) \quad \Phi_r(S) &= F_0(S) + rG(S). \end{aligned}$$

Summation by parts gives the finite positive-mixture identity

$$(5) \quad \Phi_r(S) = (1 - r\beta_1)F(S_0) + r \sum_{d=1}^{D-1} (\beta_d - \beta_{d+1})F(S_d) + r\beta_D F(S_D).$$

If $0 \leq r\beta_1 \leq 1$, its coefficients are nonnegative and sum to one. Thus Φ_r is increasing and supermodular whenever F is. Lemma 3.1, with labels sorted by depth, shows that

$$(6) \quad \mathbb{E}F(X) \leq \Phi_r(S)$$

whenever $X \subseteq S$ and every positive-depth label has probability at most $r\beta_{e(x)}$. Depth-zero labels need only the bound one.

Lemma 3.2. *The function G in (3) is nonnegative and increasing. It need not be assumed supermodular.*

Proof. For fixed $I \subseteq J$, supermodularity and increasingness give, whenever $A \subseteq B$,

$$F(A \cap J) - F(A \cap I) \leq F(B \cap J) - F(B \cap I).$$

Indeed, apply supermodularity to $A \cap J$ and $B \cap I$; their intersection is $A \cap I$ and their union is contained in $B \cap J$. Apply this to successive depth sets and sum with the nonnegative coefficients β_d . Nonnegativity follows from increasingness of F . $\square$

4. THE CAPPED-GAIN LIFT

Let $q > t \geq 0$, put $c = q - t$, and suppose $c \geq 1$. For $L \geq 0$ set

$$r(L) = \frac{1}{q - \min(t, L)}, \quad Q(L) = \frac{(L - t)_+}{c}, \quad (x)_+ = \max(x, 0).$$

Assume that $G \geq 0$ is increasing, and that $F_0 + rG$ is increasing and supermodular for every $0 \leq r \leq 1/c$. Define

$$(7) \quad H(L, S) = Q(L) + \max \left\{ F_0(S) + r(L)G(S), F_0(S) + \frac{G(S)}{c} - (1 - cr(L)) \right\}.$$

Theorem 4.1 (Capped-gain lift). *Partition a finite label universe into ending labels and future labels. Assign nonnegative rational weights to the ending labels, and let $L(A)$ be their total weight in A . Then*

$$A \longmapsto H(L(A), A \cap \{\text{future labels}\})$$

is increasing and supermodular. It majorizes $Q(L) + F_0(S) + r(L)G(S)$. If $L \geq t$, it simplifies to

$$(8) \quad H(L, S) = \frac{L - t}{c} + F_0(S) + \frac{G(S)}{c}.$$

Proof. Majorization and (8) follow directly from the definition. The identity

$$(9) \quad H(L, S) = Q(L) + F_0(S) + G(S)/c - (1/c - r(L)) \min(G(S), c)$$

follows by splitting at $G(S) = c$. Put $A_0 = \min(G(S), c)$. Up to a constant independent of L , the load dependence is

$$f(L) = \begin{cases} A_0/(q - L), & L \leq t, \\ (L - t + A_0)/c, & L \geq t. \end{cases}$$

This is increasing and convex, as can be proved over $\mathbb{Q}$ without derivatives. For $x, y \leq t$ the supporting-line remainder is

$$\frac{A_0}{q - y} - \frac{A_0}{q - x} - \frac{A_0(y - x)}{(q - x)^2} = \frac{A_0(y - x)^2}{(q - x)^2(q - y)} \geq 0.$$

Its supporting slope is nonnegative and at most $1/c$, because $0 \leq A_0 \leq c \leq q - x$. For $y \leq t$ the high-side line lies below the low branch, since

$$\frac{A_0}{q - y} - \frac{A_0 + y - t}{c} = \frac{(t - y)(q - y - A_0)}{c(q - y)} \geq 0.$$

These inequalities also supply supporting lines across the join; hence f is convex. Convexity implies increasing fixed-width increments. This proves the required ending–ending increments.

For fixed L , the two branches in (7) are increasing supermodular functions of S . Their second-minus-first difference is

$$(1/c - r(L))(G(S) - c),$$

which is increasing. The maximum of two supermodular functions with increasing difference is supermodular. To see the only nontrivial case, let f dominate at A and g at B . Supermodularity of g , followed by increasingness of $g - f$ from $A \cap B$ to A , gives $f(A) + g(B) \leq f(A \cap B) + g(A \cup B)$; bound the right side by the maxima. This proves the future–future increments and increasingness in future labels.

Finally, for $L_2 \geq L_1$, (9) gives

$$H(L_2, S) - H(L_1, S) = Q(L_2) - Q(L_1) + (r(L_2) - r(L_1)) \min(G(S), c).$$

The right side is increasing in S . This proves the mixed increments. For the resulting set function write $\Delta_x(A)$ for the increment on adjoining label x , and let U_F be the future-label universe. If $A \subseteq B$ and $x \notin B$, the three cases together give

$$\Delta_x(A) \leq \Delta_x(A \cup (B \cap U_F)) \leq \Delta_x(B) :$$

the first comparison adds future labels and the second adds ending labels, using the appropriate pure or mixed increment inequality. Thus the increments increase on the whole label set, proving supermodularity. $\square$

Remark 4.2 (Why the compensation is needed). The first branch alone need not be supermodular. Take $q = 3, t = 1, F_0 = 0, G = 4$, and two ending labels of weight $4/5$ each. The function $Q(L) + r(L)G$ has second difference

$$\frac{4}{3} + \frac{23}{10} - 2\frac{20}{11} = -\frac{1}{330}.$$

The compensated lift instead has second difference $49/330$. These data arise from $p = 5$ and a future label of depth one with $F(T) = 5|T|$. Thus even physical first-depth weights do not justify omitting the second branch. At the first two stages of our proof, $t = 0$ and (8) always applies.

5. THE NORMALIZED ONE-FIBRE INEQUALITY

Consider a safe fibre from Lemma 2.2, with actual alphabet size $P \geq p \geq 3$. Fix $0 \leq t \leq p - 3$, and put $q = p - 2, c = p - 2 - t \geq 1$. Some labels end in this coordinate, and the others have a nonzero future suffix. Let $E \subseteq Y$ be the union of the ending-prefix events, let $\alpha = \mu(E)$, and put

$$L = \sum_{x \text{ ending}} \beta_P(e(x)), \quad \delta = t/\theta.$$

All ending depths are positive. The union bound gives $\theta\alpha \leq L$; also, δ lies in $[0, 1)$.

Define a probability law ν by its density relative to μ :

$$(10) \quad \frac{d\nu}{d\mu}(y) = \begin{cases} 0, & y \in E, \alpha \leq \delta, \\ (\alpha - \delta)/(\alpha(1 - \delta)), & y \in E, \alpha > \delta, \\ 1/(1 - \alpha), & y \notin E, \alpha \leq \delta, \\ 1/(1 - \delta), & y \notin E, \alpha > \delta. \end{cases}$$

Every denominator used in its branch is nonzero. Direct summation shows

$$(11) \quad \nu(Y) = 1, \quad \nu(E) = \frac{(\alpha - \delta)_+}{1 - \delta}, \quad \frac{d\nu}{d\mu} \leq \frac{1}{1 - \min(\alpha, \delta)}.$$

These statements include $\alpha = 0$ and $\alpha = 1$.

Lemma 5.1 (Uniform actual-fibre bound). *Let S be the future labels and let $S(y)$ be those whose current prefix matches y . For every fixed increasing supermodular continuation F on these labels, with depths at most D , the lift formed from (3)–(4) at reference alphabet p satisfies*

$$(12) \quad \nu(E) + \mathbb{E}_\nu F(S(y)) \leq H(L, S).$$

Proof. By (11),

$$\nu(E) = \frac{(\theta\alpha - t)_+}{\theta - t} \leq \frac{(L - t)_+}{p - 2 - t} = Q(L).$$

A future prefix of positive depth d has ν -probability at most

$$\frac{\beta_P(d)}{\theta - \min(\theta\alpha, t)} \leq \frac{\beta_P(d)}{P - 2 - \min(L, t)} \leq r(L)\beta_p(d).$$

For the last inequality put $z = \min(L, t)$. For $P \geq p$ both positive factors in

$$(13) \quad \frac{P-1}{P^d(P-2-z)} = P^{-d} \left(1 + \frac{1+z}{P-2-z} \right)$$

are nonincreasing in P . This proves the comparison without changing the ending weights $\beta_P(e)$. Since $r(L) \leq 1/c$ and $\beta_p(1) < 1$, Lemma 3.1 applies and bounds the continuation by $F_0(S) + r(L)G(S)$. The lift majorizes this plus $Q(L)$. $\square$

The actual ending load is important: $\beta_P(1) = 1 - 1/P$ increases with P . Replacing all actual weights by those for p is not justified. Only the combined ratio in (13) is compared.

6. FIXED FULL-LABEL CONTINUATION FUNCTIONALS

Choose one integer $D \geq 9$ at least as large as every coordinate height in the fixed finite family. This is a bookkeeping parameter, not an upper bound in the theorem. Each mixed label retains its full type and its residue prefixes. Its *ending coordinate* is its last positive depth. At stage j , labels ending before j are no longer present.

Define F_7 on subsets of terminal labels by

$$(14) \quad F_7(A) = \frac{1}{8} \left(\sum_{x \in A} 17 \cdot 18^{-e_7(x)} - 9 \right)_+.$$

For $j = 6, \dots, 1$, partition A into its ending labels $E_j(A)$ and future labels $S_j(A)$. Set

$$(15) \quad (t_1, \dots, t_6) = (0, 0, 1, 3, 5, 7), \quad (c_1, \dots, c_6) = (1, 3, 4, 6, 6, 8).$$

Let $L_j(A) = \sum_{x \in E_j(A)} \beta_{P_j}(e_j(x))$. Use $F = F_{j+1}$ and the j th depths in (3), and define

$$(16) \quad F_j(A) = H_{p_j-2, t_j}(L_j(A), S_j(A)).$$

All these functions are fixed by the labelled family and its depth data. Their definitions do not select a continuation after observing a random comparison outcome. Increasing D above the actual maximum depth does not change them, since the added differences in (3) vanish.

Proposition 6.1. *Each F_j is nonnegative, increasing and supermodular, vanishes at the empty set, and is unchanged by a bijective relabelling preserving depths.*

Proof. The terminal function is an increasing convex function of a nonnegative modular sum, so it has these properties. For the induction step, (5) supplies increasing supermodular Φ_r for all $0 \leq r \leq 1/c_j$. Lemma 3.2 and Theorem 4.1 apply. The first branch is nonnegative. At the empty set it is

zero, and the second branch is nonpositive. The formulas depend only on sets of labels, their partitions and their depths, which proves the relabelling assertion by induction. $\square$

Proposition 6.2 (Covering forces value at least one). *If the active labels A cover the product of the remaining safe spaces at stage j , then $F_j(A) \geq 1$.*

Proof. At $j = 7$, the terminal prefix probability bound and the union bound give $\sum_{x \in A} 18^{-e_7(x)} \geq 1$. The sum in (14) is therefore at least 17, so $F_7(A) \geq 1$.

Suppose the assertion holds at $j + 1$. For each actual safe word y in coordinate j , either an ending label covers the entire suffix, or the active future labels $S_j(A)(y)$ cover it. Thus, pointwise,

$$1 \leq \mathbf{1}_E(y) + F_{j+1}(S_j(A)(y)).$$

Integrate against the normalized law (10) for this actual state. Lemma 5.1 bounds the right side by $F_j(A)$. This is a finite backward induction, including states of zero probability under any other law. It needs neither a conditioning on an artificial history nor a probability measure on infinitely many digits. $\square$

7. A SCALAR ENVELOPE WITHOUT MERGING LABELS

Definition 7.1. At stage j , a finite label family is *K-admissible* if it has an index map into a finite set I of size $K \geq 1$, with distinguished element 0, such that the map

$$x \mapsto (i(x), e_j(x), \dots, e_7(x))$$

is injective, every displayed suffix is nonzero, and a label with index 0 has at least two positive coordinates in that suffix.

The index records a complete preceding depth vector. It does not count OR-activated projections. Unused indices are permitted. The root mixed family is 1-admissible with index set $\{0\}$.

Let $u : \mathbb{N}_{\geq 1} \rightarrow \mathbb{Q}$ have an affine tail $u(K) = aK + b$ for $K \geq 10$, where $a \geq 0$. Define the finite scalar gain

$$(17) \quad G_{p,D}[u](K) = \sum_{d=1}^D \beta_p(d)(u(K(d+1)) - u(Kd)),$$

$$(18) \quad \overline{G}_p[u](K) = G_{p,9}[u](K) + aK/p^9.$$

For every $D \geq 9$ and $K \geq 1$,

$$(19) \quad G_{p,D}[u](K) + aK/p^D = \overline{G}_p[u](K).$$

Indeed, for $d \geq 10$ both arguments lie in the affine tail, so their difference is aK ; the added term is $aK(p-1)/p^d = aK/p^{d-1} - aK/p^d$. Finite telescoping proves the identity. In particular $G_{p,D}[u](K) \leq \overline{G}_p[u](K)$.

For $c = p - 2 - t$ and $r_K = (p - 2 - \min(t, K - 1))^{-1}$ set

$$(20) \quad (T_{p,t}u)(K) = \frac{(K - 1 - t)_+}{c} + \max \left\{ u(K) + r_K \overline{G}_p[u](K), \right. \\ \left. u(K) + \overline{G}_p[u](K)/c - (1 - cr_K) \right\}.$$

Proposition 7.2 (Uniform full-type envelope). *Suppose $u_7(K) = (K - 10)_+/8$, each u_j has a nonnegative affine-tail slope for $K \geq 10$, and*

$$(21) \quad T_{p_j,t_j}u_{j+1}(K) \leq u_j(K) \quad (1 \leq j \leq 6, K \geq 1).$$

For every finite K -admissible family at stage j , $F_j(A) \leq u_j(K)$.

Proof. At stage 7, no label has distinguished index zero. For each other index, injectivity permits at most one label at each positive depth. For any finite depth bound a , $\sum_{d=1}^a 17 \cdot 18^{-d} = 1 - 18^{-a} \leq 1$. Hence the terminal load is at most $K - 1$, proving the assertion there.

At stage $j \leq 6$, ending labels have all later depths zero, positive current depth, and nonzero preceding index. For each of the $K - 1$ possible nonzero indices, their total ending weight is at most $\sum_{d=1}^D (P_j - 1)P_j^{-d} = 1 - P_j^{-D} \leq 1$. Thus $L_j(A) \leq K - 1$. Increasingness in the load allows replacement by $K - 1$ in the lift.

For the future labels of current depth at most d , use the new index set $I \times \{0, \dots, d\}$, distinguished element $(0, 0)$, and the map

$$(22) \quad (i, e_j, e_{j+1}, \dots, e_7) \mapsto ((i, e_j), e_{j+1}, \dots, e_7).$$

This is injective, has $K(d+1)$ possible indices, and preserves admissibility. In particular the distinguished index leaves the original two-positive-coordinate condition entirely in the later suffix. These are future labels, so all later suffixes are nonzero. No labels are identified by this reindexing.

By induction and Proposition 6.1, $F_{j+1}(S_d) \leq u_{j+1}(K(d+1))$ for $0 \leq d \leq D$. Substitute these bounds into the *positive* mixture (5). Its scalar value is $u_{j+1}(K) + rG_{p_j,D}[u_{j+1}](K)$, which is at most $u_{j+1}(K) + r\overline{G}_{p_j}[u_{j+1}](K)$ by (19). Apply this for $r = r_K$ and $r = 1/c_j$ in the two branches. The result is exactly $F_j(A) \leq T_{p_j,t_j}u_{j+1}(K) \leq u_j(K)$. Pointwise bounds have not been substituted separately into signed differences in (17). $\square$

8. THE RATIONAL CERTIFICATE AND PROOF OF THE THEOREM

Set $u_7(K) = (K - 10)_+/8$. For $j \leq 6$ and $1 \leq K \leq 9$, define $u_j(K)$ to be the entry of Table 1 divided by 10^6 . For $K \geq 10$ set $u_j(K) = a_j K + b_j$ with Table 2.

Lemma 8.1 (Exact certificate). *These functions satisfy (21) for every $K \geq 1$, and*

$$u_1(1) = \frac{124847}{125000}, \quad 1 - u_1(1) = \frac{153}{125000} > \frac{1}{1000}.$$

K	$10^6 u_1$	$10^6 u_2$	$10^6 u_3$	$10^6 u_4$	$10^6 u_5$	$10^6 u_6$
1	998776	53684	2444	101	3	1
2	4711146	810347	66470	4978	362	1
3	9003873	2047552	542310	41666	5144	69
4	13502517	3464004	1155363	157161	20437	1298
5	18113840	5064174	1975441	492712	71069	3343
6	22797470	6759597	2889210	910486	188307	27942
7	27507176	8533245	3909651	1464023	465550	58007
8	32217918	10310004	4934332	2022956	749153	95589
9	36944897	12135471	6025435	2666427	1132390	250920

TABLE 1. The 54 rational upper-envelope entries. Column j uses reference prime p_j and threshold t_j ; $u_7(K) = 0$ for $1 \leq K \leq 9$.

j	1	2	3	4	5	6	7
a_j	607/128	479/256	1181/1024	185/256	61/128	17/64	1/8
b_j	-23/4	-19/4	-53/12	-47/12	-13/4	-9/4	-5/4

TABLE 2. Exact affine tails $u_j(K) = a_j K + b_j$ for every $K \geq 10$.

Proof. For $K \geq 10$, all gain differences lie in the affine tail, and (18) gives $\overline{G}_p[u](K) = aK$. Moreover $K - 1 \geq t_j$, so the two branches coincide. The six identities to check are simply

$$(23) \quad a_j = a_{j+1} + \frac{1 + a_{j+1}}{c_j}, \quad b_j = b_{j+1} - \frac{t_j + 1}{c_j},$$

which Table 2 satisfies.

For $1 \leq K \leq 9$, set $m = \lfloor 9/K \rfloor$. Then $m \leq 9$ and $K(m+1) \geq 10$, so the same finite telescoping gives the shorter formula

$$(24) \quad \overline{G}_p[u](K) = \sum_{d=1}^m \beta_p(d)(u(K(d+1)) - u(Kd)) + \frac{aK}{p^m}.$$

Substitution in (20) checks the 54 inequalities by rational arithmetic. Appendix A supplies the complete standalone checker; the Lean companion proves these inequalities with kernel-checked rational normalization, including a separate proof using (24).¹ The nine cutoffs are 9, 4, 3, 2, 1, 1, 1, 1, 1, giving 138 finite summands for the six stages. The root value and margin follow from the first entry of Table 1. $\square$

Proof of Theorem 1.1. Suppose $\omega(N) \leq 7$. Apply Lemma 2.1 and remove the pure cylinders on the nonempty product of safe spaces. The mixed family

¹Among the 54 small-state inequalities, exactly $(j, K) = (1, 4)$ and $(1, 8)$ hold with equality; the other 52 are strict. The entries of Table 1 are exact integers divided by 10^6 , not rounded decimal approximations.

A still covers this product, so Proposition 6.2 gives $F_1(A) \geq 1$. Its full types are distinct and mixed, so it is 1-admissible. Proposition 7.2 and Lemma 8.1 give $F_1(A) \leq 124847/125000 < 1$, a contradiction. Thus $\omega(N) \geq 8$. $\square$

9. THE EXACT SCALAR METHOD

The preceding argument uses scalar functions only to bound a fixed full-label construction. We now study the limitation of this scalar envelope itself. This is a different question from whether a covering exists on a particular prime support.

For an integer $p \geq 3$, put $q = p - 2$ and choose a real threshold $0 \leq t \leq q - 1$. For a function $h : \mathbb{N}_{\geq 1} \rightarrow \mathbb{R}$, define

$$(25) \quad G_p h(K) = \sum_{d=1}^{\infty} \beta_p(d) (h(K(d+1)) - h(Kd)), \quad \beta_p(d) = \frac{p-1}{p^d},$$

whenever the series is absolutely convergent. Set $L = K - 1$, $c = q - t$, and extend (20) by

$$(26) \quad (T_{p,t}h)(K) = \frac{(L-t)_+}{c} + \max \left\{ h(K) + \frac{G_p h(K)}{q - \min(t, L)}, \right. \\ \left. h(K) + \frac{G_p h(K)}{c} - \left(1 - \frac{c}{q - \min(t, L)} \right) \right\}.$$

For an ordered list $\mathbf{p} = (p_1, \dots, p_r)$ and thresholds $\mathbf{t} = (t_1, \dots, t_r)$, its *scalar value* is

$$(27) \quad V(\mathbf{p}, \mathbf{t}) = (T_{p_1, t_1} \cdots T_{p_r, t_r} 0)(1).$$

The leftmost operator is outermost. In particular, $T_{p,t}0(K) = (K - 1 - t)_+ / (p - 2 - t)$: the last operator is precisely the terminal hinge. A subunit upper bound for r reference primes yields a support bound of $r + 1$, not r .

9.1. Positive completion and exact affine tails. The gain in (25) is written as signed differences. To justify order comparisons, including comparisons with artificial lower bounds that need not be increasing, we use a positive completion. For $D \geq 1$ define

$$(28) \quad G_p^{[D]} h(K) = \sum_{d=1}^D \beta_p(d) (h(K(d+1)) - h(Kd)) + \frac{h(K(D+2)) - h(K(D+1))}{p^D}.$$

Write $T_{p,t}^{[D]}$ for (26) with this gain. The completion is defined for every real-valued function.

Lemma 9.1 (Order preservation). *For $0 \leq t \leq p - 3$, the operator $T_{p,t}^{[D]}$ preserves pointwise inequalities. It maps increasing functions to increasing functions, and $T_{p,t}^{[D]}h \geq h$ when h is increasing.*

Proof. For $0 \leq \rho \leq 1/c$, finite summation by parts gives

$$\begin{aligned}
 & h(K) + \rho G_p^{[D]} h(K) \\
 (29) \quad &= (1 - \rho \beta_p(1))h(K) + \rho \sum_{d=1}^{D-1} (\beta_p(d) - \beta_p(d+1))h(K(d+1)) \\
 & \quad + \rho \beta_p(D) \left(\frac{p-2}{p-1} h(K(D+1)) + \frac{1}{p-1} h(K(D+2)) \right).
 \end{aligned}$$

The coefficients are nonnegative and sum to one, since $c \geq 1$ and $\beta_p(1) < 1$. Apply this identity to both branches of (26). Their other terms do not depend on h , proving pointwise order preservation.

For increasing h , the gain is nonnegative and the mixture is increasing in K at fixed ρ . The first branch remains increasing when $\rho = (q - \min(t, K - 1))^{-1}$ increases with K . The second branch is an increasing mixture plus the increasing term $-1 + c/(q - \min(t, K - 1))$. The paid term is also increasing. Finally, the first branch and nonnegative gain give $T_{p,t}^{[D]} h(K) \geq h(K)$. $\square$

If $h(K) = aK + b$ for all $K \geq H$ and $K(D+1) \geq H$, then

$$(30) \quad G_p h(K) = G_p^{[D]} h(K) = \sum_{d=1}^D \beta_p(d) (h(K(d+1)) - h(Kd)) + \frac{aK}{p^D}.$$

Indeed, every remaining increment is aK and the geometric tail has total weight p^{-D} . This also proves absolute convergence. For sufficiently large K , (26) has the affine tail

$$(31) \quad a' = a + \frac{1+a}{p-2-t}, \quad b' = b - \frac{1+t}{p-2-t}.$$

Starting from zero, every successive continuation is increasing, nonnegative and eventually affine. Thus all series in (27) converge, and a sufficiently long finite completion computes the exact value. For primes at most 23, the common affine height $H = 25$ and completion $D = 24$ suffice for every ordering and every admissible threshold. This is an exact tail formula, not a bound on exponents in the covering problem.

9.2. A sharper seven-prime value. The small certificate in Section 8 was selected for ease of inspection. The exact scalar recurrence admits a slightly better value without altering the first six thresholds.

Proposition 9.2. For $\mathbf{p} = (3, 5, 7, 11, 13, 17, 19)$ and $\mathbf{t} = (0, 0, 1, 3, 5, 7, 8)$,

$$(32) \quad V(\mathbf{p}, \mathbf{t}) < 0.998419 < 1.$$

Proof. Evaluate (26) backwards from zero, using the finite formula (30) and the affine update (31). All thresholds and intermediate values are rational. Table 3 in Appendix B gives the exact value. The companion's two seven-prime witness theorems verify that value and the strict rational comparison with $998419/1000000 = 0.998419$. $\square$

This is a verified witness, not a claim of global optimality. The end-to-end integer-covering formalization continues to use the 54-entry certificate and root value 124847/125000 of Section 8.

10. A CONTINUOUS, ALL-ORDERS SCALAR CEILING

The obstruction on eight reference primes requires a lower certificate, in contrast to the upper certificate used for the covering theorem. We first reduce each continuum of threshold choices to an exact algebraic check.

Lemma 10.1 (Exact threshold selector). *Fix real f, g, L, q, ℓ, u with $L \geq 0$ and $0 \leq \ell \leq u < q$. Let $H(t)$ be the right side of (26) with $h(K) = f$ and $G_p h(K) = g$. Its minimum on $[\ell, u]$ is attained at*

$$(33) \quad t_* = \begin{cases} \ell, & q \leq L + g, \\ \max\{\ell, \min\{u, L\}\}, & L + g < q. \end{cases}$$

No sign condition on g is required.

Proof. For $t \leq L$, both branches coincide and $H(t) - f = (L + g - t)/(q - t)$. For $a \leq b \leq L$ in the interval,

$$H(b) - H(a) = \frac{(b - a)(L + g - q)}{(q - a)(q - b)}.$$

For $L \leq t$, write $H(t) - f = \max\{A, B(t)\}$, where

$$A = \frac{g}{q - L}, \quad B(t) = \frac{g}{q - t} - \frac{t - L}{q - L}, \quad B(t) - A = \frac{(t - L)(g - (q - t))}{(q - t)(q - L)}.$$

This maximum is nondecreasing on the part of the interval above L . Indeed, for $L \leq a \leq b$, the assertion is immediate if $B(a) \leq A$. Otherwise $g > q - a$ and

$$B(b) - B(a) = (b - a) \left(\frac{g}{(q - a)(q - b)} - \frac{1}{q - L} \right) \geq 0.$$

All denominators used here are positive. Combining this monotonicity with the sign of $L + g - q$ below L gives (33), including intervals that do not contain L . $\square$

Let $P_8 = \{3, 5, 7, 11, 13, 17, 19, 23\}$ and $\mathcal{B} = \prod_{p \in P_8} [0, p - 3]$. For a rational subbox $B = \prod_{p \in P_8} [\ell_p, u_p]$ and a subset $S \subseteq P_8$, a lower-certificate row is a function $R_{S,B}$ represented by its values at $K = 1, \dots, 24$ and an affine tail for every $K \geq 25$. The certificate conditions are

$$(34) \quad \begin{aligned} R_{\emptyset, B}(K) &= 0, \\ R_{S, B}(K) &\leq T_{p, t}^{[24]} R_{S \setminus \{p\}, B}(K) \quad (p \in S, t \in [\ell_p, u_p], K \geq 1). \end{aligned}$$

The implementation shares rows whenever their remaining interval specifications agree; it need not duplicate all subsets for every box.

Lemma 10.2 (Soundness of lower certificates). *If (34) holds, then for every ordering of S and every fixed threshold choice in B , the corresponding completed composition from zero dominates $R_{S,B}$ pointwise.*

Proof. Induct on $|S|$. If the first prime is p , the induction hypothesis bounds its actual scalar suffix below by $R_{S \setminus \{p\}, B}$. Apply Lemma 9.1 and then the edge inequality in (34). The base case is the zero function. $\square$

In constructing these lower rows one may minimize separately at each state K and each remaining subset. This enlarges the class of choices and therefore gives a *lower relaxation* of every fixed schedule. It does not define a physical threshold policy, and it is never used as a full-label continuation functional. In particular, no admissibility of history-dependent decisions is being inferred from the certificate.

Lemma 10.3 (Kernel-checked eight-prime certificate). *There is a finite rational box partition of $\mathcal{B}$ with rows satisfying (34) and $R_{P_8, B}(1) \geq 1001/1000$ at every leaf box.*

Proof. The companion supplies the complete certificate and a Lean proof of its soundness. Here we specify the finite checks that discharge the claim. Every row has 24 integer entries, an integer slope and an integer intercept, all divided by 10^{12} . Each edge first checks the interval endpoints, positive denominators, and nonnegative incoming and outgoing tail slopes.

For each $1 \leq K \leq 24$, put $m = \lfloor 24/K \rfloor$. Since $K(m+1) \geq 25$, the incoming gain is computed exactly by (30) with $D = m$. The checker evaluates the edge at the single rational threshold in (33). Lemma 10.1 proves that this one inequality covers the whole real interval, including when the incoming row has a negative gain at a small state. Clearing denominators is justified by the previously checked positivity conditions.

For $K \geq 25$, the incoming gain equals $aK \geq 0$ and $K - 1 \geq p - 2$, since $p \leq 23$. The minimizing threshold is therefore ℓ_p . The outgoing affine slope and intercept are checked separately to be no greater than the two coefficients in (31) at $t = \ell_p$. These coefficient inequalities imply the required bound for every $K \geq 25$; no assertion that a pointwise minimum commutes with taking affine coefficients is needed.

There are 367366 checked cells in 1441 numerical modules. Each nonempty cell includes an edge for every possible first prime and a certified child with exactly that prime's interval erased. Lean's declaration dependencies provide the well-founded subset recursion. The geometric partition has 15951 nodes and 7976 leaves, with its coverage proof in 125 modules. At an internal node the literal child boxes are checked to cover the parent, including the splitting boundary; the root box is checked to be exactly $\mathcal{B}$. Each leaf carries the root-value inequality $1001/1000 \leq R_{P_8, B}(1)$.

All these arithmetic and finite-structure obligations are discharged by kernel reduction, using `decide +kernel`. The symbolic soundness proofs transport the executable checks to inequalities of real functions. Neither a

floating-point solver nor an external program's success flag is a premise of the resulting theorem. $\square$

Proof of the eight-prime assertion in Theorem 1.3. Every threshold vector belongs to a leaf of the partition in Lemma 10.3. Lemma 10.2 bounds its completed value below by $1001/1000$, in every ordering. The exact affine-tail identity (30) identifies this value with (27). $\square$

The rational number $1001/1000$ is a certified lower bound, not an asserted exact minimum. In particular, further numerical optimization or a more elaborate upper-envelope table cannot produce a subunit bound within this same exact scalar model.

11. PERSISTENCE AND SCOPE OF THE CEILING

11.1. Adding blocks or terminal continuation.

Proposition 11.1 (No recovery by insertion). *Insert any finite number of blocks with integer parameters $p \geq 3$ and thresholds $0 \leq t \leq p - 3$ into an ordering of the eight reference blocks. The resulting zero-terminal scalar value is at least $1001/1000$. For a finite positive completion of sufficient height, the same bound holds with any nonnegative terminal continuation. It also holds for the actual infinite-depth operators with a nonnegative eventually affine terminal continuation, or more generally when every successive depth series is absolutely convergent.*

Proof. Choose a common completion height sufficient for the zero-terminal compositions under consideration. Replacing a nonnegative terminal continuation by zero cannot increase the value, by pointwise order preservation. Every zero-terminal suffix is increasing. Consequently, an inserted block applied to such a suffix is at least that suffix, by Lemma 9.1. Pointwise order preservation by the preceding blocks propagates this inequality to the root. Delete the extra blocks successively and apply the eight-prime bound to the retained subsequence.

For eventually affine continuations, (30) justifies every infinite series internally and identifies the finite and infinite calculations. For general continuations with absolutely convergent successive depth series, the infinite version of the positive mixture follows by summation by parts and passage to the limit; hence the same pointwise comparison applies. $\square$

This proves the insertion assertion of Theorem 1.3. It also shows that arbitrary pointwise upper-certificate rows cannot evade the ceiling: backward order preservation forces each such row to dominate the exact zero-terminal composition. In the general infinite-continuation statement, convergence is required at every state of every actual suffix, not merely for the initial terminal function. No totalized value assigned to a divergent series is used.

Thus the largest *uniform reference-support bound* certified by this scalar architecture is eight: seven reference blocks have a subunit witness, whereas

the first eight and every finite extension containing them do not. This is not a claim about every individual eight-prime support with larger primes, nor a proof that rank nine is false; in particular, it does not cover larger reference primes obtained by a case split on the small actual primes. The result does not exclude improvements retaining residue geometry, multiple state statistics, joint-block interactions, or information specific to exponent patterns.

The saturation $L \leq K - 1$ and reduction to the single index count in Proposition 7.2 are natural candidates for improvement, because they discard load deficits and information linking the different depth sets. This is a structural diagnosis, not a sharpness theorem: the ceiling certificate does not separate those losses from the one-fibre estimates in Lemma 5.1, and does not prove that either step alone is the binding obstruction at rank nine.

11.2. Two natural enlargements.

Proposition 11.2 (Wider thresholds). *The eight-prime ceiling still holds when the threshold at p may range over $0 \leq t \leq p - 3 + 1/p$.*

Proof. The positive-mixture proof remains valid because now $c \geq (p-1)/p = \beta_p(1)$. For an increasing continuation, clip t to $t' = \min\{t, p-3\}$. At an integer load L , either $L \leq p-3$ or $L \geq p-2$. In the first case the low-load monotonicity proved in Lemma 10.1 applies; in the second case the high-load difference has sign $L + G_p h(K) - (p-2) \geq 0$. Thus $T_{p,t'} h \leq T_{p,t} h$. Clip successive thresholds, using pointwise order preservation to propagate the comparison, and apply the original ceiling. Zero-terminal continuations have exact affine tails, so all sums converge. $\square$

Proposition 11.3 (Convex terminal penalties). *Let the terminal prime be $p \in P_8$, put $q = p - 2$, and replace its hinge by $h(K) = \psi(K - 1)$, where $\psi : [0, \infty) \rightarrow [0, \infty)$ is convex, $\psi(0) = 0$, and $\psi(q) \geq 1$. With the other seven blocks in any order and with admissible thresholds, the ceiling holds for every sufficient positive completion. It holds for the actual infinite-depth chain whenever its successive depth series are absolutely convergent.*

Proof. Normalize $\varphi = \psi/\psi(q)$ and let $s = 1 - \varphi(q - 1)$. Convexity and nonnegativity give $1/q \leq s \leq 1$. The line through $(q - 1, \varphi(q - 1))$ and $(q, 1)$ lies below $\varphi(n)$ at every nonnegative integer n : convex secant slopes give this for $n \leq q - 1$ and $n \geq q$. There is no other integer between these two endpoints. Set $t = q - 1/s$, so $0 \leq t \leq q - 1$. Since the line is $1 + s(n - q) = (n - t)/(q - t)$, nonnegativity yields

$$(35) \quad \frac{(n - t)_+}{q - t} \leq \varphi(n) \leq \psi(n) \quad (n \in \mathbb{N}).$$

Replace the terminal penalty by this admissible hinge and propagate the pointwise inequality through the earlier blocks. The terminal prime occurs only once, so this choice does not change any earlier threshold. Apply the

eight-prime ceiling. The passage to actual infinite operators uses the stated convergence hypothesis. $\square$

The last-unit secant is important: an arbitrary convex function need not majorize that secant line between $q - 1$ and q , but only integer loads occur here. Positive mixtures of normalized convex penalties are included in Proposition 11.3.

11.3. A load-shaped interpolation. A further verified comparison concerns a particular load-shaped scalar expression. For $q > u \geq 0$, $0 \leq \alpha \leq 1$, $L \geq 0$, and arbitrary real f, g , set $c = q - \alpha u$, $m = q - u$, and

$$(36) \quad H_\alpha = \frac{L - \alpha \min(L, u)}{q - \alpha \min(L, u)} + f + \frac{g}{c} - \left(\frac{1}{c} - \frac{1}{q - \alpha \min(L, u)} \right) \min(g, m).$$

Let $H(t)$ denote the fixed-threshold expression of Lemma 10.1 at the same q, L, f, g .

Proposition 11.4 (Load-shaped dominance). *Under these assumptions, $H_\alpha \geq H(\alpha u)$. Consequently, the chain obtained from (36) with $q = p - 2$, $0 \leq u \leq p - 3$, and $0 \leq \alpha \leq 1$ at each reference prime has value at least $1001/1000$.*

Proof. The fixed expression has the equivalent subtractive form

$$H(t) = \frac{(L - t)_+}{q - t} + f + \frac{g}{q - t} - \left(\frac{1}{q - t} - \frac{1}{q - \min(t, L)} \right) \min(g, q - t).$$

If $L \geq u$, the expressions agree. Suppose $L \leq u$ and put $A = 1/c - 1/(q - \alpha L) \geq 0$. For $\alpha u \leq L \leq u$ the difference is $A(q - L - \min(g, m)) \geq 0$. For $L \leq \alpha u$, put $B = 1/c - 1/(q - L)$ and $Q = (1 - \alpha)L/(q - \alpha L)$. Then $0 \leq B \leq A$, $m \leq c$, and

$$Q - (A - B)m = \frac{(1 - \alpha)L(u - L)}{(q - \alpha L)(q - L)} \geq 0.$$

Thus the difference in this case satisfies

$$Q - A \min(g, m) + B \min(g, c) \geq Q - (A - B)m \geq 0.$$

No sign assumption on g was used. Apply the pointwise inequality backwards through the chain, using order preservation of the fixed-threshold operators. These zero-terminal shaped continuations also have eventual affine tails, which justify their infinite gains. The fixed chain has thresholds αu in the original range. $\square$

This last proposition compares two precisely specified scalar expressions. It does not assert that the shaped expression is an admissible full-label lift, or rule out arbitrary state-dependent methods.

12. FORMALIZATION AND REPRODUCIBILITY

There are two separate Lean 4 [8] projects, both using mathlib [16] and the pinned toolchain `v4.32.0-rc1`. The versioned companion supplies the exact dependency revisions, paper-to-Lean theorem maps, source manifests and reproduction instructions.

12.1. The integer-covering theorem. The paper-facing theorem for the first project is

`Rank8Paper.eight_prime_support_of_integer_cover.`

Its hypotheses are coverage of every integer by a finite family of integer residue classes, natural moduli greater than one, oddness, and injectivity of the modulus map. Its conclusion is $8 \leq \omega(\text{lcm}_i m_i)$, using Lean’s finite-set `lcm` and `primeFactors.card`. Coverage is expressed with `Int.ModEq`. There is no residual-model, exponent-cutoff, numerical-oracle or certificate hypothesis. The formal proof includes the full-label argument and the arithmetic bridge. Zero-height padding handles all ranks at most seven; the audit excludes the earlier prime-support theorems from the final proof’s transitive dependencies, although they remain in the shared library.

12.2. The scalar ceiling. The second project’s principal closed theorem is

`RankScaling.eight_prime_scalar_ceiling.`

It bounds the actual infinite-depth scalar value below by 1001/1000 for every permutation of the first eight odd reference primes and every real threshold function satisfying $0 \leq \theta(p) \leq p - 3$. Convergence of the eventually affine depth series, the exact finite completion, the lower certificate and coverage of the continuous parameter box are proved internally. This is a ceiling for the specified scalar architecture, not an obstruction to a rank-nine theorem by other methods. The companion records the additional hypotheses of each extension in Section 11; in particular, general infinite continuations require explicit convergence.

12.3. Trust boundary and availability. Both projects’ public theorem audits report exactly the standard axioms `propext`, `Classical.choice` and `Quot.sound`, with no admitted goals, custom mathematical axioms or native-computation proof shortcuts. Certificate arithmetic is checked by Lean’s kernel; the generators and supplementary Python regressions are not proof oracles. Finite tests do not establish universal quantifiers. The integer-covering project additionally has a fresh imported-proof replay using `leanchecker-fresh`: this runs Lean’s own kernel, not an independently implemented checker. That separate replay claim does not apply to the scalar project.

The compact companion contains both projects, exact inputs and deterministic generators. It regenerates 1,568 scalar Lean files and verifies their preserved source hashes. Its report separates edition checks from historical scalar replay. Expanded sources and build caches are omitted. Checksums identify bytes, not mathematical correctness.

APPENDIX A. STANDALONE RATIONAL VERIFICATION

The following Python 3 program uses only exact fractions and reads no external data. Run it without optimization flags, since its assertions are verification conditions. The affine identities and (24) explain why its finite computation is sufficient. The companion additionally tests rejection of each of the 54 single-entry downward mutations.

```

from fractions import Fraction as Q
if not __debug__:
    raise RuntimeError("Assertions must be enabled")
P = (3, 5, 7, 11, 13, 17)
T = (0, 0, 1, 3, 5, 7)
V = (
    (998776, 53684, 2444, 101, 3, 1),
    (4711146, 810347, 66470, 4978, 362, 1),
    (9003873, 2047552, 542310, 41666, 5144, 69),
    (13502517, 3464004, 1155363, 157161, 20437, 1298),
    (18113840, 5064174, 1975441, 492712, 71069, 3343),
    (22797470, 6759597, 2889210, 910486, 188307, 27942),
    (27507176, 8533245, 3909651, 1464023, 465550, 58007),
    (32217918, 10310004, 4934332, 2022956, 749153, 95589),
    (36944897, 12135471, 6025435, 2666427, 1132390, 250920))
A = tuple(Q(*x) for x in
    ((607, 128), (479, 256), (1181, 1024), (185, 256),
     (61, 128), (17, 64), (1, 8)))
B = tuple(Q(*x) for x in
    ((-23, 4), (-19, 4), (-53, 12), (-47, 12),
     (-13, 4), (-9, 4), (-5, 4)))
small = {k: Q(0) for k in range(1, 10)}
a, b = A[6], B[6]
count = terms = 0
for j in reversed(range(6)):
    p, t = P[j], T[j]
    c = p-2-t
    def u(k):
        return small[k] if k < 10 else a*k+b
    new = {}
    for k in range(1, 10):
        m = 9//k
        assert k*(m+1) >= 10
        g = sum((Q(p-1, p**d)*(u(k*(d+1))-u(k*d))
                 for d in range(1, m+1)), Q(0))
        g += a*k/p**m
        r = Q(1, p-2-min(t, k-1))
        v = Q(max(k-1-t, 0), c) + max(
            u(k)+r*g, u(k)+g/c-(1-c*r))
        new[k] = Q(V[k-1][j], 10**6)
        assert v <= new[k]
        count += 1
        terms += m
    a, b = a+(1+a)/c, b-Q(t+1, c)
    assert (a, b) == (A[j], B[j]) and a >= 0
    small = new
assert count == 54 and terms == 138
assert small[1] == Q(124847, 125000)
assert 1-small[1] == Q(153, 125000) > Q(1, 1000)
print("PASS: 54 inequalities, six tails, exact root margin")

```

APPENDIX B. EXACT SEVEN-PRIME SCALAR WITNESS

For the schedule in Proposition 9.2, the exact root value is $V(\mathbf{p}, \mathbf{t}) = A/B$, with the integers in Table 3. They satisfy $10^6 A < 998419B$ by integer arithmetic. This is the unrounded witness from zero terminal data, not the 54-entry upper certificate used in Theorem 1.1.

Symbol	Exact integer
A	881004897188574708854993369659077464836643
B	882400271037897294391745131024934604600000

TABLE 3. Numerator and denominator of the exact seven-prime scalar value.

ACKNOWLEDGMENTS

This work benefited from research assistance by AI systems developed by OpenAI and Anthropic, including support for proof exploration, proof development, and exact computational checks among others.

AVAILABILITY AND RESPONSIBILITY

The paper and companion are available at doi:10.5281/zenodo.22747841. The compact source package supplies both Lean projects, their theorem maps, reproduction instructions, exact checkers and verification records. All certificate data are included; neither an external solver nor a search transcript is needed. AI output is not a proof oracle. The mathematical argument is given above, and Section 12 sets out the machine-checking boundary. The author is responsible for the statements and their presentation.

REFERENCES

- [1] S. Adenwalla, *A question of Erdős and Graham on covering systems*, preprint, 2025, arXiv:2501.15170v3.
- [2] P. Balister, *Erdős covering systems*, in *Surveys in Combinatorics 2024*, London Mathematical Society Lecture Note Series **493**, Cambridge University Press, 2024, 31–54. doi:10.1017/9781009490559.003.
- [3] P. Balister, B. Bollobás, R. Morris, J. Sahasrabudhe and M. Tiba, *On the Erdős covering problem: the density of the uncovered set*, *Invent. Math.* **228** (2022), 377–414. arXiv:1811.03547.
- [4] P. Balister, B. Bollobás, R. Morris, J. Sahasrabudhe and M. Tiba, *The Erdős–Selfridge problem with square-free moduli*, *Algebra Number Theory* **15** (2021), 609–626. doi:10.2140/ant.2021.15.609.
- [5] P. Balister, B. Bollobás, R. Morris, J. Sahasrabudhe and M. Tiba, *The structure and number of Erdős covering systems*, *J. Eur. Math. Soc.* **26** (2024), 75–109. doi:10.4171/JEMS/1357.
- [6] C. Bispels, M. Cohen, J. Harrington, J. Lowrance, K. Pontes, L. Schaumann and T. W. H. Wong, *A further investigation on covering systems with odd moduli*, *Discrete Math.* **349** (2026), article 115013. doi:10.1016/j.disc.2026.115013.

- [7] B. Côté and R. Wang, *On convex order and supermodular order without finite mean*, preprint, 2026, arXiv:2502.17803v3.
- [8] L. de Moura and S. Ullrich, *The Lean 4 theorem prover and programming language*, in Automated Deduction—CADE 28, Lecture Notes in Computer Science **12699**, Springer, 2021, 625–635. doi:10.1007/978-3-030-79876-5_37.
- [9] P. Erdős, *On integers of the form $2^k + p$ and some related problems*, Summa Brasil. Math. **2** (1950), 113–123. Author’s collected papers.
- [10] M. Filaseta and A. Kalogirou, *Covering systems with the sum of the reciprocals of the moduli close to 1*, preprint, 2024, arXiv:2407.15280.
- [11] S. Guo and Z.-W. Sun, *On odd covering systems with distinct moduli*, Adv. Appl. Math. **35** (2005), 182–187. doi:10.1016/j.aam.2005.01.004.
- [12] J. Harrington, J. Klein, J. Lowrance and O. Trifonov, *Covering systems where the prime divisors of all moduli are only 2, 3, or 5*, preprint, 2026, arXiv:2605.18644.
- [13] J. Harrington, Y. Sun and T. W. H. Wong, *Covering systems with odd moduli*, Discrete Math. **345** (2022), article 112936. doi:10.1016/j.disc.2022.112936.
- [14] B. Hough, *Solution of the minimum modulus problem for covering systems*, Ann. of Math. (2) **181** (2015), 361–382. doi:10.4007/annals.2015.181.1.6.
- [15] R. D. Hough and P. P. Nielsen, *Covering systems with restricted divisibility*, Duke Math. J. **168** (2019), 3261–3295. doi:10.1215/00127094-2019-0058.
- [16] The mathlib Community, *The Lean mathematical library*, in Proceedings of the 9th ACM SIGPLAN International Conference on Certified Programs and Proofs (CPP 2020), ACM, 2020, 367–381. doi:10.1145/3372885.3373824.
- [17] I. Mian and S. Siddique, *Kernel-checked exclusions for the Erdős–Selfridge odd covering problem: any odd covering of $\mathbb{Z}$ has lcm exceeding 10000*, preprint, 2026, arXiv:2607.25628.
- [18] G. Puccetti and R. Wang, *Extremal dependence concepts*, Statist. Sci. **30** (2015), 485–517. doi:10.1214/15-STS525.
- [19] M. Schroeder, *Seven prime divisors in odd distinct covering systems*, preprint, 1 September 2026, Zenodo, version 1.0.1. doi:10.5281/zenodo.22739035.
- [20] P. Erdős, *Some recent advances and current problems in number theory*, in Lectures on Modern Mathematics, Vol. III (T. L. Saaty, ed.), Wiley, New York, 1965, 196–244; see p. 235. Author’s collected papers.
- [21] T. F. Bloom, *Erdős Problem #7*, erdosproblems.com/7 (accessed 6 September 2026).

INDEPENDENT RESEARCHER