

NINE PRIME DIVISORS IN ODD DISTINCT COVERING SYSTEMS

MICHAEL SCHROEDER

ABSTRACT. We prove that every finite covering system with pairwise distinct odd moduli greater than one has at least nine distinct prime divisors in the least common multiple of its moduli. There is no restriction on the prime-power exponents. The proof combines completion of the congruence family with conditional reweighting that retains the joint geometry of the two smallest prime coordinates. Two deletion budgets reduce to a finite set of vertices by an explicit convex interpolation, and a prescribed mixed-modulus projection closes the remaining cases. Exact geometric sums and first moments control every omitted exponent range. A certificate of 28,001 integer inequalities completes the argument. We also prove that the least common multiple of an odd distinct cover is at least 9,704,539,845, and that every finite distinct odd family with moduli greater than one supported on at most eight primes leaves natural density at least $1/1,002,375$ uncovered. A Lean formalization proves the rank theorem and both quantitative corollaries.

1. INTRODUCTION

A *covering system* is a finite family

$$\mathcal{C} = \{a_m \pmod{m} : m \in \mathcal{M}\}$$

whose union is $\mathbb{Z}$. It is *distinct* if its moduli are pairwise different, and *odd* if every modulus is odd. Throughout, every modulus is greater than one. Write

$$N = \text{lcm}_{m \in \mathcal{M}} m, \quad \omega(N) = |\{p : p \text{ prime, } p \mid N\}|.$$

We call $\omega(N)$ the *rank* of the family.

Theorem 1.1 (Nine-prime support theorem). *Every finite odd distinct covering system with moduli greater than one satisfies*

$$\omega(N) \geq 9.$$

Date: 13 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, distinct odd moduli, prime divisors, conditional distortion, convex comparison, computer-assisted proof.

Corollary 1.2. *The least common multiple of every such covering system satisfies*

$$N \geq 3 \cdot 5 \cdot 7 \cdot 11 \cdot 13 \cdot 17 \cdot 19 \cdot 23 \cdot 29 = 3,234,846,615.$$

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

Corollary 1.3 (Strengthened LCM bound). *The least common multiple of every such covering system satisfies*

$$N \geq 9,704,539,845.$$

Proof. Put $P = 3 \cdot 5 \cdot 7 \cdot 11 \cdot 13 \cdot 17 \cdot 19 \cdot 23 \cdot 29$. If $\omega(N) \geq 10$, then $N \geq 31P > 3P$. By Theorem 1.1, it remains to consider $\omega(N) = 9$.

Suppose first that N is square-free, with prime divisors $q_1 < \dots < q_9$. In each coordinate $\mathbb{Z}/q_i\mathbb{Z}$, delete the residue of the pure q_i -class if present, and otherwise delete any one residue. Give the remaining $q_i - 1$ residues uniform measure and take the product. All pure classes now have measure zero. A class with support $S \subseteq \{1, \dots, 9\}$ of size at least two has measure at most $\prod_{i \in S} (q_i - 1)^{-1}$. Distinct square-free moduli have distinct supports. Since $q_i \geq p_i$, where $(p_1, \dots, p_9) = (3, 5, 7, 11, 13, 17, 19, 23, 29)$, the union bound gives

$$\begin{aligned} \mathbb{P}(\text{covered}) &\leq \sum_{|S| \geq 2} \prod_{i \in S} \frac{1}{p_i - 1} \\ &= \prod_{i=1}^9 \left(1 + \frac{1}{p_i - 1}\right) - 1 - \sum_{i=1}^9 \frac{1}{p_i - 1} = \frac{295112669}{340623360} < 1, \end{aligned}$$

contrary to coverage.

Thus N is not square-free. Its radical satisfies $\text{rad}(N) \geq P$, and $N/\text{rad}(N)$ is an odd integer greater than one. Hence $N \geq 3P = 9,704,539,845$. $\square$

The exponents and the residues are unrestricted. There is no square-free, minimality, disjointness, or complete-divisor-list hypothesis. The theorem excludes rank at most eight; it neither constructs a rank-nine covering system nor resolves the odd covering problem. The measure construction also gives a uniform lower bound for the actual uncovered density; see Corollary C.2.

1.1. Context. Distinct covering systems were introduced by Erdős in 1950 [7]. The odd covering problem of Erdős and Selfridge asks whether such a system can have all its moduli odd; see [8, 6]. Balister's survey [1] gives a broad account of the classical questions and modern methods.

Berger, Felzenbaum and Fraenkel obtained the unrestricted necessary condition [4, 5]

$$\prod_{p|N} \frac{p-1}{p-2} - \sum_{p|N} \frac{1}{p-2} > 2;$$

see also the explicit account in [9]. Writing $x_p = (p - 2)^{-1}$, the left side is $\prod_p(1 + x_p) - \sum_p x_p$, which is nondecreasing in each variable and under the addition of a positive variable. For at most four odd primes it is at most its value at 3, 5, 7, 11, namely $86/45 < 2$. Thus this condition already implies $\omega(N) \geq 5$.

The probabilistic study of covering systems includes Hough's solution of the minimum-modulus problem [10] and the theorem of Hough and Nielsen that every distinct covering has a modulus divisible by 2 or 3 [11]. In an odd cover this forces $3 \mid N$, without requiring 3 itself to occur as a modulus. The distortion method of Balister, Bollobás, Morris, Sahasrabudhe and Tiba gives lower bounds for uncovered density [2]. Their Theorem 1.4 also gives $2 \mid N$, or $9 \mid N$, or $15 \mid N$; the last alternative does not assert that one modulus is divisible by 15. The present normalized conditional reweighting belongs to this line of methods.

The square-free case has a different status. Simpson and Zeilberger [18] and Guo and Sun [9] obtained restrictions on the prime support of square-free odd coverings. Balister et al. subsequently excluded that case altogether [3]. Here arbitrary powers of each prime remain possible at every stage.

The author's preceding papers establish the unrestricted bounds seven and eight [16, 17]. The latter also identifies a ceiling for a specified scalar comparison architecture. That ceiling does not apply to the present argument: we retain the spatial anchor geometry, two deletion budgets, and a prescribed mixed-modulus projection. The proof of Theorem 1.1 is self-contained apart from the explicitly specified finite integer and rational computations. Neither preceding rank theorem is assumed in that written proof; the accompanying Lean formalization reuses the proved rank-eight theorem for smaller prime supports, as explained in Section 11.

1.2. Architecture and the computational assertion. Prime replacement reduces the problem to the fixed support

$$(1.1) \quad \mathcal{P} = \{3, 5, 7, 11, 13, 17, 19, 23\}.$$

A completion construction then supplies every pure prime power and the mixed moduli

$$(1.2) \quad 15, 21, 35, 45, 63, 75, 105, 165.$$

The completed family covers every point covered by the original family. It has at most one class per modulus. We may therefore use all these selected classes in a single argument.

The full 3- and 5-coordinates are exposed first. Their retained structure is organized on 135 coarse cells, together with two continuous budgets for the deeper pure-3 deletions and the overlapping pure-5/75 deletions. The later primes use the single threshold sequence

$$(1.3) \quad (\tau_7, \tau_{11}, \tau_{13}, \tau_{17}, \tau_{19}, \tau_{23}) = (2, 4, 4, 8, 8, 12).$$

At prime 7 we enlarge each forbidden fibre by a fully charged amount. The resulting bound depends only on the sum of four projection indicators, without distinguishing which target digits coincide. At prime 11 we retain only the projection of the 165-class, holding it fixed before taking the auxiliary expectation.

The resulting upper loss is a convex functional of the comparison weights. An explicit vertex decomposition reduces the continuous parameters to a finite domain. Proposition 10.1 asserts the required finite inequalities. Its complete certificate has 28,001 integer inequalities, with minimum surplus four in units of mass $1/135000$. The four terminal states share the simpler inequality

$$\frac{5310}{1000} - \frac{5299}{1000} > 0$$

in 135-cell units. The geometry algorithm and all infinite remainders are specified below and implemented in the accompanying standalone checker. The Lean development proves the rank theorem, both LCM bounds, and the uncovered-density corollary. The density statement concerns arbitrary families on at most eight primes and has no covering hypothesis; Section 11 gives the exact formal scope.

2. PRIME REPLACEMENT AND COMPLETION

A class modulo $\prod_p p^{e_p}$ is a product of prefix cylinders in the prime-adic coordinates: the coordinate at p fixes its first e_p digits, read from least significant to most significant. Distinct moduli have distinct complete exponent vectors. In particular, there is at most one class for any given complete vector.

Lemma 2.1 (Prime replacement). *Let $p_1 < \dots < p_r$ and $q_1 < \dots < q_r$ be primes with $p_i \leq q_i$. A finite distinct covering family supported on the q_i pulls back to a finite distinct covering family supported on the p_i .*

Proof. Let A_i be the largest exponent of q_i in the original period. Inject the digit strings of $\mathbb{Z}/p_i^{A_i}\mathbb{Z}$ into those of $\mathbb{Z}/q_i^{A_i}\mathbb{Z}$ by keeping each digit unchanged. The inverse image of a depth- e cylinder is either empty or a depth- e cylinder. Apply the product injection to the finite CRT space. Discard the empty inverse images. The remaining cylinders cover, and their complete exponent vectors are still distinct and nonzero. $\square$

This is the prime-replacement principle of [18, Theorem 1 and Corollary 1, pp. 59–61]; those results already allow arbitrary prime powers. If a cover has rank at most eight, take the p_i to be the first r odd primes. It therefore suffices to exclude covers supported on (1.1), allowing some of those primes to be absent.

Let

$$X = \prod_{p \in \mathcal{P}} \mathbb{Z}_p$$

with product Haar probability H . A finite covering of $\mathbb{Z}$ supported on $\mathcal{P}$ covers X : its classes depend on a finite common modulus, and every residue of that modulus occurs among the integers. We make this passage before introducing any infinite family.

Lemma 2.2 (Countable completion). *Let D be a finite or countable set of moduli greater than one, all supported on $\mathcal{P}$, such that*

$$(2.1) \quad \sum_{\substack{d \in D \\ d|m, d < m}} \frac{1}{d} < 1 \quad (m \in D).$$

A finite cylinder family on X , with all moduli supported on $\mathcal{P}$ and at most one class per modulus, can be enlarged, preserving its covered subset of X , to a family containing one class of every modulus in D . The selected classes at $d, m \in D$ are disjoint whenever $d \mid m$ and $d < m$. Only selected moduli are inserted or have their residues changed.

Proof. Process the elements of D in increasing order. At modulus m , the already fixed proper-divisor classes occupy at most

$$m \sum_{\substack{d \in D \\ d|m, d < m}} \frac{1}{d} < m$$

residues modulo m . Hence a free residue exists. Keep an existing m -class if it is disjoint from these classes; insert a class if it is absent; otherwise move it to a free residue. If an m -class meets a d -class with $d \mid m$, it is contained in that smaller class. Moving it therefore loses no covered point.

Every smaller class used in this argument is already fixed permanently. When a later selected multiple of m is processed, the fixed m -class is among its proper-divisor classes, so the resulting class is disjoint from the m -class. At the countable limit, each original class either remains or was contained in a permanently fixed smaller class when changed. Thus the final family still covers the original covered set. This last observation proves preservation at the limit, rather than merely at each finite stage. $\square$

Apply the lemma with

$$(2.2) \quad D = \{p^e : p \in \mathcal{P}, e \geq 1\} \cup \{15, 21, 35, 45, 63, 75, 105, 165\}.$$

For p^e the sum in (2.1) is less than $1/(p-1) \leq 1/2$. For the mixed moduli the sums are

$$(2.3) \quad \frac{m}{\sum_{d|m, d < m, d \in D} d^{-1}} \left| \begin{array}{cccccccc} 15 & 21 & 35 & 45 & 63 & 75 & 105 & 165 \\ \hline \frac{8}{15} & \frac{10}{21} & \frac{12}{35} & \frac{32}{45} & \frac{40}{63} & \frac{16}{25} & \frac{86}{105} & \frac{38}{55} \end{array} \right.$$

They are all below one. Write P_p for the union of the completed pure p -power classes. They are disjoint, so

$$(2.4) \quad H_p(P_p) = \sum_{e \geq 1} p^{-e} = \frac{1}{p-1}.$$

All mixed classes remain finite in number. Only the auxiliary pure families have become countable.

Remark 2.3. Completion preserves coverage, not the individual original residues. The proof uses the completed family and constructs a positive measure outside it. Such a measure is also outside the original covered set. There is no inference that an arbitrary infinite covering of the integers covers X .

3. CONDITIONAL KERNELS AND ABSOLUTE LOSSES

Expose the full coordinates in the order 3, 5, 7, 11, 13, 17, 19, 23. A later conditional kernel may depend on the complete fixed family and all earlier actual coordinates. The auxiliary comparison variables introduced below are not information available to that kernel.

Lemma 3.1 (Capped current-loss kernel). *On a current coordinate let P be a pure forbidden set and B a mixed bad set. Suppose $CH(P^c) \geq 1$. There is a normalized measurable density k , with $0 \leq k \leq C$ and $k = 0$ on P , whose loss on B is*

$$(3.1) \quad \int_B k dH = [1 - CH((P \cup B)^c)]_+.$$

Proof. Put $G = (P \cup B)^c$. If $H(G) \geq 1/C$, take density $1/H(G)$ on G . Otherwise take density C on G and constant density

$$\frac{1 - CH(G)}{H(B \setminus P)}$$

on $B \setminus P$, with zero density on P . The denominator is positive in this case, and feasibility gives $1 - CH(G) \leq CH(B \setminus P)$. The density is normalized, is at most C , and loses precisely the stated mass. The formulas are measurable in the earlier history. $\square$

For a current prime $q > 5$, call a mixed class *ending at q* if its largest prime divisor is q . Retain distinct labels for its current exponent and all earlier exponents. For a class $q^e n$ with $n > 1$, let $I_{e,n}$ indicate its earlier-coordinate cylinder and set

$$(3.2) \quad Z_q = (q - 1) \sum_{\substack{e \geq 1 \\ n > 1}} q^{-e} I_{e,n}, \quad \beta_q = q - 1 - \tau_q, \quad C_q = \frac{q - 1}{\beta_q}.$$

The current pure and mixed union has Haar mass at most $(1 + Z_q)/(q - 1)$. Lemma 3.1 therefore gives

$$(3.3) \quad \ell_q \leq \frac{[1 + Z_q - \tau_q]_+}{\beta_q}.$$

The right side is an unclipped convex majorant. Although an actual conditional loss is at most one, clipping this majorant at one would invalidate the convex comparison.

The values used throughout the proof are

$$(3.4) \quad \begin{array}{c|cccccc} q & 7 & 11 & 13 & 17 & 19 & 23 \\ \hline \tau_q & 2 & 4 & 4 & 8 & 8 & 12 \\ \beta_q & 4 & 6 & 8 & 8 & 10 & 10 \\ C_q & \frac{3}{2} & \frac{5}{3} & \frac{3}{2} & 2 & \frac{9}{5} & \frac{11}{5} \end{array}$$

They satisfy $1 < C_q < q$ and $C_q(1 - 1/(q - 1)) \geq 1$.

Let A be the anchor set avoiding every completed class supported on $\{3, 5\}$. Initialize with H restricted to A , then use normalized kernels and delete each current bad union. Denote the resulting live submeasure after stage i by ν_i .

Lemma 3.2 (Domination and first-hit accounting). *If L_i is the absolute mass deleted at stage i , then*

$$(3.5) \quad \nu_k(X) = H(A) - \sum_{i=1}^k L_i.$$

For a cylinder with anchor part E and constrained later depths e_p , the live mass is at most

$$(3.6) \quad H(A \cap E) \prod_{p: e_p > 0} \frac{C_p}{p^{e_p}}.$$

Both conclusions hold with the kernels extended to histories already deleted. Enlarging the initial measure can only increase its absolute loss under these fixed kernels and deletions.

Proof. A normalized kernel preserves the incoming total mass; deletion subtracts exactly the current loss. This proves (3.5). The live measure is dominated by the law with the same kernels before deletions. Integrate that law in reverse coordinate order. At an unconstrained coordinate its kernel integrates to one; at a constrained coordinate the integral is bounded by C_p/p^{e_p} . Continue to the anchor coordinates to obtain (3.6). Finally, transitions and deletions are positive linear operators on the initial measure. $\square$

Survivors are never renormalized. Conditional bounds on entire prime-power coordinates, rather than marginal first-digit bounds, are what license the reverse integration. We call a lower bound for the initial anchor mass a *reserve*; its difference from the sum of the upper loss bounds is the *ledger*. A positive ledger guarantees a surviving point.

4. CONVEX COMPARISON AND THE ALL-HEIGHT INVENTORY

Lemma 4.1 (Ordered increments). *Let f be increasing and convex, let $b, w_1, \dots, w_n \geq 0$, and let events $E_1, \dots, E_n$ satisfy $\mathbb{P}(E_i) \leq \lambda_i$, where $1 \geq \lambda_1 \geq \dots \geq \lambda_n \geq 0$. For U uniformly distributed on $[0, 1]$,*

$$(4.1) \quad \mathbb{E}f\left(b + \sum_i w_i \mathbf{1}_{E_i}\right) \leq \mathbb{E}f\left(b + \sum_i w_i \mathbf{1}_{\{U \leq \lambda_i\}}\right).$$

Proof. Let d_i be the increment of f when w_i is added after all preceding weights. For a subset of the indices, add its weights in their original order. Each increment starts from a base no larger than the one defining d_i , so convexity bounds it by d_i . Hence pointwise

$$f\left(b + \sum_i w_i \mathbf{1}_{E_i}\right) \leq f(b) + \sum_i d_i \mathbf{1}_{E_i}.$$

Take expectations and use $d_i \geq 0$. The resulting expression $f(b) + \sum_i \lambda_i d_i$ is the expectation on the right of (4.1), because its events are nested. $\square$

This is a finite form of the classical comonotone extremal principle for convex functions of nonnegative weighted sums; see [15]. The countable inventories below are treated by nonnegative truncation and the explicit first-moment bounds of Section 8.

Apply the lemma conditionally at the last exposed coordinate and then iterate backwards. A fixed value of the auxiliary variable leaves an increasing convex function of a nonnegative sum of earlier cylinders, so the iteration is legitimate. A cap C_p gives the independent auxiliary comparison run

$$(4.2) \quad \mathbb{P}(J_p = 0) = 1 - \frac{C_p}{p}, \quad \mathbb{P}(J_p = j) = \frac{C_p(p-1)}{p^{j+1}} \quad (j \geq 1).$$

These runs do not assert independence of the actual distorted coordinates. Repeated labels at the same depth remain separate in the application of the lemma and share only the corresponding upper probability bound.

Lemma 4.2 (Subprobability comparison). *A measure of total mass $A \geq 1/p$ with each depth- e cylinder of mass at most p^{-e} is dominated, for nonnegative increasing convex costs of the indicated cylinders, by the nested comparison masses*

$$(4.3) \quad \pi(0) = A - \frac{1}{p}, \quad \pi(j) = \frac{p-1}{p^{j+1}} \quad (j \geq 1).$$

More generally, a submeasure of mass at most S with depth- e caps C/p^e admits a nested upper comparison of total mass S and tails $\min(S, C/p^e)$. If $S \geq C/p$, only the zero atom changes from (4.2): it becomes $S - C/p$.

Proof. For the first statement normalize by A , use the bounds p^{-e}/A in Lemma 4.1, and multiply back. For the second, the constant term costs at most $Sf(b)$ and each nonnegative increment costs at most both the total mass and its cylinder cap. The same chain-increment proof gives the stated nested measure. $\square$

In particular, deleting mass does not multiply every positive-depth atom by the surviving mass. A deep surviving cylinder may still attain its full Haar cap.

4.1. Aggregation of positive types. For a fixed current prime, distinctness gives at most one ending class for each complete exponent vector. Enlarge its nonnegative inventory to all such vectors. At fixed auxiliary runs, the earlier nonanchor exponents have

$$(4.4) \quad M_q = \prod_{5 < p < q} (1 + J_p), \quad \mathbb{E}M_q = \prod_{5 < p < q} \left(1 + \frac{C_p}{p-1}\right)$$

possible choices in the comparison. The empty product is one. Within a cell modulo $135 = 3^3 \cdot 5$, let u, v be the extra 3- and 5-depth runs. The seven anchor categories have moduli

$$(4.5) \quad \mathcal{G} = (3, 9, 27, 5, 15, 45, 135)$$

and coefficients

$$(4.6) \quad w(u, v) = (1, 1, 1 + u, 1 + v, 1 + v, 1 + v, (1 + u)(1 + v)).$$

For example, the 27-category represents depths at least three at 3 and depth zero at 5; the 135-category represents depths at least three at 3 and at least one at 5. The low-depth categories separate the remaining possibilities.

The term $1 + Z_q$ in (3.3) has baseline M_q and coefficients $M_q w(u, v)$. Indeed, the nonanchor exponent choices with no anchor component contribute $M_q - 1$, and the initial one in $1 + Z_q$ restores the missing constant. The normalized current depth weights satisfy

$$(q-1) \sum_{e \geq 1} q^{-e} = 1.$$

This identity is applied after the labelled convex comparison; different current exponents need not have the same residue.

Lemma 4.3 (Weighted aggregation). *Let f be increasing and convex. For a fixed background b_x , cell weights $a_x \geq 0$, and one projected category with nonnegative weights c_i of total $C > 0$,*

$$(4.7) \quad \begin{aligned} \sum_x a_x f\left(b_x + \sum_i c_i \mathbf{1}_{x \equiv r_i} (g)\right) &\leq \sum_i \frac{c_i}{C} \sum_x a_x f\left(b_x + C \mathbf{1}_{x \equiv r_i} (g)\right) \\ &\leq \max_r \sum_x a_x f\left(b_x + C \mathbf{1}_{x \equiv r} (g)\right). \end{aligned}$$

Proof. For each cell, the argument on the left is the convex combination of the arguments on the right. Apply Jensen pointwise and then sum. $\square$

Iterate this operation over unknown categories, leaving prescribed projections in the background. The geometric maximum may be taken over all nonempty projected residues in the core: an empty residue cannot improve the maximum when its coefficient is nonnegative. At modulus 135 the residue is a single cell.

All uses of infinite type inventories follow by monotone convergence from finite truncations. Section 8 proves finiteness of the resulting expectations and evaluates their remainders.

5. THE ANCHOR GEOMETRY AND ITS DELETION BUDGETS

All masses in the numerical formulas from this point are multiplied by 135. One *135-cell unit* is Haar mass $1/135$; no normalization by the number of retained cells is used.

Permutations of children at prime-adic tree nodes preserve Haar measure and carry every depth- e cylinder to another such cylinder. Normalize the selected pure 3-, 9-, and 5-classes to

$$0 \pmod{3}, \quad 1 \pmod{9}, \quad 0 \pmod{5}.$$

The 15-class is disjoint from 3 and 5. According as its ternary first digit agrees with that of the 9-class or not, it has representative $a = 1$ or $a = 2$ modulo 15. The remaining ternary symmetries give a representative $b = 2$ or $b = 4$ for the 27-class. The first digit c of the 25-class is either the 15-column or a different nonzero column; choose $c \in \{a, 3 - a\}$.

Thus there are eight coarse triples

$$(5.1) \quad a \in \{1, 2\}, \quad b \in \{2, 4\}, \quad c \in \{a, 3 - a\}.$$

Define

$$(5.2) \quad K_{a,b} = \{x \pmod{135} : x \not\equiv 0 \pmod{3}, 1 \pmod{9}, b \pmod{27}, 0 \pmod{5}, a \pmod{15}\}.$$

The choices $b = 2, 4$ respectively put the 27-class in the other surviving mod-3 branch, or in the same mod-3 branch as the 9-class but in a different mod-9 node. The normalizations preserve arbitrary deeper digits.

5.1. The basic reserve. The pure 3- and 5-survivors have product mass $(1/2)(3/4) = 3/8$. The total reciprocal mass of all mixed 3–5 types is

$$\sum_{i,j \geq 1} 3^{-i} 5^{-j} = \frac{1}{8}.$$

A union bound therefore gives the basic anchor reserve $135/4$. Write D_h for the relative pure-5 deletion inside nonzero column h , beyond the 5-class. Pure completion gives

$$(5.3) \quad D_h = \frac{1}{5} \mathbf{1}_{h=c} + t_h, \quad t_h \geq 0, \quad \sum_{h=1}^4 t_h = \frac{1}{20}.$$

The last identity scales the absolute pure-5 mass beyond depth two by the column mass $1/5$.

Put

$$\gamma_{15} = 3\mathbf{1}_{a=1} + \mathbf{1}_{b \equiv a \pmod{3}}.$$

These are the intersections of the selected 15-class with the disjoint 9- and 27-classes, in 135-cell units. The 15-class has total mass nine in these units. Its already deleted portion gives the reserve

$$(5.4) \quad R_{\text{basic}} = \frac{135}{4} + \gamma_{15} + (9 - \gamma_{15})D_a.$$

For upper costs, ignore deeper pure-3 exclusions and retain the pure-5 restriction. On a coarse cell, the zero extra-depth atoms are $2/3$ at 3 and $4/5 - D_{x \bmod 5}$ at 5. Their positive atoms are as in (4.3).

The t -simplex has vertices $t = (1/20)e_j$, $1 \leq j \leq 4$. This gives the 32 basic evaluations used in the certificate.

Lemma 5.1 (A uniform overlap credit). *When $b = 2$, the selected 75-class meets the completed pure-3 union in Haar mass at least $1/675$. In particular, $1/5$ may be added to (5.4) without increasing any upper loss bound.*

Proof. The 75-class is disjoint from the 3-class, so its residue modulo 3 is 1 or 2. In the first case it meets the 9-class in a cylinder of mass $1/\text{lcm}(75, 9) = 1/225$; in the second it meets the 27-class in a cylinder of mass $1/\text{lcm}(75, 27) = 1/675$. The basic union bound charged the full reciprocal $1/75$, so this already deleted part is a valid credit. For upper costs we may continue to ignore the additional hole, by Lemma 3.2. $\square$

This elementary credit is used for $(a, b, c) = (2, 2, 1)$. Its worst uncredited basic ledger under (1.3) is $-75775881/5000000000$. After adding $1/5$ it is

$$(5.5) \quad \frac{924224119}{5000000000} > 0.$$

Together with the other positive basic cases, this leaves only

$$(5.6) \quad (a, b, c) = (1, 4, 2), (2, 4, 2), (2, 4, 1).$$

These numerical assertions are included in the single finite proposition of Section 10.

5.2. The 45- and 75-classes. The selected 45-class is disjoint from its selected proper-divisor classes. Its residue r is therefore nonzero modulo 3 and 5, different from 1 modulo 9, and different from a modulo 15. Remove it from the core:

$$(5.7) \quad K = K_{a,b,r} := K_{a,b} \setminus \{x : x \equiv r \pmod{45}\}.$$

Representatives for r are selected by the signature

$$(5.8) \quad (r \bmod 3, \mathbf{1}_{r \equiv b \pmod{9}}, \mathbf{1}_{r \equiv a \pmod{5}}, \mathbf{1}_{r \equiv c \pmod{5}}).$$

Take the least permitted integer in each signature. These signatures are exactly the orbits under the remaining ternary and quinary tree permutations fixing the previously normalized data. In the ternary coordinate, a mod-9 branch is distinguished only by its first digit and whether it contains the 27-node. In the quinary coordinate, the distinguished columns are a and c .

Let the 75-class project to d modulo 3 and k modulo 5. Then

$$(5.9) \quad d \in \{1, 2\}, \quad k \in \{1, 2, 3, 4\}, \quad (d, k) \neq (a, a).$$

Set $B = \{x \in K : x \equiv d \pmod{3}, x \equiv k \pmod{5}\}$. In each such cell the 75-class removes a child of relative extra-5 mass $1/5$. If $k = c$, this child differs from the 25-child, since 25 divides 75 and completion made their classes disjoint.

Let e be the part of the higher pure-5 mass in that child. The joint budget is

$$(5.10) \quad t_h \geq 0, \quad \sum_h t_h = \frac{1}{20}, \quad 0 \leq e \leq t_k.$$

The retained extra-5 mass above $x \in K$ is

$$(5.11) \quad A_5(x) = 1 - D_{x \bmod 5} - (1/5 - e)\mathbf{1}_B(x).$$

This retains the possible overlap; the two deletions are not assumed disjoint. The five vertices of (5.10) are

$$(5.12) \quad (t, e) = ((1/20)e_j, 0) \ (1 \leq j \leq 4), \quad (t, e) = ((1/20)e_k, 1/20).$$

Their coefficients are $20(t_j - e\mathbf{1}_{j=k})$ and $20e$.

Ignoring deeper pure-3 restrictions for now, put $\gamma_{45} = \mathbf{1}_{r \equiv b \pmod{9}}$. A valid reserve is

$$(5.13) \quad R = \frac{135}{4} + \gamma_{15} + (9 - \gamma_{15})D_a + \gamma_{45} + (3 - \gamma_{45})D_{r \bmod 5} + \frac{9}{5} - |B|(1/5 - e).$$

The 15- and 45-classes are disjoint because 15 divides 45. The first two credits remove their already purely excluded portions from the reciprocal charges. The last credit replaces the full charge $135/75 = 9/5$ by an upper bound on the remaining 75-mass. Since B was defined after removing 45, overlap with that class is counted once. These same deletions determine both the reserve and the upper-cost measure.

At a vertex write $i = 0$ for the first four possibilities and $(j, i) = (k, 1)$ for the last. The extra-5 zero atom is $w_5(x)/20$, where

$$(5.14) \quad w_5(x) = 16 - 4\mathbf{1}_{x \equiv c \pmod{5}} - \mathbf{1}_{x \equiv j \pmod{5}} - (4 - i)\mathbf{1}_B(x).$$

Positive atoms remain $4/5^{v+1}$, $v \geq 1$.

The three domains in (5.6) have respectively $7 \cdot 7 \cdot 5 = 245$, $5 \cdot 7 \cdot 5 = 175$, and $8 \cdot 7 \cdot 5 = 280$ vertices. All d, k and all budget vertices are enumerated after selecting the representative r . The first two domains have positive certificates. Hence the remaining case has

$$(5.15) \quad (a, b, c) = (2, 4, 1), \quad r \in \{4, 7, 8, 11, 16, 22, 31, 34\}.$$

5.3. The deeper pure-3 budget. The pure 3-, 9-, and 27-classes have total mass $13/27$. Let z_h be the relative additional pure-3 mass in a surviving row h modulo 27. There are fourteen such rows, and

$$(5.16) \quad z_h \geq 0, \quad \sum_{\substack{h \pmod{27} \\ h \not\equiv 0 \pmod{3}, 1 \pmod{9}, 4 \pmod{27}}} z_h = \frac{1}{2}.$$

The last value is $27(1/2 - 13/27)$. In a row h the retained extra-3 mass is $1 - z_h$, so its zero comparison atom is $2/3 - z_h$. At a vertex supported in

row z , this atom is $w_3(x)/6$, where

$$(5.17) \quad w_3(x) = 4 - 3\mathbf{1}_{x \equiv z \pmod{27}}.$$

The positive atoms are $2/3^{u+1}$, $u \geq 1$.

In the canonical case, the refined reserve is obtained from (5.13) by setting

$$(5.18) \quad \begin{aligned} \gamma_{15} &= \sum_{h \equiv 2 \pmod{3}} z_h, \\ \gamma_{45} &= \mathbf{1}_{r \equiv 4 \pmod{9}} + \sum_{h \equiv r \pmod{9}} z_h, \\ |B| &\longmapsto \sum_{x \in B} (1 - z_{x \bmod 27}). \end{aligned}$$

The formula follows by computing the same intersections with the larger pure-3 union. It is affine separately in z and (t, e) . The reserve increases when ignored pure-3 exclusions are restored.

Above each retained coarse cell, the selected extra-3 and extra-5 restrictions factor into a product. Other mixed anchor classes are ignored in the upper-cost measure and are paid for in the reserve. Thus Lemmas 3.2 and 4.2 apply with exactly the weights just described.

6. CHARGED FORBIDDEN FIBRES AT PRIME SEVEN

Normalize the pure 7-class to target zero. The classes with selected moduli 21, 35, 63, 105 are disjoint from it. In the canonical anchor case, their four anchor projections form a tuple ξ in

$$(6.1) \quad \Xi = \{1, 2\} \times \{1, 2, 3, 4\} \times \{2, 4, 5, 7, 8\} \times \{1, 4, 7, 8, 11, 13, 14\}.$$

The restrictions follow from their selected proper divisors: 63 avoids the 3- and 9-classes, and 105 avoids the 3-, 5-, and 15-classes. Thus $|\Xi| = 280$.

For $x \in K$ define the number of active projections

$$(6.2) \quad s_\xi(x) = \mathbf{1}_{x \equiv \xi_1 \pmod{3}} + \mathbf{1}_{x \equiv \xi_2 \pmod{5}} + \mathbf{1}_{x \equiv \xi_3 \pmod{9}} + \mathbf{1}_{x \equiv \xi_4 \pmod{15}}.$$

Their active target digits modulo 7 may coincide. We will need neither their equality pattern nor a global relabeling of them.

Lemma 6.1 (Charged fibre enlargement). *There is an admissible deletion process, containing all actual mixed-7 deletions, with conditional current-loss bound*

$$(6.3) \quad \ell_7(x) \leq \frac{1}{4} \left[R_7(x) + \frac{6}{7} s_\xi(x) - 1 \right]_+,$$

where $R_7/6$ bounds the remaining mixed-7 Haar union after the four selected classes are extracted. Its live 7-fibre has mass at most

$$(6.4) \quad S_{s_\xi(x)} = \min \left(1, \frac{3}{2} \frac{6 - s_\xi(x)}{7} \right).$$

Consequently its comparison zero atom at 7 may be taken to be

$$(6.5) \quad \rho_s = S_s - \frac{3}{14} = \frac{1}{14}(11, 11, 9, 6, 3)_s, \quad 0 \leq s \leq 4,$$

with positive atoms $9/7^{j+1}$, $j \geq 1$.

Proof. Let $T(x)$ be the set of active nonzero target digits. Then $|T(x)| \leq s_\xi(x) \leq 4$. Enlarge it to a subset $T^+(x)$ of $\{1, \dots, 6\}$ of size exactly $s_\xi(x)$, for example by adding the smallest unused digits. This is a measurable choice depending on the coarse cell. Declare those entire first-digit cylinders forbidden, in addition to the remaining actual bad set.

Their Haar mass is $s_\xi(x)/7$. The pure union has mass $1/6$. Lemma 3.1, with $C_7 = 3/2$, bounds the loss by

$$\left[1 - \frac{3}{2} \left(1 - \frac{1}{6} - \frac{R_7(x)}{6} - \frac{s_\xi(x)}{7} \right) \right]_+,$$

which is (6.3). Every survivor avoids zero and the $s_\xi(x)$ added digits, so the density cap and normalization give (6.4). Since $S_s \geq 3/14$, the last assertion follows from Lemma 4.2. $\square$

The extra forbidden digits need not define globally fixed congruences. They are part of the measure construction. They contain the actual forbidden fibres, and their full volume has been charged in (6.3). The current charge and the later surviving mass refer to this same enlarged process.

After the four first-current-power types are extracted, the coefficients for the remaining inventory at 7 are

$$(6.6) \quad c(u, v) = \left(\frac{1}{7}, \frac{1}{7}, 1 + u, v + \frac{1}{7}, v + \frac{1}{7}, 1 + v, (1 + u)(1 + v) \right).$$

Each extraction removes $6/7$ once from its category, and the four removed terms reappear as the fixed offset $6s_\xi(x)/7$. The geometric numerator for the current bound is therefore

$$(6.7) \quad \max_{\mathbf{r}} \sum_{x \in K} a_x \left[-1 + \frac{6}{7}s_\xi(x) + \sum_{g \in \mathcal{G}} c_g(u, v) \mathbf{1}_{x \equiv r_g \pmod{7}} \right]_+.$$

Here and below a_x denotes the relevant comparison weight, not the residue of a congruence. The original residues are fixed; allowing the unknown $\mathbf{r}$ to vary inside the auxiliary expectation is a conservative upper bound.

6.1. A coarse screening inequality. For $s \geq 0$ and $R \geq 0$,

$$(6.8) \quad \frac{1}{4} \left[R + \frac{6}{7}s - 1 \right]_+ \leq \frac{1}{4} \left[R - \frac{1}{7} \right]_+ + \frac{3}{14}(s - 1)_+.$$

For an unrefined canonical vertex, the total retained anchor comparison mass at cell x is $W(x)/20$, where

$$W(x) = 20 - 4\mathbf{1}_{x \equiv 1 \pmod{5}} - \mathbf{1}_{x \equiv j \pmod{5}} - (4 - i)\mathbf{1}_B(x).$$

Let C_{common} be the upper expectation of the first term in (6.8), with unknown residues maximized and coefficients (6.6). A valid current bound is

$$(6.9) \quad C_{\text{common}} + \frac{3}{280} \sum_{x \in K} W(x)(s_{\xi}(x) - 1)_+.$$

For its omitted anchor depths the checker uses the nonnegative linear load with baseline $13/7$ and threshold 2 ; its finite hinge is exactly $[R - 1/7]_+$. This specifies the harmless upward slack in the coarse remainder.

For later stages, replace only the zero atom of J_7 by (6.5). The positive-depth cylinder caps remain unchanged. This retained first-hit information is the principal improvement over the ordinary scalar inventory.

7. ONE PRESCRIBED PROJECTION AT PRIME ELEVEN

Completion supplies the 165-class, disjoint from its selected 3-, 5-, 11-, and 15-divisor classes. Its anchor projection satisfies

$$(7.1) \quad \eta = a_{165} \bmod 15 \in \mathcal{E} := \{1, 4, 7, 8, 11, 13, 14\}.$$

Its target modulo 11 need not be specified. Its additive current load is a valid upper bound irrespective of other targets.

At stage 11 put $m = 1 + J_7$. The selected 165-class is one first-11-power type with zero earlier 7-exponent. Its normalized weight is $10/11$, not $10m/11$. Extract it from the 15-category and retain it as a fixed offset. The remaining coefficients are

$$(7.2) \quad b_g(m, u, v) = mw_g(u, v) - \frac{10}{11} \mathbf{1}_{g=15}.$$

They are nonnegative, including $m = 1, v = 0$. The relevant numerator is

$$(7.3) \quad F_{\eta}(a; m, u, v) = \max_r \sum_{x \in K} a_x \left[m - 4 + \frac{10}{11} \mathbf{1}_{x \equiv \eta \pmod{15}} + \sum_{g \in \mathcal{G}} b_g(m, u, v) \mathbf{1}_{x \equiv r_g \pmod{11}} \right]_+.$$

The current-loss factor is $1/\beta_{11} = 1/6$.

Proposition 7.1 (The fixed-label upper bound). *For each fixed physical projection η , averaging (7.3) with the anchor comparison and the post-first-hit 7-comparison gives a valid stage-11 upper loss. A uniform upper bound is obtained by maximizing over η after that expectation.*

Proof. Keep the selected first-current-power term as the fixed nonnegative background $(10/11) \mathbf{1}_{x \equiv \eta \pmod{15}}$. Apply Lemma 4.1 to the remaining labelled cylinders in reverse coordinate order, using (6.5) at 7. Then apply Lemma 4.3 only to the remaining nonnegative weights. This gives (7.3) with its prescribed η unchanged. The same kernel of Lemma 3.1 is being bounded, so the earlier current loss and all later caps remain compatible. $\square$

The order is

$$(7.4) \quad \max_{\eta \in \mathcal{E}} \mathbb{E}_{\text{aux}} F_{\eta}(a; m, u, v),$$

with the maximum over unknown residues already inside F_{η} . One does not reselect η at each auxiliary state. The latter is a valid but weaker bound. The improvement comes from preserving a residue that belongs to one actual selected class throughout the calculation.

8. EXACT GEOMETRIC FUNCTIONALS AND INFINITE REMAINDERS

We specify the finite computation independently of a particular implementation. Given nonnegative cell weights $a = (a_x)_{x \in K}$, define

$$(8.1) \quad \Phi_t(a; u, v) = \max_r \sum_{x \in K} a_x \left[1 - t + \sum_{g \in \mathcal{G}} w_g(u, v) \mathbf{1}_{x \equiv r_g(g)} \right]_+.$$

For each of the first six moduli in (4.5), the maximum ranges over its nonempty projections on K . The last residue is a cell of K . This is a finite maximum of linear functions of a , with nonnegative hinge values.

Put $A(a) = \sum_x a_x$, $a_* = \max_x a_x$, and

$$h_g(a) = \max_r \sum_{\substack{x \in K \\ x \equiv r(g)}} a_x.$$

The maximum full-cell positive linear load is

$$(8.2) \quad \begin{aligned} \Lambda(a; u, v) = & A(a) + h_3(a) + h_9(a) + (1 + u)h_{27}(a) \\ & + (1 + v)(h_5(a) + h_{15}(a) + h_{45}(a)) + (1 + u)(1 + v)a_*. \end{aligned}$$

For a full linear sum, each residue can maximize its contribution independently. We do not use this separation for the nonlinear hinge in (8.1).

8.1. The four anchor regions. Let $\sigma_3(u) = 2/3^{u+1}$ and $\sigma_5(v) = 4/5^{v+1}$ on positive integers. Use these masses on positive-depth coordinates and a unit atom at zero on zero-depth coordinates. At a vertex, the four region contributions use the weights and divisors in Table 1. The mass in the last column is the product of the two scalar depth masses divided by the divisor; it does not include the cell weight.

u	v	Integer cell weight	Divisor	Total scalar mass
> 0	> 0	1	1	1/15
0	> 0	w_3	6	1/30
> 0	0	w_5	20	1/60
0	0	$w_3 w_5$	120	1/120

TABLE 1. The anchor comparison. For an unrefined row, $w_3 = 4$. For a basic vertex, omit the B term in w_5 .

Write $\mathcal{A}[G]$ for the sum of these four weighted expectations of a geometric functional $G(a; u, v)$. This notation incorporates the cell-dependent zero atoms. It is not an expectation under a cell-independent normalized anchor law.

To bound $\mathcal{A}[\Phi_t]$, retain $u < 12$, $v < 9$ and use Λ outside that rectangle. Denote the resulting upper value by $F(t)$. Also set

$$(8.3) \quad \mathcal{M} = \mathcal{A}[A], \quad \mathcal{W} = \mathcal{A}[\Lambda].$$

The same construction applies to (6.7), the coarse hinge, and (7.3), using their own nonnegative full linear loads. All omitted-region terms are positive upper sums.

Lemma 8.1 (Exact anchor remainders). *Let $\sigma_p(j) = (p-1)/p^{j+1}$ for $j \geq 1$. For every integer $L \geq 1$,*

$$(8.4) \quad \begin{aligned} \sum_{j \geq L} \sigma_p(j) &= p^{-L}, & \sum_{j \geq L} j \sigma_p(j) &= p^{-L} \left(L + \frac{1}{p-1} \right), \\ \sum_{j \geq 1} \sigma_p(j) &= \frac{1}{p}, & \mathbb{E}(j \mid j > 0) &= \frac{p}{p-1}. \end{aligned}$$

Consequently every omitted-region linear load used here has an exact rational expectation.

Proof. Sum the geometric series and its first derivative. The loads are affine separately in u and v , with at most a uv term. Their product-measure expectations depend only on the masses and first moments in (8.4). Partition the omitted region into $u \geq 12$ with arbitrary v , and $u < 12$ with $v \geq 9$. Zero-depth slices are treated by their fixed weights. These two pieces are disjoint. $\square$

Equivalently, one may subtract the retained positive linear sum from the full linear expectation. Positivity and convexity are justified by the omitted-region representation, not by an arbitrary difference of upper bounds.

8.2. Multiplier distributions. Store the exact probabilities $\pi(m)$ for $1 \leq m < 32$ and the full first moment $\bar{m}$. Initially $M = 1$. When adjoining a prime p with cap C , update by

$$(8.5) \quad \begin{aligned} \pi'(n) &= \sum_{m(j+1)=n} \pi(m) \alpha_p(j) \quad (n < 32), \\ \bar{m}' &= \bar{m} \left(1 + \frac{C}{p-1} \right), \end{aligned}$$

where α_p is (4.2). Every multiplier factor is at least one; a discarded value cannot return below 32.

Lemma 8.2 (Exact multiplier remainder). *For threshold $t \leq 12$, a valid ordinary numerator expectation is*

$$(8.6) \quad \begin{aligned} \mathcal{C}_t = & \sum_{m < t} \pi(m) m F(t/m) \\ & + \left(\bar{m} - \sum_{m < t} m \pi(m) \right) \mathcal{W} - t \left(1 - \sum_{m < t} \pi(m) \right) \mathcal{M}. \end{aligned}$$

All sums and remainders in this expression are exact rationals.

Proof. If $m < t$, factor m from the hinge and use $F(t/m)$. If $m \geq t$, the baseline $m - t$ is nonnegative at every cell, so the hinge is linear and its exact maximum is $m\Lambda - tA$. Its remaining expectation uses only the remaining mass and first moment. All $m < t$ are stored below 32. $\square$

Divide (8.6) by β_q for stage q , using the distribution of M_q . The required ratios are

$$(8.7) \quad \mathcal{T} = \left\{ \frac{t}{m} : t \in \{2, 4, 8, 12\}, 1 \leq m < t \right\}; \quad |\mathcal{T}| = 15.$$

The all-height calculation therefore uses a finite set of weighted hinge queries, together with the exact formulas above.

8.3. Retaining the first-hit zero component. For $q > 7$, the ordinary comparison can be partitioned according to $J_7 = 0$ or $J_7 > 0$. Let $Z_q(a)$ be the numerator computed from (8.6) with 7 omitted from the multiplier product; the other caps and the current threshold are unchanged. Let $\kappa_\xi(x) = 14\rho_{s_\xi(x)}$, so that $\kappa_\xi(x) \in \{11, 9, 6, 3\}$. The live numerator upper bound is

$$(8.8) \quad \mathcal{C}_q^{\text{live}} = \mathcal{C}_q - \frac{11}{14} Z_q(a) + \frac{1}{14} Z_q(\kappa_\xi a).$$

The ordinary zero atom $11/14$ is replaced by the spatial zero atom (6.5). This is an identity between specified positive comparison components. In particular, the subtraction removes exactly the same finite and omitted-height zero component that occurred in $\mathcal{C}_q$. It is not subtraction of unrelated upper estimates. Equivalently one can compute the positive-depth 7-components and the new zero component directly.

8.4. The fixed-165 remainder. At stage 11 the multiplier law is

$$(8.9) \quad \mathbb{P}(M = 1) = \frac{11}{14}, \quad \mathbb{P}(M = m) = \frac{9}{7^m} \quad (m \geq 2),$$

with its first term replaced by the spatial weights $\kappa_\xi/14$. Thus for $m = 1$, use cell weights $\kappa_\xi a$ and factor $1/14$ in (7.3); for $m = 2, 3$, use ordinary weights and factors $9/49, 9/343$.

The maximum positive linear load for the fixed label is

$$(8.10) \quad \begin{aligned} \Lambda_\eta(a; m, u, v) = & mA(a) + \sum_{g \in \{3, 9, 27, 5, 15, 45\}} b_g(m, u, v) h_g(a) \\ & + m(1+u)(1+v)a_* + \frac{10}{11} \sum_x a_x \mathbf{1}_{x \equiv \eta} \pmod{15}. \end{aligned}$$

For $m = 1, 2, 3$ use the retained rectangle and this load outside it. For $m \geq 4$, every hinge is linear, so the exact value is $\Lambda_\eta - 4A$. The remaining multiplier mass and first moment are

$$(8.11) \quad \sum_{m \geq 4} \frac{9}{7^m} = \frac{3}{686}, \quad \sum_{m \geq 4} \frac{9m}{7^m} = \frac{25}{1372}.$$

Equations (8.4), (8.10), and (8.11) specify the full stage-11 expectation. Divide the total numerator by six.

9. ONE VERTEX ARGUMENT FOR THE COMMON FUNCTIONAL

Fix the completed family, its physical anchor data, the projections ξ, η , and the caps (3.4). Define the common upper functional by the exact spatial stage-7 comparison, the positive first-hit continuation comparison, and the fixed-165 stage-11 comparison, with the positive anchor-tail majorants specified above.

Each geometric value, at fixed auxiliary parameters, is a maximum of linear functions of its nonnegative cell weights. The expectations are positive sums of these values. For continuous budgets, a zero extra-3 weight is affine in z and a zero extra-5 weight is affine in (t, e) . Their product gives the weights on the doubly zero region. The reserve has the same separate affinity.

Lemma 9.1 (Explicit vertex interpolation). *For z, t, e satisfying (5.16) and (5.10), set*

$$(9.1) \quad \lambda_{h,j} = 40z_h(t_j - e\mathbf{1}_{j=k}), \quad \lambda_{h,*} = 40z_h e.$$

These coefficients are nonnegative and sum to one. At each auxiliary region, the comparison weights and the reserve are the corresponding convex combinations of their joint vertex values. Hence the actual lower ledger is at least the same convex combination of its vertex lower bounds.

Proof. The coefficients $2z_h$ decompose the pure-3 simplex. The mixed-hole budget has coefficients $20(t_j - e\mathbf{1}_{j=k})$ and $20e$. Multiply these identities. The resulting coefficients are (9.1) and have total $40(1/2)(1/20) = 1$. Every comparison weight is the product of the two affine coordinate factors, and the reserve is separately affine. Thus the interpolation identities follow. For each geometric maximum G ,

$$G\left(\sum_v \lambda_v a^{(v)}\right) \leq \sum_v \lambda_v G(a^{(v)}).$$

Positive expectations and the positive omitted-region bounds preserve this inequality. Subtract from the affine reserve. $\square$

The simpler one-budget statements used in the earlier screens are the same argument with the other budget ignored. It remains to justify using different upper estimates at different vertices.

Lemma 9.2 (Compatible screening bounds). *Every earlier screen used below gives a lower bound for the same fixed-family lower ledger to which Lemma 9.1 applies. It is valid to stop at any vertex whose screen is positive and to sharpen the bound only at the other vertices.*

Proof. Ignoring anchor holes enlarges the comparison weights and reduces the credited reserve. At stage 7, the ordinary inventory dominates the comparison with prescribed projections by Lemma 4.3: the four selected contributions are nonnegative terms in that inventory. Inequality (6.8) gives another valid stage-7 upper bound. For later stages, ignoring the first-hit reduction replaces $\kappa_\xi \leq 11$ by the larger ordinary zero weight 11. At stage 11, releasing the prescribed 165-residue to the geometric maximum is again the weighted aggregation inequality. The corresponding positive linear tail loads have the same dominations. All bounds use the same caps.

Thus every stopping value is a lower bound for the common functional at the indicated vertex, or uniformly over the budget that has been ignored. Lemma 9.1 may interpolate these certified vertex numbers. There is no need for the minimum of the screening upper functions to be convex. $\square$

The caps are fixed across the entire calculation. Each vertex therefore bounds the same prescribed deletion process. Likewise, convexity of (8.8) is justified through its positive component representation, not inferred from its subtraction form.

10. THE FINITE CERTIFICATE AND THE MAIN THEOREM

We now state the sole computational proposition. Every input domain, coefficient, geometric maximum, and infinite remainder in its statement has been defined above.

Proposition 10.1 (Finite positivity certificate). *Use the caps (3.4), the anchor domains of Section 5, and the comparisons of Sections 6–8. At every required vertex the screening procedure in Table 2 terminates with a positive lower ledger. If R is its reserve and L the sum of its six certified loss bounds, both in 135-cell units, then every recorded terminal inequality satisfies*

$$(10.1) \quad R^- := \lfloor 1000R \rfloor, \quad L^+ := \lceil 1000L \rceil, \quad R^- - L^+ \geq 4.$$

There are 28,001 distinct recorded inequalities. The minimum integer surplus is exactly four.

Evaluation	Number	Positive entries	Continuation
Basic anchor vertices	32	20	Three coarse triples
Mixed anchor vertices	700	604	96 canonical vertices
Projection coarse screen	26,880	25,840	1,040 spatial states
Exact spatial comparison	1,040	1,003	37 states
Pure-3 refinement	518	506	12 residual vertices
Fixed-165 closing bounds	28	28	None

TABLE 2. Complete screening domain. The 20 basic entries include four using Lemma 5.1. The mixed positive entries comprise 420 noncanonical and 184 canonical vertices.

Verification. Enumerate (5.1) and their four t -vertices. Four triples close directly; the triple $(2, 2, 1)$ closes with (5.5). For each of the three remaining triples enumerate the least representatives of (5.8), all seven pairs in (5.9), and all five vertices in (5.12). This gives 700 evaluations. The two noncanonical triples close, as do 184 canonical vertices.

For the other 96 canonical vertices enumerate every $\xi \in \Xi$. Use (6.9) at 7 and the ordinary bounds at later primes. Of these $96 \cdot 280 = 26880$ comparisons, 1040 require the exact spatial stage-7 bound and the first-hit correction (8.8). Exactly 37 of those require pure-3 refinement. Enumerate all fourteen rows in (5.16) at each of these states, giving 518 evaluations. Twelve residual vertices reduce to the four states in Table 4 by the symmetries proved below. For each representative, enumerate all seven $\eta \in \mathcal{E}$ in (7.3). All 28 bounds close.

All arithmetic uses integers and rational numbers. Ordinary and current stage bounds may first be rounded upward to 10^{-10} in 135-cell units. For (8.8), only the ordinary component is rounded upward; its subtracted zero component and replacement are exact. The fixed-label stage bound is rounded upward after its full expectation. Finally, the total of the six valid stage bounds is rounded once as in (10.1). Independent rounding of the six stages to thousandths is not used.

The standalone program `checks/verify.py` implements exactly these loops and formulas. It compiles the finite enumerator of Appendix A and regenerates the certificate without importing earlier result tables. The provided certificate records each phase, its configuration, and the two integers R^-, L^+ . Its counts and minimum surpluses are given in Table 3. $\square$

10.1. The four terminal states. Every residual vertex lies in the case $r = 8$. Its remaining data satisfy $(d, k, i) = (2, 1, 0)$, $j \in \{1, 3\}$, and $z \in \{13, 22\}$ for the pure-3 row. Its projection tuple is one of

$$(10.2) \quad (1, 4, 7, 14), \quad (2, 4, 2, 14), \quad (2, 4, 5, 14).$$

Positive certificate phase	Entries	Minimum $R^- - L^+$
Basic, without the extra 75-credit	16	221
Basic, with the extra 75-credit	4	184
Noncanonical mixed anchors	420	182
Canonical anchors	184	5
Projection coarse screen	25,840	4
Exact spatial comparison	1,003	15
Pure-3 refinement	506	6
Fixed-165 comparison	28	12

TABLE 3. Integer witnesses, each unit corresponding to Haar mass $1/135000$. The uniform closing bound below has surplus 11 because it separately rounds the available stage-11 budget.

All remaining configurations are thus on the 44-cell core

$$K_{2,4,8} = \{x \pmod{135} : x \not\equiv 0 \pmod{3}, 1 \pmod{9}, 4 \pmod{27}, 0 \pmod{5}, 2 \pmod{15}, 8 \pmod{45}\}.$$

Swap rows 13 and 22 within their common mod-9 node to normalize $z = 13$. When the third projection is 5, swap the mod-9 branches 2 and 5, carrying all their descendants with them, to change it to 2. The swaps act in different mod-3 branches. They preserve the core, the weights, and the relevant projection field, and permute each category in (4.5) into itself. They fix residues modulo 3 and modulo 5, and therefore fix each 165-projection modulo 15 individually. Both permutations extend to the full prime-adic trees by leaving all later relative digits unchanged.

Write A for $(1, 4, 7, 14)$ and B for $(2, 4, 2, 14)$. Together with $j = 1, 3$, these give A1, B1, A3, B3. Let B_{11} be the reserve after subtracting the five losses other than stage 11. Table 4 gives its downward rounding and the worst upward-rounded fixed-label stage-11 loss L_{11} .

State	$\lfloor 1000B_{11} \rfloor$	$\max_{\eta} \lceil 1000L_{11}(\eta) \rceil$	Difference
A1	5310	5299	11
B1	5418	5281	137
A3	5522	5292	230
B3	5591	5282	309

TABLE 4. The four terminal states, with pure-3 row 13. All seven fixed projections are evaluated at each state.

In particular, the single inequality

$$(10.3) \quad B_{11} - L_{11} \geq \frac{5310 - 5299}{1000} = \frac{11}{1000} > 0$$

closes every terminal state. The smaller global surplus four comes from an earlier coarse screen.

Proof of Theorem 1.1. Suppose an odd distinct cover has rank at most eight. Lemma 2.1 gives a finite cover supported on $\mathcal{P}$, and its lift covers X . Apply Lemma 2.2 with (2.2). The completed family still covers X .

Normalize its anchor data and fix its actual continuous budgets and its projections ξ, η . Use the normalized kernels of Lemma 3.1, with the charged enlargement at 7 from Lemma 6.1. Every completed class is either excluded at the anchor stage, avoided as a pure class, or deleted when its largest prime is exposed. The final submeasure therefore lies outside the completed covered set.

Proposition 10.1 and Lemmas 9.1–9.2 give, for the actual budgets and fixed projections,

$$135H(A) - 135 \sum_i L_i \geq \frac{4}{1000}.$$

First-hit accounting now gives

$$(10.4) \quad \nu_{\text{final}}(X) \geq \frac{4}{135000} = \frac{1}{33750} > 0.$$

A positive measure cannot be supported outside a set covering all of X . This contradiction proves the theorem. $\square$

The mass in (10.4) is that of the constructed submeasure. Appendix C converts it to a bound for the Haar density of the original uncovered set and extends that bound to arbitrary odd prime support.

Prospects for rank ten. A rank-ten lower bound would require excluding nine-prime support, so the next stage is $q = 29$. At the most difficult terminal state A1 with $\eta = 14$, the reserve is 33.75 and the six upper charges, in increasing prime order, are approximately

$$(10.406267, 5.298044, 6.175771, 3.906187, 4.365028, 3.586210).$$

Their exact sum leaves 0.012492546... units. With these caps fixed, the cell $x = 23$ has zero-depth anchor weight $8/15$ and $s_\xi(23) = 0$. Aligning all seven categories there bounds the added *geometric comparison charge*, for every $1 \leq \tau < 28$, by

$$\hat{L}_{29}(\tau) \geq \frac{8}{15} \frac{\mathbb{E}(8M_{29} - \tau)_+}{28 - \tau} \geq \frac{2180671}{9302400} > 0.234420.$$

Here $\mathbb{E}M_{29} = 22869/10240$ and $\mathbb{P}(M_{29} = 1) = 688/1615$. Checking the breakpoints 8, 16, 24 gives the minimum at $\tau = 16$; the companion supplies the endpoint checks. More strongly, its exact 13-interval certificate allows all seven caps to vary within the comparison laws: with the same anchor geometry, padding rule, and prescribed projections, the total charge is at least $34.099 > 33.75$ at this state. Thus cap retuning alone cannot complete this extension; rank ten needs a sharper comparison or additional geometric information. These are bounds on comparison charges, not on actual deleted mass. The diagnostic is separate from the Lean formalization of the rank-nine theorem and its corollaries.

11. VERIFICATION AND REPRODUCIBILITY

The computational companion contains a Python checker using only the standard library, a C++17 integer geometry routine, the integer certificate, and the 28 rational closing records. The Python program derives its finite domains from the formulas in the paper. Its input is neither an earlier certificate nor a numerical optimization output. Running

```
python3 checks/verify.py -fresh
```

from the companion directory removes its geometry cache, recompiles the enumerator, regenerates all required maxima, and checks the complete hierarchy. A run without `-fresh` reuses geometry keyed by the full integer input. Floating-point numbers are used only to display elapsed time.

A fresh reconstruction passed all domain assertions and recovered the 28,001 published integer inequalities. The same integer certificate was also obtained through the earlier, separately organized geometric implementation. Boundary expectations were compared using independent geometric and tail summations during the audit. The final companion removes the earlier software dependencies; their agreement supplies an additional check rather than an input to its proof calculation.

The computational assertion is finite: residue maxima, geometric rational sums, domain coverage, symmetries, and outward rounding. The passage from those numbers to a statement about all congruence families uses the mathematical lemmas proved above. In particular, the program does not replace the countable-completion argument, the conditional comparison, the charged-fibre construction, or the common-functional interpolation.

The accompanying development uses Lean 4 [14] and mathlib [12]. It proves Theorem 1.1 and Corollaries 1.2, 1.3, and C.2 without additional assumptions about obstructions, certificate correctness, or exponent bounds. The principal public statements are

```
Rank9.nine_prime_support_of_integer_cover,
Rank9.strengthened_lcm_lower_bound_of_integer_cover,
Rank9.uncovered_density_of_integer_family.
```

The rank theorem is in `Rank9/RankNine.lean`; the quantitative corollaries are in `Rank9/Corollaries.lean`. The formal argument uses finite rational measures on word spaces of arbitrary heights; exact mass and moment remainders control all greater exponents. For smaller prime supports, the formal arithmetic reduction reuses the proved rank-eight theorem from the preceding Lean development. This is a difference from the reduction in this manuscript and introduces no additional axiom.

The stronger LCM bound uses the elementary nine-prime square-free exclusion in the proof of Corollary 1.3, whose exact rational union bound is checked by Lean. For density, completion preserves every originally covered point without a covering assumption. Lean bounds the mass of each partially exposed word prefix and obtains the full counting-measure bound after all eight coordinates have been exposed. Finite digit-translation averaging

transfers this bound to arbitrary prime supports, and a separate limit theorem identifies the one-period fraction with the natural density. The density endpoint assumes only distinct odd moduli greater than one and at most eight prime divisors; it has no covering hypothesis. For related formal work, Mian and Siddique [13] report a Lean verification of the finite exclusion $N > 10000$ for odd distinct covers. Their result is not a dependency of this development.

The geometry inequalities and their consequences for the actual losses are proved in Lean, as is the exhaustive case assembly. The external numerical programs supply no trusted oracle. The companion command

`python3 scripts/verify.py`

is run from the companion's `formal/` directory. It rebuilds the proof, checks the public statements and the axiom dependencies of all project theorems, and performs an additional kernel replay. That replay checks the foundation from an empty environment and then checks each local module after its local imports have passed, using the pinned Lean toolchain's own replay routines. The permitted axioms are only `propext`, `Classical.choice`, and `Quot.sound`.

The accompanying archive `lean_proof.tar.gz` contains the full Lean source companion. Extracting it supplies the `formal/` directory, pinned dependency manifests, a theorem map, and a verification summary. It pins Lean to `leanprover/lean4:v4.32.0-rc1` and `mathlib` to revision

360da6fa66c1273b76b6b2d8c5666fd5ac2e3b56.

`PROOF_ACCEPTANCE.json` states the mathematical scope. The file hashes are listed in `SHA256SUMS`.

APPENDIX A. A FINITE INTEGER GEOMETRY ALGORITHM

The shared primitive receives a finite core K , nonnegative integer weights a_x , integer offsets o_x , a baseline b , six nonnegative coefficients c_g , a nonnegative singleton coefficient s , and a threshold t . After clearing a common denominator, it computes

$$(A.1) \quad \max_{\substack{r_g \ (g=3,9,27,5,15,45) \\ h \in K}} \sum_{x \in K} a_x \left[b - t + o_x + \sum_g c_g \mathbf{1}_{x \equiv r_g \ (g)} + s \mathbf{1}_{x=h} \right]_+.$$

The same routine handles ordinary, spatial, and fixed-label queries. For the spatial stage-7 query, the common denominator is 7, the baseline is zero, the threshold is 7, and the offset is $6s_\xi(x)$. The coefficient vector is

$$(1, 1, 7(1+u), 7v+1, 7v+1, 7(1+v)).$$

For the fixed-165 query, the common denominator is 11, the baseline is $11m$, the threshold is 44, and the offset is $10\mathbf{1}_{x \equiv \eta \ (15)}$. The 15-coefficient is $11m(1+v) - 10$; the other coefficients are obtained directly from (7.2).

An efficient exact enumeration avoids looping over the 27-residue and singleton independently. Fix the other five residues, and write

$$y_x = b - t + o_x + \sum_{g \in \{3,9,5,15,45\}} c_g \mathbf{1}_{x \equiv r_g (g)}.$$

Let $B_0 = \sum_x a_x [y_x]_+$. For each nonempty row r modulo 27, compute

$$\begin{aligned} \Delta_r &= \sum_{x \equiv r (27)} a_x ([y_x + c_{27}]_+ - [y_x]_+), \\ U_r &= \max_{x \equiv r (27)} a_x ([y_x + s]_+ - [y_x]_+), \\ (A.2) \quad V_r &= \max_{x \equiv r (27)} a_x ([y_x + c_{27} + s]_+ - [y_x + c_{27}]_+). \end{aligned}$$

Then the exact best value for these five fixed residues is

$$(A.3) \quad B_0 + \max_r \left\{ \Delta_r + \max \left(V_r, \max_{r' \neq r} U_{r'} \right) \right\}.$$

An empty outside maximum is interpreted as zero.

Correctness. Choosing row r adds exactly Δ_r . The singleton is either in that row, with best additional gain V_r , or outside it, with best gain $\max_{r' \neq r} U_{r'}$. Maximize over r , then enumerate the five other residues. Every possible placement in (A.1) is covered. The two largest values of U_r suffice to evaluate every outside maximum, including ties. $\square$

The singleton is optimized on the same residue pattern as the base sum. This matters: adding a singleton improvement from an independently maximizing pattern would not be the claimed exact maximum.

The supplied routine uses signed 64-bit integers. In the finite queries of this paper, $|K| \leq 135$, the integer cell weights are at most $4 \cdot 16 \cdot 11 = 704$, and a cell's positive load after clearing denominators is below 8192. Consequently a complete weighted sum is below

$$135 \cdot 704 \cdot 8192 = 778,567,680 < 2^{30}.$$

The checker asserts the input bounds. The larger rational numerators arising in the outer expectations are handled by Python's arbitrary-precision integers.

APPENDIX B. FINITE DOMAINS AND SOURCE CORRESPONDENCE

The unweighted incidences provide small checks on the geometric definitions. For the basic core $K_{a,b}$, put $n = |K_{a,b}|$ and $h_g = \max_r |\{x \in K_{a,b} : x \equiv r \pmod{g}\}|$. Direct enumeration gives Table 5. These numbers are not substituted for the simultaneous hinge maximum; they enter only the linear remainder bounds.

For the canonical triple $(2, 4, 1)$ the pure-3 budget rows are

$$\{2, 5, 7, 8, 11, 13, 14, 16, 17, 20, 22, 23, 25, 26\}.$$

TABLE 5. Basic-core cardinalities and maximal incidences.

a	b	n	h_3	h_9	h_{27}	h_5	h_{15}	h_{45}
2	2	48	24	12	4	14	8	3
2	4	47	27	12	4	14	9	3
1	2	50	32	12	4	14	8	3
1	4	51	36	12	4	14	9	3

The ordinary hinge thresholds required by the distribution of M_q are precisely

$$(B.1) \quad \left\{ \frac{12}{11}, \frac{8}{7}, \frac{6}{5}, \frac{4}{3}, \frac{3}{2}, \frac{8}{5}, \frac{12}{7}, 2, \frac{12}{5}, \frac{8}{3}, 3, 4, 6, 8, 12 \right\}.$$

They equal $\{\tau_q/m : 1 \leq m < \tau_q, q > 5\}$ with repetitions removed. Together with $0 \leq u < 12$ and $0 \leq v < 9$, these specify the finite ordinary geometric queries. Positive-depth regions use $u \geq 1$ or $v \geq 1$ as appropriate; the zero atoms are handled separately.

The checker enumerates a configuration as

$$(a, b, c, r, j, d, k, i, z).$$

The entries are respectively the 15-residue, the 27-residue, the 25-column, the 45-residue, the pure-5 budget column, the 75-projection modulo 3 and modulo 5, its overlap-vertex flag, and the pure-3 budget row. The sentinel $r = -1$ means that the 45-hole has not yet been retained; $d = 0$ means that the 75-hole is ignored; $z = -1$ means that deeper pure-3 deletions are ignored. These sentinels enlarge an upper-cost measure; they do not describe missing moduli in the completed family.

The main routines correspond directly to the proof:

Routine	Mathematical task
cells, weights	Core and four comparison regions of Sections 5 and 8.
depth	Geometric mass and first moments, including the omitted infinite tails.
geometry	Integer maximum in Appendix A, with each of the stated offsets and coefficient vectors.
envelope	Finite anchor sum plus its outward linear remainder.
distributions	Exact distribution of M_q below 32 and its unrestricted first moment.
ordinary	Six ordinary stage-loss bounds.
spatial	Charged current-7 loss and the later first-hit subtraction.
fixed165	The seven fixed-label expectations at stage 11.
reserve	The matching lower reserve for the same configuration and budget vertex.
main	Domain enumeration, the hierarchy of certificates, terminal symmetries, and exact integer assertions.

The file `integer_certificate.json` records the phase, configuration, reserve lower integer, and loss upper integer for every certified case. Its phase names `basic` and `basic75` denote, respectively, the 16 basic entries without the additional 75-credit and the four using Lemma 5.1. The closing data retain the exact rational available budget and the rational upper cost for each of the 28 final expectations. A short replay program checks these recorded integer inequalities. Such a replay does not regenerate the geometric bounds: the command `python3 checks/verify.py -fresh` performs that stronger check from the definitions.

APPENDIX C. A UNIFORM UNCOVERED-DENSITY BOUND

The proof constructs a measure for every finite distinct family on the reference support, whether or not the family covers. Its surviving mass and bounded density give a quantitative lower bound for the uncovered proportion. We first record the averaging argument needed to transfer a uniform uncovered proportion between prime supports. For a periodic set $U \subseteq \mathbb{Z}$, its natural density is

$$d(U) = \lim_{T \rightarrow \infty} \frac{|U \cap \{0, 1, \dots, T-1\}|}{T}, \quad T \in \mathbb{N}.$$

Periodicity guarantees that this limit exists and equals the uncovered proportion in any period; it also agrees with the symmetric density on $\mathbb{Z}$.

Lemma C.1 (Averaged prime replacement). *Let $p_1 < \dots < p_r$ and $q_1 < \dots < q_r$ be primes with $p_i \leq q_i$, and let $A_1, \dots, A_r$ be nonnegative integers. Suppose that every family of prefix cylinders with distinct nonzero complete exponent vectors on $X_p = \prod_{i=1}^r \mathbb{Z}/p_i^{A_i} \mathbb{Z}$ leaves a proportion at least c uncovered. Then the same is true on $X_q = \prod_{i=1}^r \mathbb{Z}/q_i^{A_i} \mathbb{Z}$.*

Proof. Identify each prime-power residue with its digit word. For each coordinate i and digit position ℓ , choose an independent uniform shift $b_{i,\ell} \in \mathbb{Z}/q_i \mathbb{Z}$. Send a p_i -digit u to $u + b_{i,\ell} \pmod{q_i}$. Applying these digit maps gives a random prefix-preserving injection F_i , and their product gives $F : X_p \rightarrow X_q$. The inverse image of a prefix cylinder is either empty or a prefix cylinder with the same complete exponent vector. Thus every realization pulls a given family on X_q back, after discarding empty inverse images, to a family satisfying the hypothesis on X_p .

For every fixed $x \in X_p$, the random variable $F(x)$ is uniform on X_q , since all output digits are independent and uniform. If $U \subseteq X_q$ is the uncovered set, then

$$\mathbb{E}_F \frac{|F^{-1}(U)|}{|X_p|} = \frac{1}{|X_p|} \sum_{x \in X_p} \mathbb{P}_F(F(x) \in U) = \frac{|U|}{|X_q|}.$$

For every realization the fraction inside the expectation is at least c , so the uncovered proportion on X_q is also at least c . $\square$

Corollary C.2 (Uniform uncovered density). *Every finite family of residue classes with pairwise distinct odd moduli greater than one, whose least common multiple has at most eight distinct prime divisors, leaves a set of integers of natural density at least*

$$\frac{1}{1,002,375}.$$

Proof. First suppose the family is supported on $\mathcal{P}$. Completion, the kernel construction, and the estimate leading to (10.4) apply without a covering assumption; that assumption was used only for the concluding contradiction. Write $U \subseteq X$ for the original uncovered set, after any Haar-preserving tree permutations used to normalize the anchors. The final measure is supported in U , since completion preserves the original covered subset and every completed forbidden class is removed.

Initially the density with respect to H is $\mathbf{1}_A$. Each subsequent kernel is at most C_q , and each deletion only decreases the density. Consequently, for every measurable $E \subseteq X$,

$$\nu_{\text{final}}(E) \leq M H(E), \quad M = \prod_{q \in \{7, 11, 13, 17, 19, 23\}} C_q.$$

The caps in (3.4) give

$$M = \frac{3}{2} \cdot \frac{5}{3} \cdot \frac{3}{2} \cdot 2 \cdot \frac{9}{5} \cdot \frac{11}{5} = \frac{297}{10}.$$

The global surplus four, rather than the terminal surplus eleven, gives (10.4). Hence

$$H(U) \geq \frac{10}{297} \nu_{\text{final}}(X) \geq \frac{10}{297 \cdot 33750} = \frac{1}{1,002,375}.$$

The original family is finite, so this Haar measure is its uncovered proportion in a finite period and hence its uncovered natural density. The normalizing tree permutations preserve that proportion. Unused reference coordinates can be added without changing it.

For arbitrary odd prime support $q_1 < \cdots < q_r$, with $r \leq 8$, take p_i to be the first r odd primes and A_i to be the largest exponent of q_i in the finite period. The reference-support bound and Lemma C.1 give the same constant. The empty family leaves every integer uncovered. $\square$

ACKNOWLEDGMENTS

The preparation of this manuscript and its computational companion benefited from assistance by AI systems developed by OpenAI and Anthropic. Responsibility for the mathematical claims, implementation, and presentation rests with the author.

REFERENCES

- [1] P. Balister, *Erdős covering systems*, in Surveys in Combinatorics 2024 (F. Fischer and R. Johnson, eds.), London Math. Soc. Lecture Note Ser. **493**, Cambridge Univ. Press, 2024, 31–54. doi:10.1017/9781009490559.003.
- [2] 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. doi:10.1007/s00222-021-01087-5. arXiv:1811.03547.
- [3] 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.
- [4] M. A. Berger, A. Felzenbaum and A. S. Fraenkel, *Necessary condition for the existence of an incongruent covering system with odd moduli*, Acta Arith. **45** (1986), 375–379. doi:10.4064/aa-45-4-375-379.
- [5] M. A. Berger, A. Felzenbaum and A. S. Fraenkel, *Necessary condition for the existence of an incongruent covering system with odd moduli II*, Acta Arith. **48** (1987), 73–79. doi:10.4064/aa-48-1-73-79.
- [6] T. F. Bloom, *Erdős Problem #7*, erdosproblems.com/7 (consulted 9 September 2026).
- [7] 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.
- [8] 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.
- [9] 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.
- [10] 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.
- [11] 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.

- [12] 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.
- [13] 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:2607.25628v1.
- [14] L. de Moura and S. Ullrich, *The Lean 4 theorem prover and programming language*, in Automated Deduction–CADE 28, Lecture Notes in Comput. Sci. **12699**, Springer, 2021, 625–635. doi:10.1007/978-3-030-79876-5_37.
- [15] G. Puccetti and R. Wang, *Extremal dependence concepts*, Statist. Sci. **30** (2015), 485–517. doi:10.1214/15-STS525. Author-maintained text, with corrections through May 2026.
- [16] M. Schroeder, *Seven prime divisors in odd distinct covering systems*, author preprint, 1 September 2026. Author’s preprint.
- [17] M. Schroeder, *Eight prime divisors in odd distinct covering systems*, author preprint, 6 September 2026. Author’s preprint.
- [18] R. J. Simpson and D. Zeilberger, *Necessary conditions for distinct covering systems with square-free moduli*, Acta Arith. **59** (1991), 59–70. doi:10.4064/aa-59-1-59-70.

INDEPENDENT RESEARCHER