

An elementary quantitative proof of Erdős Problem 126

Shenghui Bao

Tashkent Institute of Irrigation and Agricultural
Mechanization Engineers - National Research University

Corresponding author: Shenghui Bao

bsh@nooslab.cn

ORCID: 0009-0008-5462-6988

Abstract

Let A be a finite set of distinct positive integers, and let $S(A)$ be the set of primes dividing a sum of two distinct elements of A . We give an elementary quantitative proof of $|A| \leq 3|S(A)|^2$ for $|A| \geq 2$, which implies the superlogarithmic growth asked for in Erdős Problem 126. The construction originates in an AI-generated formal proof. Prime colours orient negation orbits modulo prime powers; their supports form laminar families, and logarithmic weights relate sums to differences. A positive semidefinite correction accounts for self-opposite residues and transfers conditional negative type from the kernel $\log(x+y)$ to an arithmetic kernel. Applying the signed laminar-kernel theorem from a companion paper then gives the cardinality bound. We provide the arithmetic and analytic details, explain how they fit together, and record the provenance of the formal construction in a short appendix.

Keywords: Pairwise sums; prime support; Erdős Problem 126; prime-power congruences; laminar families; conditionally negative kernels.

1 Introduction

How large can a set of positive integers be if every sum of two distinct elements has its prime factors in a fixed, small collection? The question places a multiplicative restriction on an additive operation. Prescribing only the number of primes asks more than prescribing the primes themselves: the primes can be arbitrarily large, and the integers can contain arbitrarily high powers. A bound in terms of the number of primes must therefore be independent of both kinds of arithmetic height.

For a finite set $A \subset \mathbb{N}$, put

$$S(A) = \{p : p \text{ is prime and } p \mid a + b \text{ for some distinct } a, b \in A\}.$$

Throughout, $\mathbb{N} = \{0, 1, 2, \dots\}$; when positivity is required, we state it explicitly. The extremal function

in Erdős Problem 126 is

$$f(n) = \min_{\substack{A \subset \mathbb{N} \\ |A|=n}} |S(A)|.$$

The minimum exists because the collection is nonempty and its values are nonnegative integers. The product over ordered distinct pairs used in the formal statement has the same prime support as the product over unordered pairs. Thus the order and multiplicities of the factors do not affect f .

In their first joint paper, Erdős and Turán [6] proved an elementary cardinality bound that is exponential in the number of supporting primes. In present notation, their estimate for positive sets is

$$|A| < 3 \cdot 2^{r-1}, \quad r = |S(A)|,$$

in the nontrivial range. It yields a logarithmic lower bound for the extremal prime-support function. For the upper bound, consider the interval $\{1, \dots, n\}$. Every relevant sum is at most $2n - 1$, so the prime number theorem gives $f(n) \ll n/\log n$. Erdős later asked whether $f(n)$ eventually exceeds every fixed multiple of $\log n$, that is, whether

$$\frac{f(n)}{\log n} \longrightarrow \infty. \quad (1)$$

Subsequent work developed both arithmetic and combinatorial approaches. Győry, Stewart and Tijdeman [9] treated prime factors of sums from two sets using methods involving S -unit equations. For $|A| \geq |B| \geq 2$, they obtained an effective logarithmic lower bound in $|A|$ for the number of prime factors of the full product of the sums $a + b$. Their result concerns a two-set product, which differs from the off-diagonal same-set product studied here. Wu [11] introduced p -graphs and improved the Erdős–Turán bound for positive sets to

$$|A| \leq 2^r.$$

He also proved that this estimate is sharp when $r = 2$. These results distinguish the quantitative problem from the simpler assertion that a set with prescribed supporting primes is finite.

The construction developed below originates in a formal resolution produced by GPT-6 Astra in the FrontierMath Erdős benchmark [1, 2]. That argument gives a quadratic cardinality bound. We present it as a conventional mathematical proof, with the arithmetic construction and analytic mechanism made explicit. We do not claim independent discovery of the construction. Its abstract combinatorial step is isolated and proved in the companion paper [4]; here we establish that theorem’s hypotheses from prime-power divisibility. Appendix A identifies the source and distinguishes the selected alternate resolution from the repository’s primary resolution.

Theorem 1.1 (Quadratic prime-support bound). *Let A be a finite set of distinct positive integers with $n = |A| \geq 2$, and let $r = |S(A)|$. Then*

$$n \leq 3r^2. \quad (2)$$

The estimate is independent of the values of the elements and the supporting primes. Removing a possible zero gives the uniform statement

$$|A| \leq 3|S(A)|^2 + 2 \quad (A \subset \mathbb{N} \text{ finite}). \quad (3)$$

For $n \geq 2$ this implies

$$f(n) \geq \sqrt{\frac{n-2}{3}}, \quad (4)$$

and consequently (1). Section 6 gives the zero reduction and the deduction for f . The restriction $n \geq 2$ in Theorem 1.1 is essential because a singleton has empty prime support. The same restriction ensures that the square root in (4) is defined over the reals.

The proof combines valuation counts, refinement of residue classes, and signs. The exponent of a prime in a sum counts the prime-power levels at which the two residues are opposite. Negation orbits at successive powers refine one another, so their supports form a laminar family. A single binary colour for each prime consistently distinguishes the two sides of every represented non-self-opposite orbit. Logarithmic weights retain the valuation counts, laminarity records refinement, and a signed kernel encodes the colours. Together with conditional negative type, these are the ingredients required by the abstract theorem.

Two familiar reductions help explain why this structure is useful. Choose distinct anchors $a_1, a_2 \in A$ and let $x \in A \setminus \{a_1, a_2\}$. If G is the multiplicative group generated by $S(A)$, then $u = a_1 + x$ and $v = a_2 + x$ belong to G and satisfy

$$\frac{u}{a_1 - a_2} - \frac{v}{a_1 - a_2} = 1.$$

Both terms are nonzero, so these are nondegenerate solutions of a fixed two-variable equation in G^2 , which has rank $2r$. Different x give different solutions. Uniform results on equations in finite-rank multiplicative groups, such as the theorem of Evertse, Schlickewei and Schmidt [7], therefore bound $|A|$ in terms of r . Applied in this way, their general estimate is exponential in r and gives a logarithmic prime-support lower bound. This reduction, however, no longer records the relations between two freely varying elements of A .

A matrix formulation retains more of the pairwise information. For this comparison, consider the sum matrix $L_{ij} = a_i + a_j$. It has rank at most two, since

$$L = a\mathbf{1}^\top + \mathbf{1}a^\top,$$

and its off-diagonal entries lie in G . This connects the problem to the results of Alon and Solymosi [3] on matrices with entries in a multiplicative group. The argument here uses further properties of this matrix: positivity, symmetry, and the simultaneous refinement of congruences modulo all powers of every supporting prime. This comparison identifies the additional structure used below; it does not rule out refinements of the S -unit or low-rank approaches.

Section 2 states the abstract input, and Section 3 constructs the prime colours and proves laminarity. Section 4 derives the exact sum identity and the difference inequality, including the correction for self-opposite residues. Section 5 proves the conditional negative type of $\log(x + y)$ and transfers it to the arithmetic cross kernel. Section 6 applies the companion theorem without repeating its matching and energy arguments. Section 7 illustrates the construction with explicit matrices.

2 Notation and the abstract input

We also use the inverse extremal function

$$M(r) = \sup\{|A| : A \subset \mathbb{N} \text{ is finite and } |S(A)| \leq r\}.$$

Equivalently, one may choose a set of exactly r primes and maximize the cardinality of a set whose distinct pairwise sums are supported on those primes, then take the supremum over the choice of primes. A smaller supporting set can always be enlarged with unused primes. Taking subsets shows that $f(n) \leq r$ exactly when $n \leq M(r)$. The classical exponential estimate, after removal of a possible zero, ensures that $M(r)$ is finite. Inverting this monotone relation shows that (1) is equivalent to $\log M(r) = o(r)$, or $M(r) = \exp(o(r))$. The uniform bound proved here gives the stronger estimate $M(r) \leq 3r^2 + 2$.

Let V be a finite index set of size $n \geq 2$, and let

$$a : V \longrightarrow \mathbb{N}_{>0}$$

be injective. Write $a_i = a(i)$, $A = a(V)$, $S = S(A)$, and $r = |S|$. If p is prime and m is a positive integer, $v_p(m)$ denotes its p -adic valuation. Unique factorization gives the identity

$$\log m = \sum_{p|m} v_p(m) \log p. \quad (5)$$

The empty sum for $m = 1$ is zero. Every logarithm is the natural logarithm.

Choose an integer $K \geq 1$ such that

$$v_p(a_i + a_j) \leq K \quad (p \in S, i, j \in V). \quad (6)$$

There are finitely many pairs and supporting primes, so a common cutoff exists. We include diagonal pairs for convenience, although S is defined using distinct pairs; this does not assert that every prime divisor of $2a_i$ belongs to S . The cutoff makes the construction finite. Its size does not enter the final estimate, which is independent of the largest valuation.

For a symmetric real kernel $H = (H_{ij})_{i,j \in V}$ define

$$q_H(z) = \sum_{i,j \in V} z_i z_j H_{ij} \quad (z \in \mathbb{R}^V).$$

The kernel is *positive semidefinite* if $q_H(z) \geq 0$ for every z . It is *conditionally negative semidefinite*, abbreviated *CND*, if $q_H(z) \leq 0$ whenever $\sum_i z_i = 0$. Only finite-dimensional kernels are needed here; general references are [5, 10].

A family of subsets of V is *laminar* if any two members are disjoint or one contains the other. We allow a finite collection of labelled nodes with repeated underlying subsets. Laminarity concerns the subsets, while nodes with the same support may carry different weights. This convention is necessary because equal supports at different prime-power levels encode separate units of valuation.

The following result from the companion paper is our combinatorial input.

Theorem 2.1 (Signed laminar-kernel theorem, [4, Theorem 2.1]). *Let P be a finite label set with $|P| = r$. For each $p \in P$ let $\mathcal{F}_p$ be a finite collection of labelled subsets $U_{p,t} \subseteq V$, each of cardinality at least two,*

which is laminar for fixed p . Assign weights $w_{p,t} \geq 0$ and signs $\sigma_p : V \rightarrow \{-1, 1\}$. Set

$$M_{ij}^p = \sum_{t \in \mathcal{F}_p} w_{p,t} \mathbf{1}_{U_{p,t}}(i) \mathbf{1}_{U_{p,t}}(j),$$

and define

$$C_{ij} = \sum_{\substack{p \in P \\ \sigma_p(i) \neq \sigma_p(j)}} M_{ij}^p, \quad R_{ij} = \sum_{\substack{p \in P \\ \sigma_p(i) = \sigma_p(j)}} M_{ij}^p.$$

If $n = |V| \geq 2$, C is CND , and $R_{ij} < C_{ij}$ for all $i \neq j$, then $n \leq 3r^2$.

The companion paper proves this theorem using a signed Gram estimate and a matching estimate for laminar families. We recall only its statement; the arithmetic construction below can be checked without those proofs. We must construct the families, weights and signs, then establish conditional negative type and strict separation. These are the points at which the arithmetic argument meets its combinatorial input.

Strict separation requires more than a positive semidefinite family kernel. Membership vectors readily provide such a kernel, but positive semidefiniteness alone does not force a cardinality bound. After the signs split each local kernel, the cross-colour contribution must strictly exceed the same-colour contribution for every distinct pair. We will obtain this comparison from two logarithmic expressions with a common correction term.

3 Prime colours and mixed negation orbits

3.1 A colour that works at every prime-power level

For a positive integer a and a prime p , write

$$a = p^{v_p(a)} u_p(a), \quad p \nmid u_p(a).$$

The integer $u_p(a)$ is the unit part of a . For odd p , let $\bar{u}_p(a) \in \{1, \dots, p-1\}$ be its least positive residue modulo p , and set

$$\sigma_p(a) = \begin{cases} +1, & 2\bar{u}_p(a) < p, \\ -1, & 2\bar{u}_p(a) > p. \end{cases} \quad (7)$$

There is no equality case. At $p = 2$, define instead

$$\sigma_2(a) = \begin{cases} +1, & u_2(a) \equiv 1 \pmod{4}, \\ -1, & u_2(a) \equiv 3 \pmod{4}. \end{cases} \quad (8)$$

The colour depends only on the integer and the prime, so the same choice applies at every level. For odd primes, the leading nonzero digit determines which half of the nonzero residue system contains the unit residue. At the prime two, all odd parts are congruent modulo two; reduction modulo four supplies the distinction we need.

Lemma 3.1 (Prime-colour separation). *Let p be prime, $m \geq 1$, and $a, b > 0$. If*

$$p^m \mid a + b \quad \text{and} \quad p^m \nmid 2a,$$

then $\sigma_p(a) \neq \sigma_p(b)$.

Proof. Suppose first that p is odd, and put $\alpha = v_p(a)$. The second hypothesis gives $\alpha < m$. If $v_p(b) \neq \alpha$, the valuation of $a + b$ is the smaller of the two valuations and hence is less than m , a contradiction. Thus $v_p(b) = \alpha$. Write $a = p^\alpha u$ and $b = p^\alpha v$ with $p \nmid uv$. Cancelling p^α gives $u + v \equiv 0 \pmod{p}$. The least positive residues of u and v add to p , so exactly one is smaller than $p/2$. Their colours are therefore opposite.

Now let $p = 2$ and put $\alpha = v_2(a)$. The hypothesis $2^m \nmid 2a$ gives $m > \alpha + 1$. Again, the valuation of b must equal α . Write $a = 2^\alpha u$ and $b = 2^\alpha v$ with u, v odd. Divisibility of the sum gives $2^{m-\alpha} \mid u + v$. Since $m - \alpha \geq 2$, the odd parts are opposite modulo four: one is congruent to 1 and the other to 3. This proves the assertion. $\square$

The nondivisibility condition excludes exactly the obstruction to orienting a negation orbit. Modulo p^m , a residue x is its own negative precisely when $2x = 0$. The lemma applies when the residues are opposite and the first is not self-opposite. Since $a \equiv -b \pmod{p^m}$, either both residues are self-opposite or neither is. The statement is therefore symmetric in a and b .

3.2 The represented orbits

For $p \in S$ and $1 \leq m \leq K$, write

$$\rho_{p,m}(i) = a_i \pmod{p^m} \in \mathbb{Z}/p^m\mathbb{Z}.$$

The involution $x \mapsto -x$ partitions this residue group into one-element and two-element orbits. A two-element orbit $\{x, -x\}$ is called *mixed* if both fibres $\rho_{p,m}^{-1}(x)$ and $\rho_{p,m}^{-1}(-x)$ are nonempty. Its support is

$$U_{p,m,\{x,-x\}} = \rho_{p,m}^{-1}(x) \cup \rho_{p,m}^{-1}(-x). \quad (9)$$

Each unordered orbit is included once, regardless of the choice of representative. A node records both its level m and its orbit. Every retained support has at least two vertices, one from each fibre.

Mixedness ensures the size condition in Theorem 2.1 and supplies a witness on the opposite side of each fibre. The prime-colour lemma can then be applied to every vertex of the fibre. If only one side were present, that residue class could contribute family mass without any corresponding opposite-residue pair. We therefore omit such classes from the laminar families.

Lemma 3.2 (Laminarity). *For a fixed prime p , the supports of the mixed orbits over all levels $1 \leq m \leq K$ form a laminar family. Repeated supports with distinct level labels are allowed.*

Proof. At a fixed level, distinct negation orbits are disjoint. Now suppose $m \leq \ell$, let U be the support of a mixed orbit modulo p^ℓ , and let W be the support of a mixed orbit modulo p^m . If $U \cap W$ is empty, there is nothing to prove. Otherwise choose $k \in U \cap W$. For $i \in U$, we have $\rho_{p,\ell}(i) = \pm \rho_{p,\ell}(k)$. Reducing modulo p^m gives $\rho_{p,m}(i) = \pm \rho_{p,m}(k)$. Since $k \in W$, both possibilities belong to the lower-level orbit defining W . Hence $U \subseteq W$. $\square$

The proof uses only the compatibility of reduction with negation. A higher-level orbit need not have a retained mixed ancestor at every lower level: it may reduce to a self-opposite residue, which is omitted

from the mixed family. Laminarity concerns two retained nodes when they intersect, and the reduction argument covers exactly that case.

The supports can be pictured as a forest ordered by containment. Higher-level orbits may split a lower-level support into disjoint groups, but cannot merge groups that were already disjoint. Some levels contribute no retained node; successive levels may also give the same underlying support. These repetitions must remain in the weighted collection. Whenever a pair remains opposite through several prime-power levels, each level contributes one unit of valuation, whether or not the support has changed.

The signs orient every retained node consistently. Opposite fibres have opposite colours by Lemma 3.1. If i, j lie in the same fibre, choose k in the other fibre. Both have colour opposite to k , so their colours agree. Thus, within a mixed support, two vertices have the same colour exactly when they lie in the same fibre, and opposite colours exactly when they lie in opposite fibres. The single function σ_p on all vertices therefore produces compatible bipartitions of the nodes, rather than an arbitrary bipartition of each node separately.

3.3 Family kernels

Let $\mathcal{F}_p$ be the finite collection of level-tagged mixed nodes, with support $U_{p,t}$ for $t \in \mathcal{F}_p$. Give each node for the prime p the weight $\log p$, and define

$$M_{ij}^p = \log p \sum_{t \in \mathcal{F}_p} \mathbf{1}_{U_{p,t}}(i) \mathbf{1}_{U_{p,t}}(j). \quad (10)$$

The entries are nonnegative, and for every real vector z ,

$$q_{M^p}(z) = \log p \sum_{t \in \mathcal{F}_p} \left(\sum_{i \in U_{p,t}} z_i \right)^2 \geq 0. \quad (11)$$

Each support's membership vector contributes a rank-one positive semidefinite kernel. Since $\log p > 0$, their weighted sum is also positive semidefinite.

Using the signs $\sigma_p(a_i)$, split these local kernels into

$$C_{ij} = \sum_{\substack{p \in S \\ \sigma_p(a_i) \neq \sigma_p(a_j)}} M_{ij}^p, \quad (12)$$

$$R_{ij} = \sum_{\substack{p \in S \\ \sigma_p(a_i) = \sigma_p(a_j)}} M_{ij}^p. \quad (13)$$

Both kernels are symmetric and entrywise nonnegative. The cross-colour kernel has zero diagonal, while the same-colour kernel carries all the family diagonal mass:

$$C_{ii} = 0, \quad R_{ii} = \sum_{p \in S} M_{ii}^p.$$

These are the kernels required by Theorem 2.1. Conditional negative type and strict separation do not follow from laminarity or positive semidefiniteness alone; we establish them using the arithmetic identities below.

4 The baseline and the two logarithmic comparisons

4.1 Self-opposite residues

At a fixed level, $p^m \mid a_i + a_j$ means $\rho_{p,m}(i) = -\rho_{p,m}(j)$. For a two-element orbit, the vertices i, j witness mixedness, and the cross-colour kernel detects the event. For a self-opposite residue, both vertices lie in one fibre, so there is no two-sided orbit to retain. We restore these missing events with a correction whose quadratic form has the required sign.

Set

$$H_{p,m} = \{x \in \mathbb{Z}/p^m\mathbb{Z} : 2x = 0\},$$

and define

$$b_{ij}^{p,m} = \mathbf{1}_{\{\rho_{p,m}(i)=\rho_{p,m}(j) \in H_{p,m}\}}. \quad (14)$$

For odd p , zero is the only self-opposite residue. For $p = 2$, the self-opposite residues modulo 2^m are 0 and 2^{m-1} . The formula covers both cases. Define the global baseline by

$$B_{ij} = \sum_{p \in S} \log p \sum_{m=1}^K b_{ij}^{p,m}. \quad (15)$$

Lemma 4.1. *The kernel B is positive semidefinite.*

Proof. For each level,

$$b_{ij}^{p,m} = \sum_{x \in H_{p,m}} \mathbf{1}_{\{\rho_{p,m}(i)=x\}} \mathbf{1}_{\{\rho_{p,m}(j)=x\}},$$

so

$$q_{b^{p,m}}(z) = \sum_{x \in H_{p,m}} \left(\sum_{\rho_{p,m}(i)=x} z_i \right)^2 \geq 0.$$

Taking the nonnegative weighted sum in (15) proves the assertion. $\square$

Because the correction is a sum of same-fibre kernels, its positive semidefiniteness follows directly from a sum of squares. It is not an arbitrary matrix chosen to repair the arithmetic identity. We will subtract it from an analytic CND kernel, preserving the required inequality for zero-sum vectors.

4.2 A levelwise identity for sums

At level (p, m) , let $c_{ij}^{p,m}$ count retained mixed supports containing i, j with different p -colours. At most one orbit contains the pair, so this count is zero or one. The orientation of each mixed orbit gives

$$c_{ij}^{p,m} + b_{ij}^{p,m} = \mathbf{1}_{\{p^m \mid a_i + a_j\}}. \quad (16)$$

Indeed, opposite non-self-opposite residues contribute to $c^{p,m}$, while opposite self-opposite residues are equal and contribute to $b^{p,m}$. Conversely, either contribution implies opposite residues. The two cases are disjoint.

The identity also holds on the diagonal. A vertex has the same colour as itself, so $c_{ii}^{p,m} = 0$. The condition $p^m \mid 2a_i$ on the right says exactly that its residue is self-opposite, as detected by the baseline.

Weighting and summing (16) yields

$$C_{ij} + B_{ij} = \sum_{p \in S} \log p \sum_{m=1}^K \mathbf{1}_{\{p^m | a_i + a_j\}}. \quad (17)$$

For $i \neq j$, every prime factor of the sum lies in S , and the cutoff captures its full valuation. Equation (5) therefore gives

$$C_{ij} + B_{ij} = \log(a_i + a_j) \quad (i \neq j). \quad (18)$$

This is where the prime-support hypothesis enters; constructing the local kernels did not require it.

4.3 An inequality for differences

Let $r_{ij}^{p,m}$ count retained mixed supports containing i, j with the same p -colour. Within a mixed support, equal colours mean equal residues. Thus a contribution from either $r^{p,m}$ or the baseline forces equal residues, and the two possibilities are disjoint. It follows that

$$r_{ij}^{p,m} + b_{ij}^{p,m} \leq \mathbf{1}_{\{a_i \equiv a_j \pmod{p^m}\}}. \quad (19)$$

The inequality can be strict. Two vertices may share a non-self-opposite residue whose opposite fibre is absent from A . They are congruent, but no mixed node is retained and the baseline does not include their residue. The difference relation is therefore an inequality rather than an identity.

For $i \neq j$ put $d_{ij} = |a_i - a_j| > 0$. Summing (19) gives

$$\begin{aligned} R_{ij} + B_{ij} &\leq \sum_{p \in S} \log p \sum_{m=1}^K \mathbf{1}_{\{p^m | d_{ij}\}} \\ &\leq \sum_{p \in S} v_p(d_{ij}) \log p \leq \log d_{ij}. \end{aligned} \quad (20)$$

The difference need not be supported on S . Omitting primes outside S and powers above the cutoff can only decrease the nonnegative valuation sum. Injectivity of a ensures $d_{ij} > 0$, so its logarithm is defined and its valuations are finite.

The diagonal requires a separate, one-sided estimate. Since $C_{ii} = 0$, the diagonal case of (17) gives

$$B_{ii} = \sum_{p \in S} \log p \sum_{m=1}^K \mathbf{1}_{\{p^m | 2a_i\}} \leq \log(2a_i). \quad (21)$$

All summands are nonnegative, so omitting prime divisors of $2a_i$ outside S preserves the inequality. The resulting diagonal slack has the required sign in the CND argument.

Proposition 4.2 (Strict separation). *For all distinct $i, j \in V$, one has $R_{ij} < C_{ij}$.*

Proof. The elements are distinct and positive, hence

$$0 < |a_i - a_j| < a_i + a_j.$$

Using strict monotonicity of the logarithm, (20), and (18), we obtain

$$R_{ij} + B_{ij} \leq \log |a_i - a_j| < \log(a_i + a_j) = C_{ij} + B_{ij}.$$

Cancellation of the common baseline proves the assertion. $\square$

Strictness comes from comparing the sum and difference as positive real numbers; no estimate of an individual valuation is needed. The common baseline allows the comparison to pass to R and C . Discarding the self-opposite levels would remove that common correction and prevent this direct use of the inequality between sum and difference.

The logarithmic weights serve a precise purpose. Valuations add under multiplication; weighting each $v_p(m)$ by $\log p$ and summing over the prime divisors gives exactly $\log m$. Unit weights would instead count prime factors with multiplicity, without retaining their sizes. Such a count has no corresponding strict monotonicity under $|a_i - a_j| < a_i + a_j$. The weights are therefore chosen to produce the logarithmic comparison used in the proof.

5 Conditional negative type and the Möbius substitution

With strict separation established, it remains to show that C is CND. We do not need to establish negative type for its individual prime components. Instead, we use the logarithmic expression obtained by combining all the prime-power contributions.

Lemma 5.1 (Additive logarithm kernel). *For positive real numbers $x_i, i \in V$, the kernel $L_{ij} = \log(x_i + x_j)$ is conditionally negative semidefinite.*

Proof. Put

$$t_i = \frac{x_i - 1}{x_i + 1}.$$

Positivity gives $-1 < t_i < 1$. A direct calculation shows

$$1 - t_i t_j = \frac{2(x_i + x_j)}{(x_i + 1)(x_j + 1)}.$$

Every factor is positive, so taking logarithms gives

$$\log(x_i + x_j) = \log(x_i + 1) + \log(x_j + 1) - \log 2 + \log(1 - t_i t_j). \quad (22)$$

Let $z \in \mathbb{R}^V$ satisfy $\sum_i z_i = 0$. The first three terms on the right contribute zero to $q_L(z)$. For example,

$$\sum_{i,j} z_i z_j \log(x_i + 1) = \left(\sum_i z_i \log(x_i + 1) \right) \left(\sum_j z_j \right) = 0.$$

The column term and constant term vanish in the same way. Thus

$$q_L(z) = \sum_{i,j} z_i z_j \log(1 - t_i t_j).$$

Use the absolutely convergent expansion

$$\log(1 - u) = - \sum_{\ell=1}^{\infty} \frac{u^\ell}{\ell} \quad (|u| < 1).$$

Since there are finitely many t_i , we have $\rho = \max_i |t_i| < 1$. The terms are bounded by a fixed multiple of $\rho^{2\ell}/\ell$, so we may interchange the finite sum over the indices with the series. This gives

$$q_L(z) = - \sum_{\ell=1}^{\infty} \frac{1}{\ell} \left(\sum_i z_i t_i^\ell \right)^2 \leq 0. \quad (23)$$

□

The substitution maps the positive half-line into $(-1, 1)$, where a power series separates the two variables. The zero-sum condition removes every term depending on only one index. What remains is $\log(1 - t_i t_j)$, whose coefficients produce negative squares. This explicit decomposition is an elementary instance of the classical theory of negative-type kernels [5, 10]; no general representation theorem is required.

The scale 1 in the Möbius substitution is arbitrary. Replacing it by any positive s gives $t_i = (x_i - s)/(x_i + s)$. In (22), the row and column terms then use $x_i + s$ and $x_j + s$ in place of $x_i + 1$ and $x_j + 1$, and the constant $\log 2$ becomes $\log(2s)$. The same zero-sum cancellation applies. The analytic statement therefore has no preferred arithmetic scale: the substitution places the finitely many positive inputs inside the interval of absolute convergence.

We now specialize to $x_i = a_i$. Define

$$D_{ij} = \begin{cases} \log(2a_i) - B_{ii}, & i = j, \\ 0, & i \neq j. \end{cases}$$

By (21), D is a nonnegative diagonal matrix, hence positive semidefinite. Together with $C_{ii} = 0$, equation (18) gives the full matrix identity

$$C = L - B - D, \quad L_{ij} = \log(a_i + a_j). \quad (24)$$

The explicit diagonal correction allows us to use the off-diagonal identity in the full quadratic form.

Proposition 5.2. *The cross kernel C is symmetric, has zero diagonal, and is CND .*

Proof. Symmetry and the zero diagonal follow from (12). If $\sum_i z_i = 0$, then Lemma 5.1, Lemma 4.1, and positivity of D give

$$q_C(z) = q_L(z) - q_B(z) - q_D(z) \leq 0.$$

□

The two corrections account for different omissions. The baseline restores self-opposite events off the diagonal. The diagonal correction accounts for prime divisors of $2a_i$ outside $S(A)$, where no exact factorization identity is available. Subtracting these positive semidefinite corrections preserves the required negative inequality. Entrywise inequalities would not suffice off the diagonal, since $z_i z_j$ can have either

sign. We use exact equality there and allow an inequality only on the diagonal, where the coefficients are squares.

6 Applying the combinatorial theorem

Proof of Theorem 1.1. Apply Theorem 2.1 with $P = S(A)$. For each prime p , use the level-tagged mixed supports from Section 3, give each node weight $\log p$, and take the sign $i \mapsto \sigma_p(a_i)$. Mixedness ensures that every support has at least two elements. The weights are positive, and Lemma 3.2 gives laminarity. The family kernels and their same-colour and cross-colour parts are exactly (10)–(13). Proposition 4.2 gives $R_{ij} < C_{ij}$ for distinct indices, and Proposition 5.2 gives conditional negative type. All hypotheses hold, so $n \leq 3|S(A)|^2$. $\square$

This is the only application of the companion theorem. The arithmetic work is complete at this point: the abstract result does not use the fact that the labels are primes or the particular formula for their weights. Conversely, the arithmetic construction does not require a choice of the matching used to prove that result. This separation allows the theorem about signed laminar families to be stated and proved independently of the number-theoretic problem.

Corollary 6.1 (Uniform natural-number bound). *Every finite $A \subset \mathbb{N}$ satisfies (3).*

Proof. Let $A^+ = A \setminus \{0\}$ and put $r = |S(A)|$. If $|A^+| \leq 1$, then $|A| \leq 2 \leq 3r^2 + 2$. Otherwise Theorem 1.1 applies to A^+ , and $S(A^+) \subseteq S(A)$ gives

$$|A| \leq |A^+| + 1 \leq 3|S(A^+)|^2 + 1 \leq 3r^2 + 1.$$

This is stronger than the asserted uniform inequality in that case. $\square$

At most one zero is removed. The allowance of two in the uniform formula covers both this removal and the small exceptional sets; it does not count two deleted vertices. For example, $A = \{0, 1\}$ has cardinality two and no supporting prime, so an additive allowance is necessary. A positive singleton is excluded from Theorem 1.1 by the hypothesis $|A| \geq 2$.

Corollary 6.2. *For the function f defined in the introduction and all $n \geq 2$, inequality (4) holds. In particular, $f(n)/\log n \rightarrow \infty$ as $n \rightarrow \infty$.*

Proof. For every n -element set $A \subset \mathbb{N}$, Corollary 6.1 gives $n - 2 \leq 3|S(A)|^2$. Since $n \geq 2$, we may take nonnegative square roots. Taking the minimum over A then gives (4). Finally, $\sqrt{(n-2)/3}/\log n$ tends to infinity, so $f(n)/\log n$ does as well. $\square$

If $f_+(n)$ is defined using only positive sets, the same argument gives $f_+(n) \geq \sqrt{n/3}$ for $n \geq 2$. The interval example still gives $f_+(n) \ll n/\log n$, leaving a substantial gap between the square-root lower bound and the interval upper bound. The proof does not establish optimality of the exponent $1/2$ for prime support. Sharpness in the more general signed laminar framework is discussed in [4]; results about that framework need not settle the arithmetic extremal problem.

Corollary 6.3 (A refinement from fractional matching). *Under the hypotheses of Theorem 1.1, one also has*

$$|A| \leq \frac{5}{2}|S(A)|^2.$$

Proof. Apply the refined signed laminar-kernel bound in [4, Proposition 5.2] to the same families, weights and signs. The hypotheses verified above are unchanged. $\square$

The refinement comes from the companion paper's analysis of the abstract constant. It leaves the arithmetic construction unchanged, and we do not repeat its fractional-matching proof. We retain $3r^2$ as the principal estimate to present the quantitative route in the originating source. The refined constant improves the resulting lower bound without changing its square-root scale.

7 A three-element example

Take $A = \{1, 2, 3\}$, with vertices indexed in that order. The off-diagonal sums are 3, 4, 5, so $S = \{2, 3, 5\}$. The cutoff $K = 2$ also covers the valuations in the diagonal sums 2, 4, 6.

For $p = 2$, the signs of 1, 2, 3 are +, +, −, respectively, since their odd parts are 1, 1, 3. At level one, all residues modulo two are self-opposite, so there is no mixed node. At level two, the orbit $\{1, 3\}$ is mixed and its support consists of the first and third vertices. The residue of 2 is self-opposite modulo four.

For $p = 3$, the mixed orbit at level one is $\{1, 2\}$, supported by the first and second vertices; the vertex 3 has zero residue. There is no mixed orbit modulo nine. For $p = 5$, the level-one orbit $\{2, 3\}$ is supported by the second and third vertices, and there is no mixed orbit modulo twenty-five. These are the only retained nodes. Each consists of a pair whose vertices have opposite colours for the corresponding prime.

Writing $\ell_p = \log p$, the family kernels therefore give

$$C = \begin{pmatrix} 0 & \ell_3 & \ell_2 \\ \ell_3 & 0 & \ell_5 \\ \ell_2 & \ell_5 & 0 \end{pmatrix}, \quad R = \begin{pmatrix} \ell_2 + \ell_3 & 0 & 0 \\ 0 & \ell_3 + \ell_5 & 0 \\ 0 & 0 & \ell_2 + \ell_5 \end{pmatrix}.$$

The baseline is

$$B = \begin{pmatrix} \ell_2 & 0 & \ell_2 \\ 0 & 2\ell_2 & 0 \\ \ell_2 & 0 & \ell_2 + \ell_3 \end{pmatrix}.$$

The nonzero off-diagonal baseline entry records that 1 and 3 have the same self-opposite residue modulo two. Their sum is divisible by two, but this level contributes no mixed orbit. At the next level, their residues are opposite and non-self-opposite, so the cross kernel contributes a second ℓ_2 . Thus $(C + B)_{13} = 2\ell_2 = \log 4$, as required.

For the other two pairs, the baseline vanishes and the cross entries are $\log 3$ and $\log 5$. The diagonal entries of the baseline are exactly $\log 2$, $\log 4$, $\log 6$, so $D = 0$ in this example. Hence $C + B$ is the full kernel ($\log(a_i + a_j)$). Strict separation is immediate: every off-diagonal entry of R is zero, while every off-diagonal entry of C is positive.

The baseline need not be a small error term: here it contains half the logarithmic mass of the sum $1 + 3$. The argument uses its exact arithmetic meaning and positive semidefiniteness, without requiring a bound on its size. For larger sets, mixed supports may contain many vertices and recur at several levels. The identities still hold because they were proved separately for each prime-power event.

Funding

This research received no specific funding.

Competing interests

The author declares no relevant financial or nonfinancial competing interests.

Data and code availability

No empirical dataset was generated or analysed for this theoretical study. The public formal source and its fixed revision are identified in Appendix A and reference [2]. The abstract combinatorial input is given in the companion manuscript [4]. The scope of source inspection and formal verification is stated in Appendix A.

Ethics statement

This work is a mathematical study and involves no human participants, animals, or personal research data. Ethics approval and participant consent do not apply.

Generative AI disclosure

OpenAI’s Codex assistant was used in preparing this manuscript, including organising the exposition, assisting with mathematical reasoning and checks, checking references and source provenance, drafting and revising English prose, and preparing the LaTeX files. AI assistance was used during September and October 2026. This assistance is distinct from the AI-generated formal proof on which the arithmetic construction is based; that source is attributed in the text and Appendix A. Responsibility for the manuscript’s accuracy, originality, interpretation, and attribution rests with the author.

References

- [1] Adamczewski T, Bloom TF. *FrontierMath Erdős*. 2026. arXiv:2609.25050.
<https://arxiv.org/abs/2609.25050>.
- [2] Adamczewski T (repository maintainer). *Erdős problem #126: proof*. Public Lean source repository; 2026. Commit abd42394bc47d58440dd3db4c3bd1ffa19b403de.
<https://github.com/tadamcz/erdos126>. Selected quadratic source at the pinned revision.
- [3] Alon N, Solymosi J. Rank of matrices with entries from a multiplicative group. *International Mathematics Research Notices*. 2023;2023(14):12383–12399.
<https://doi.org/10.1093/imrn/rnac183>.
- [4] Bao S. *Signed laminar kernels and a quadratic cardinality bound*. Companion manuscript; 2026. Theorem 2.1.

- [5] Berg C, Christensen JPR, Ressel P. *Harmonic Analysis on Semigroups: Theory of Positive Definite and Related Functions*. Graduate Texts in Mathematics; vol. 100. New York: Springer; 1984.
<https://doi.org/10.1007/978-1-4612-1128-0>.
- [6] Erdős P, Turán P. On a problem in the elementary theory of numbers. *The American Mathematical Monthly*. 1934;41(10):608–611. <https://doi.org/10.2307/2301909>.
- [7] Evertse J-H, Schlickewei HP, Schmidt WM. Linear equations in variables which lie in a multiplicative group. *Annals of Mathematics*. 2002;155(3):807–836.
<https://doi.org/10.2307/3062133>.
- [8] Firsching M, Lezeau P, Mercuri S, Horváth MZ, Dillies Y, Sönne C, Wieser E, Zhang F, Hubert T, Agüera y Arcas B, Kohli P. *Formal Conjectures: An Open and Evolving Benchmark for Verified Discovery in Mathematics*. 2026. arXiv:2605.13171. <https://arxiv.org/abs/2605.13171>.
- [9] Győry K, Stewart CL, Tijdeman R. On prime factors of sums of integers I. *Compositio Mathematica*. 1986;59(1):81–88. https://www.numdam.org/item/CM_1986__59_1_81_0/.
- [10] Schoenberg IJ. Metric spaces and positive definite functions. *Transactions of the American Mathematical Society*. 1938;44(3):522–536.
<https://doi.org/10.1090/S0002-9947-1938-1501980-0>.
- [11] Wu B-L. Sumsets with restricted number of prime factors. *Lithuanian Mathematical Journal*. 2019;59(2):251–260. <https://doi.org/10.1007/s10986-019-09441-0>.

A Source and verification scope

The construction originates in

```
Erdos126/Resolutions/Erdos126_104usd_15h.lean
```

in Tom Adamczewski’s public repository [2], at commit

```
abd42394bc47d58440dd3db4c3bd1ffa19b403de.
```

The repository and benchmark attribute the proof to GPT-6 Astra [1, 2]. In this module, `card_le_three_sq` gives the bound for positive sets and `quadraticBound` gives the uniform natural-number bound. The declaration `signed_family_card_bound` is the abstract input, whose conventional proof appears separately in [4].

The selected module is an alternate resolution. At the pinned revision, `Solution.lean` imports `Erdos126_132usd_25h`; the quadratic module belongs to the separate `Erdos126Alternates` build target. The alternate’s local theorem `Erdos126.erdos_126` proves the same terminal statement in its own environment. Identical declaration names do not imply an import dependency. The target statement comes from the Formal Conjectures collection [8].

The source files and import configuration were inspected at the pinned revision. The Lean build and external verification toolchain were not independently run for this paper; compilation and formal-verification claims remain attributed to the repository and benchmark. The mathematical argument above rests on the explicit lemmas and the stated companion theorem, independently of those external checks.