

FINITE FRANKL–WILSON CERTIFICATES AND NEWTON ZERO COVERS

NEAL SHANDILYA
NEAL@MILLENNIUMRESEARCH.AI

ABSTRACT. We compare prime and prime-power versions of two normalized Frankl–Wilson certificate families and study their extension to Newton zero covers. In the Ramsey family, every tuned integer modulus is uniquely isolated; proper prime powers give gains exceeding every fixed power of the logarithm along a subsequence. In the Euclidean family, the verified finite-range maximum is 12597/8815, while the infinite problem reduces asymptotically to the product of the two neighboring-prime distances. A signed local expansion includes the integer-rounding loss and explains why the error-free quadratic envelope fails.

For Newton words, sparse Fourier lifts give a sharp degree-product bound and an exact all-characteristic envelope. Fixed-characteristic results include density and integer-attainment theorems, ordinary binary profiles, a Han–Teixeira two-support formula, and the universal power-plus-one profile. Actual-degree reductions lead to finite graph and hypergraph algorithms. Coprime-support geometry yields strict-sixth rectangle rigidity; in a first-essential single-leading-digit family the sharp nonrectangle thresholds are $2/9$ and, outside one binary row, $5/22$, while normalized maximizers are unique through relative deficit $1/3$. We give the main proofs and references for longer arguments. We distinguish mathematical proofs from external computational inputs and state the remaining open problems and questions of prior literature.

Results to look at first. For a quick overview of the main results, see Theorems 2.2, 3.2, 4.1, 5.2, 7.1, and 8.4. Section 9 records references for longer proofs, computational inputs, attribution, and open questions.

1. SCOPE, DEFINITIONS AND THE MAIN QUESTIONS

The Frankl–Wilson intersection theorem and its prime-power extension provide algebraic certificates for large Ramsey graphs and high-dimensional distance graphs [6, 7]. Their asymptotic strengths are classical. The questions here concern exact optimization, the effect of allowing prime powers, and the algebraic structure of more flexible separators. The graph comparisons below apply only to the stated certificate families, with no claim about unrestricted Ramsey numbers, all distance graphs, or the best currently known constructions.

Prime powers include primes; a *proper* prime power has exponent at least two. All logarithms are natural. Binomial coefficients are computed in $\mathbb{Z}$ before reduction into a field, including the usual integer-valued convention at negative arguments. Strict budget inequalities, origin normalization and upper versus actual degree are part of each definition.

For $S \subseteq \mathbb{Z}_{\geq 2}$ and integer $k \geq 2$, put

$$F_S(k) = \max \left(\{k-1\} \cup \left\{ \binom{m}{r} : \begin{array}{l} a \in S, \quad a \leq r \leq \lfloor m/2 \rfloor, \\ \max\{\binom{m}{a-1}, \binom{m}{\lfloor r/a \rfloor}\} < k \end{array} \right\} \right). \quad (1.1)$$

Here m is the ground-set size, r is the full uniform layer, and a is the modulus. Write F_P , F_{PP} and $F_{\mathbb{Z}}$ for primes, prime powers and all integers at least two. Only the first two are the stated Frankl–Wilson constructions; $F_{\mathbb{Z}}$ is a numerical relaxation. The first cost comes from the prime-power intersection bound and the second from the rank bound for the permitted intersection sizes.

For the Euclidean family set

$$B_m(t) = \frac{\binom{m}{2t-1}}{\binom{m}{t-1}}, \quad G_E(m) = \frac{\max_{q \in \text{PP}} B_m(q)}{\max_{p \in \text{P}} B_m(p)}, \quad (1.2)$$

where eligible parameters satisfy $2t - 1 \leq m$. These are unrounded quotients; the ambient Euclidean dimension is $m - 1$, not m .

A *Newton word* of upper budget A over a field K of characteristic p is

$$f(n) = \sum_{i=0}^{A-1} a_i \binom{n}{i}, \quad a_i \in K.$$

It is normalized or anchored when $a_0 = f(0) = 1$. Its actual dimension is one plus its largest nonzero Newton index. Thus budget A means degree at most $A - 1$; an optimum may have smaller actual dimension. For two independent coefficient fields of distinct prime characteristics p, q , define

$$G(f, g) = \min\{n \geq 1 : f(n) \neq 0, g(n) \neq 0\}, \quad H_{p,q}(A, B) = \max_{f,g} G(f, g).$$

The tests occur in the two fields separately. The zero-covered interval is $\{1, \dots, G - 1\}$, of length $G - 1$. A characteristic-power period at least the budget is a period of the word, so a common positive meeting always exists. The maximum H is independent of the coefficient extensions, although particular support patterns and the number of optimizing pairs need not be. Indeed, imposed zero conditions and $f(0) = 1$ form an affine system over the prime field; descending a feasible pair cannot shorten its gap, and at a maximum cannot lengthen it.

The global envelope is $h(A, B) = \max_{p \neq q} H_{p,q}(A, B)$. A budget is *literally essential* when lowering it by one strictly lowers this fixed-characteristic maximum. First essentiality is the inequality $H(A, B) > H(A - 1, B)$. A nonzero top coefficient in one displayed optimizer does not establish this condition. The normalized symmetric zero-cover model is already present in Gopalan [1], following the composite-modulus Boolean representation program of Barrington, Beigel and Rudich [11].

Sections 2–3 concern the original families; Sections 4–6 develop general zero covers; Sections 7–8 give the later actual-degree and rigidity refinements. Section 9 is a guide to the longer proofs and computational dependencies.

2. EXACT OPTIMIZATION IN THE ORIGINAL RAMSEY FAMILY

For $d \geq 2$, define

$$M_d(k) = \max\{m \geq 2(d - 1) : \binom{m}{d-1} < k\}, \quad Q_S(d) = \max(S \cap [2, d]),$$

when defined, and $R_d = \min\{dQ_S(d) - 1, \lfloor M_d(k)/2 \rfloor\}$.

Theorem 2.1 (Finite reduction). *Omitting undefined rows and rows with $R_d < Q_S(d)$,*

$$F_S(k) = \max \left\{ k - 1, \max_d \binom{M_d(k)}{R_d} \right\}.$$

Only $d - 1 < \log_2 k$ can occur.

Proof. For a feasible (m, r, a) take $d = \max\{a, \lfloor r/a \rfloor + 1\}$. Then $m \leq M_d$, $a \leq Q_S(d)$ and $r \leq ad - 1 \leq dQ_S(d) - 1$. Monotonicity on the increasing half of the binomial row gives the displayed upper bound. Conversely a retained row with modulus $Q_S(d)$ has both costs at most $\binom{M_d}{d-1} < k$. Finally $2^{d-1} \leq \binom{2d-2}{d-1} < k$, making the search finite. The discarded central-cap rows cannot contain a feasible row of actual budget d . $\square$

The inverse is computed by integer doubling and bisection, with the final inequalities $\binom{M_d}{d-1} < k \leq \binom{M_d+1}{d-1}$. If x solves $\binom{x}{d-1} = k$, its strict cutoff is $\lceil x \rceil - 1$, not always $\lfloor x \rfloor$.

Theorem 2.2 (Isolation at every integer target). *For every integer $Q \geq 2$, put*

$$m_Q = \left\lfloor Q^3 + 3Q^2 \log Q + \frac{5}{2}Q^2 \right\rfloor, \quad k_Q = \binom{m_Q}{Q-1} + 1, \quad N_Q = \binom{m_Q}{Q^2-1}.$$

Then $F_{\mathbb{Z}}(k_Q) = N_Q$ and the unique maximizing triple is $(m, r, a) = (m_Q, Q^2 - 1, Q)$. Consequently $F_{\text{PP}}(k_Q) > F_{\text{P}}(k_Q)$ at every proper prime-power target, while the two are equal at prime targets.

Proof outline and finite dependency. The finite reduction leaves virtual rows

$$V_d = \binom{M_d}{\min(d^2 - 1, \lfloor M_d/2 \rfloor)}.$$

Uniform estimates first retain all central caps and strict cutoffs, exclude remote degrees, and then compare the competitive degrees. The analytic inequalities hold for $Q \geq 50$. For $2 \leq Q \leq 49$, rational logarithm intervals certify each floor defining m_Q ; exact integers verify 2,304 non-target rows, 48 target rows and 48 baseline inequalities. The single tail inequality $2^{M_{2Q}} < N_Q$ for each target excludes every larger feasible degree, since M_d is nonincreasing. At $d = Q$, $M_Q = m_Q$; equality in both monotonicity steps forces $m = m_Q$, $r = Q^2 - 1$ and $a = Q$. The proof requires both the analytic argument and the exact initial-block script. [Detailed proof: Supplement P1.] $\square$

For a fixed positive integer h , delete the numerical moduli at distance less than h from Q . The maximum loss has the asymptotic

$$\log \frac{N_Q}{F_{\{a \geq 2: |a-Q| \geq h\}}(k_Q)} = h^2(\log Q + 3/2) + o(1).$$

Elementary congruences produce arbitrarily large proper prime powers with any prescribed fixed neighborhood free of primes. Applying the loss estimate after fixing h proves

$$\limsup_{k \rightarrow \infty} \frac{F_{\text{PP}}(k)}{F_{\text{P}}(k)(\log k)^C} = \infty \quad (C > 0). \quad (2.1)$$

The hole width is fixed before taking the limit, so the argument requires no uniform growing-modulus prime-distribution theorem.

2.1. Tuning, frequency and larger certificate families. The isolation proof also determines the eventual isolation region in $m = \lfloor Q^3 + (3 + \lambda)Q^2 \log Q + \beta Q^2 \rfloor$: it is $|\lambda| < 1$, together with $\lambda = 1, \beta \leq 4$ and $\lambda = -1, \beta \geq 1$. For numerical allowed sets, the sharp leading nearest-modulus radius is $\sqrt{3Q \log Q}$; sharpness for a pair of actual primes is not established. For fixed left and right hole distances h_-, h_+ , calibrated loss is $h_- h_+ (\log Q + 3/2) + o(1)$ at retuned budgets. Effective growing-hole theorems also give explicit full-window and half-window loss bounds, with separate numerical and prime-family onset checks; their quantifiers differ from Theorem 2.2.

Let

$$\mathcal{H}_R(K) = \sum_{\substack{2 \leq k \leq K \\ F_{\text{PP}}(k) > F_{\text{P}}(k)}} \frac{1}{k}.$$

A classical prime-gap moment gives $\mathcal{H}_R(K) \ll_{\varepsilon} (\log K)^{7/8+\varepsilon}$. Thus strict advantage has logarithmic density zero, despite both advantage and equality having lower natural density zero and

upper natural density one. This oscillation persists for each fixed polylogarithmic gain threshold. Conversely there are computable, unevaluated $\tau, c > 0, K_0$ such that for $K \geq K_0$,

$$\sum_{\substack{k \leq K \\ \log(F_{\text{PP}}(k)/F_{\text{P}}(k)) \geq \tau(\log \log k)^3}} \frac{1}{k} \geq c\sqrt{\log K \log \log K}.$$

The lower theorem uses a two-sign prime-square sieve and ordinary effective prime distribution. It supplies no numerical onset.

A weighted bridge relates selected prime-gap lengths, their logarithmic weight, Euclidean winning counts and Ramsey harmonic counts. It transfers specified leading asymptotics or two-sided order estimates; it proves no universal leading constant. Replacing both charges everywhere by the same $D(m, d)$ satisfying $\binom{m}{d} \leq D(m, d) \leq \sum_{j \leq d} \binom{m}{j}$ gives finite harmonic symmetric difference and controlled relative values, but need not preserve individual optima.

For distinct prime gaps with left endpoint at most X , the number containing a proper prime power is

$$\pi(X^{1/2}) + \pi(X^{1/3}) + O(X^{17/60}/\log^2 X).$$

The potential fourth-power term is obstructed by square-cube collisions: a third term $\pi(X^{1/4})$ is equivalent to the number of gaps containing both types being $o(\pi(X^{1/4}))$. This is open. A gap count is not a count of its length or of winning dimensions.

For a proper power q between neighboring primes $u < v$, put $g = v - u$ and $P_q = (q - u)(v - q)$. The global singleton peak over all integer budgets obeys

$$\log \max_{k \geq 2} \frac{F_{\{q\}}(k)}{F_{\text{P}}(k)} = (\log u + 3/2)P_q + O\left(\frac{g + \log^2 u}{u}P_q + \frac{1}{u}\right),$$

uniformly under $g = o(u/\log u)$, which follows from the imported explicit prime estimate. A common near-crossing integer budget is $O(1/u)$ -optimal for all proper powers in one gap. Exact staircase maximizers need not satisfy a two-nearest-integers rule.

The patched-LCM prime-square construction also controls recurrent peaks. Write $R(k) = F_{\text{PP}}(k)/F_{\text{P}}(k)$ and $A(K) = \max_{2 \leq k \leq K} \log R(k)$. For a fixed admissible Linnik exponent L ,

$$\liminf_{K \rightarrow \infty} \frac{A(K)}{(\log \log K)^3} \geq \frac{1}{4L^3}.$$

For every $\varepsilon > 0$ and $c < 1/(4L^2)$, each sufficiently large K has a tuned integer budget $K \leq k \leq \exp((\log K)^{L+\varepsilon})$ with $\log R(k) \geq c(\log \log k)^3$. The conservative source-checked value $L = 26/5$ [26, Theorem 1.1] gives constants 125/70304 and 25/2704. These are running-maximum and recurrence statements; no explicit onset or gain at every individual budget is asserted.

2.2. What changing the certificate does. Two ordinary polynomial separators enlarge the family. With actual degree costs, the exact mixed ordinary/prime-power optimizer is obtained by adding the nearest-two-prime rectangle rows. For example, at $k = \binom{169678}{57} + 1$ the mixed optimum is $\binom{169678}{3360}$, over characteristics 59 and 61, while the original optimum is $\binom{147868}{3480}$; their ratio is about $7.04834 \cdot 10^7$. Maynard's dense-cluster theorem [15] and a fixed-width CRT construction give gains beyond every fixed log power over the original prime-power family. In contrast, arbitrary prime blue residues and the specified one- p -adic-separator/exact-count-blue enlargement give no gain.

Charging the fixed canonical full-layer lifts by full rational ranks gives the same original values for primes and prime powers. For arbitrary prime-power sets S , this all-budget equality holds exactly when S is empty or contains two. Mixed modular-red/rational-blue ranks also preserve values and maximizing triples when $\text{P} \subseteq S \subseteq \text{PP}$; finite interval certificates enter this theorem. Sparse allowed sets and selected vertex sets are different. At budget three a selected three-set

family has both ranks two, whereas the full-layer optimum is two; the unrestricted selected optimum remains in $\{3, 4\}$. The exact selected optima for degrees two through four are two; degree five remains unresolved.

As an external comparator, the literal projective construction of Ihringer–Mattheus [22], with its precise two strict charges, exceeds $F_{\mathbb{Z}}$ for every $k \geq 38418$, sharply. Its refined clique charge has 234 literal improvements in thirteen windows, ninety improving the combined original/literal envelope; refined and literal values permanently agree from 2097153. These comparisons do not convert any of the present family-specific barriers into unrestricted Ramsey upper bounds.

3. EUCLIDEAN QUOTIENTS AND THE OPEN INFINITE TAIL

For fixed m , successive terms in (1.2) have ratio

$$\frac{B_m(t+1)}{B_m(t)} = \frac{(m-2t+1)(m-2t)}{2(2t+1)(m-t+1)},$$

which is strictly decreasing in eligible t . The quotient is therefore strictly log-concave. Let $u < q < v$ be the neighboring primes of a proper prime power q , and let $D(q)$ be its maximum contribution against the entire eligible prime family, with baseline one.

Proposition 3.1 (Two endpoints and two integer candidates). *There is a unique real crossing m_0 of $B_m(u) = B_m(v)$ in the relevant range. Wherever q gives a strict gain, the full prime denominator is $\max(B_m(u), B_m(v))$. The integer maximum $D(q)$ is obtained from the baseline one and the eligible integers immediately on either side of m_0 . Each fixed q contributes strictly at only finitely many ground-set sizes.*

Proof sketch. The discrete mode boundary is $\mu(t) = 4t + \frac{1}{2} + \frac{1}{2}\sqrt{32t^2 + 32t + 9}$. Every strict contribution lies between $\mu(u)$ and $\mu(v-1)$. The ratio $B_m(v)/B_m(u)$ increases strictly with real m and crosses one once. The two endpoint comparisons for q respectively increase and decrease, so their minimum peaks at the crossing. Strict log-concavity excludes all more distant primes on the strict-gain region. Integer rounding leaves only the two adjacent candidates. This does not assert equality of denominators outside the strict region. $\square$

Write $a = q - u$, $b = v - q$, $g = a + b$, $W(q) = ab/q$ and $A_0 = 2(\sqrt{2} - 1)$.

Theorem 3.2 (Product law and tail equivalence). *If $1000g \leq u$, then*

$$\left| \log D(q) - A_0 \frac{ab}{q} \right| \leq \frac{2g^3}{q^2} + \frac{g}{3q}.$$

Moreover, in the extended nonnegative reals,

$$\limsup_{m \rightarrow \infty} \log G_E(m) = A_0 \limsup_{\substack{q \rightarrow \infty \\ q \text{ proper prime power}}} W(q).$$

In particular $G_E(m) \rightarrow 1$ if and only if $W(q) \rightarrow 0$.

Proof sketch. At the crossing, the negative second difference of $\log B_{m_0}(t)$ is $4(\sqrt{2}-1)/u + O(g/u^2)$. Comparing with the quadratic having zero endpoint values at u, v gives

$$|\log D(q) - A_0 ab/q| \leq (7/2 + A_0) gab/u^2 + g/(4u).$$

The last term is the integer-rounding loss, bounded from exact derivatives in m . The prime number theorem gives $g = o(u)$, so the error is $o(W) + o(1)$. Fixed powers eventually stop contributing and bounded ground sets admit only finitely many powers; these facts transfer the two limsups. The Baker–Harman–Pintz exponent $21/40$ [12] improves the first displayed remainder to $O(q^{-17/40})$, but is unnecessary for the equivalence. $\square$

3.1. A signed refinement and a failed envelope. Let $r = (u + v)/2$, $C_0 = 5\sqrt{2}/3 - 2$, $\lambda = 3\sqrt{2} - 4$ and

$$\hat{m} = (4 + 2\sqrt{2})r - \frac{3}{2} - \frac{\sqrt{2}g^2}{12r}, \quad \eta = \{\hat{m}\}.$$

A refinement, uniform under the same $1000g \leq u$ condition, is

$$\begin{aligned} \mathcal{L}(q) &= A_0 \frac{ab}{q} + C_0 \frac{ab(a-b)}{q^2} - \frac{\lambda}{q} \min\{a\eta, b(1-\eta)\}, \\ |\log D(q) - \mathcal{L}(q)| &\leq 70g^4/q^3 + 3g^2/q^2. \end{aligned} \tag{3.1}$$

The cubic term records gap asymmetry; the negative term records the integer phase. The phase can be computed exactly in $\mathbb{Q}(\sqrt{2})$ using an integer square root, and rational logarithm enclosures give rigorous endpoint-only sufficient upper tests or lower certificates. A failed test is inconclusive. BHP makes the remainder $O(q^{-9/10})$, without bounding the leading arithmetic product ab/q universally.

The error-free inequality $\log D(q) \leq A_0 W(q)$ is false. With actual neighboring primes $23 < 27 < 29$,

$$D(27) = \frac{70775}{55022} \quad \text{at } m = 176, \quad \log D(27) > \frac{1}{4} > A_0 \frac{8}{27}.$$

The ambient dimension here is 175. This counterexample disproves the proposed local envelope and leaves the global conjecture below open. The refined proof uses the exact crossing, signed Taylor remainder and periodic phase control, without differentiating an unspecified asymptotic error.

3.2. The finite record and its limits. The exact finite result is

$$\max_{3 \leq m \leq 27313708498984759620} G_E(m) = \frac{12597}{8815}, \tag{3.2}$$

uniquely at $m = 60$, hence ambient dimension 59. It joins exact checks of the 80,070 proper prime powers at most 10^{12} to the imported prime-gap computation of Oliveira e Silva–Herzog–Pardi [13]. The scan's bound $g \leq 1476$ is used with its *right-endpoint* convention through $4 \cdot 10^{18}$, not as an unrestricted claim about gaps crossing that endpoint.

The frequency horizon is smaller. If $E(X) = \#\{3 \leq m \leq X : G_E(m) > 1\}$, then $E(6828468094921) = 13881694$ and $E(10^{12}) = 5339452$. The record horizon in (3.2) is not a complete count of all dimensions up to that horizon. Independently, $E(X) \gg \sqrt{X}$ with computable but unevaluated constants; for each fixed $R > 1$, the count with $G_E(m) \geq R$ is $O_{R,\varepsilon}(X^{57/100+\varepsilon})$, using the version-specific short-interval transfer from Järvinen [14].

Whether (3.2) remains the sharp bound for all m is open. The limsup identity equates tail boundedness and asymptotic limits with the corresponding statements about W ; boundedness of W does *not* imply the exact constant $12597/8815$, nor control all finite m . That exact bound would force a nearest prime within $(0.656461573775 \dots + o(1))\sqrt{q}$, but such a nearest-prime condition alone is not sufficient. Aggregate gap estimates also do not settle the issue: a modified prime-like subset can share the stated aggregate inputs and a prescribed finite prefix while having unbounded quotient.

The normalized binary-layer enlargement, optimizing $\binom{m}{r}/\binom{m}{t-1}$ for $t \leq r \leq \min(2t-1, \lfloor m/2 \rfloor)$, equals $\max_t B_m(t)$ for $m \geq 5$. At $m = 4$ the individual maxima change but their prime-power/prime ratio agrees. This is an exact equivalence of those normalized layer families only.

4. SPARSE FOURIER LIFTS AND THE DEGREE-PRODUCT BARRIER

For $f \in \mathbb{Q}_p[x]$ with $f(0) \neq 0$, let $Z_p(f) = \{n \in \mathbb{Z} : v_p(f(n)) > v_p(f(0))\}$.

Theorem 4.1 (Sparse lift and zero-cover bound). *If f has degree d , there is a complex p -power-periodic exponential sum w , with at most $d + 1$ distinct frequencies, such that $w(0) \neq 0$ and w vanishes on $Z_p(f)$. Consequently, if r consecutive positive integers are covered by separator sets of degrees $d_1, \dots, d_h$, then*

$$r + 1 \leq \prod_{i=1}^h (d_i + 1).$$

Primes may repeat. More generally, if polynomials of degrees at most d_i over $\mathbb{Z}/m_i\mathbb{Z}$ induce symmetric functions on $\{0, 1\}^n$, are individually nonzero at zero, and every other Boolean vector is zero for at least one polynomial in its own ring, then

$$n + 1 \leq \prod_i (d_i + 1). \quad (4.1)$$

The origins may be nonunits, and no coprimality of the m_i is required.

Proof of the product step and proof architecture. Multiply the sparse complex lifts. After combining equal frequencies, the product has at most $L = \prod_i (d_i + 1)$ terms and is nonzero at zero. A sum $\sum_{j=1}^s c_j z_j^n$, with distinct nonzero z_j , cannot vanish at s consecutive positive integers unless all $c_j = 0$: the corresponding Vandermonde matrix is invertible. Thus a zero block has length at most $L - 1$.

The lift itself uses the classical p -adic split replacement at the distinguished origin [7], followed by induction on residue fibres. Its invariant bounds occupied frequency residues at every scale; finite-dimensional column-basis gluing preserves this invariant with additive, rather than multiplicative, support loss. The resulting total is $d + 1$. The all-degree lift requires a separate proof; it does not follow from Vandermonde alone.

For a symmetric Boolean function over a residue ring, multilinearization gives a unique expression $\sum_j c_j \binom{t}{j}$ on weights t . Choose a prime-power component in which its origin survives, lift coefficients integrally, and compare valuations with the origin. Every zero in that component lies in the strict separator set. Apply the preceding bound to the weights $1, \dots, n$. No division by $j!$ inside the residue ring is used. [Detailed proof: Supplement P3.] $\square$

The exponential-sum obstruction and covering implications have classical antecedents in Sun [8]. Elkies gives a closely related arbitrary-factor product bound for the coprime-support problem [23]; neither the exponential-sum argument nor the product mechanism is presented as new. The exact scope of (4.1) is residue rings and the stated symmetric zero-cover acceptance rule, not every weak representation of OR or every finite ring.

For the full-layer h -colour family obtained from strict valuation separators, with actual degrees d_i , $d = \max_i d_i < r \leq m/2$ and strict costs $\binom{m}{d_i} < k$, the sharp leading growth is

$$\log \mathcal{F}_h(k) \sim \frac{1}{h^h} \frac{(\log k)^h}{(\log \log k)^{h-1}} \quad (h \geq 2 \text{ fixed}).$$

The product bound gives the upper layer cap; binomial asymptotics optimize it. Distinct-prime constructions and the prime number theorem supply a matching row at every sufficiently large budget, in the classical multicolour tradition [9]. For two colours the leading coefficient is $1/4$. The result does not constrain nonsymmetric colourings, selected vertex families or arbitrary actual-rank charges.

5. THE GLOBAL NEWTON ENVELOPE

Theorem 5.1 (Cap, equality and characteristic scaling). *For normalized Newton pairs, $G(f, g) \leq AB$. In distinct characteristics,*

$$H_{p,q}(A, B) = AB \iff A = p^a \text{ or } B = q^b \quad (a, b \geq 0).$$

In equal characteristics $G \leq A + B - 1$. Maxima satisfy the exact identities

$$H_{p,q}(pA, B) = pH_{p,q}(A, B), \quad H_{p,q}(A, qB) = qH_{p,q}(A, B).$$

The cap follows by lifting normalized integer Newton coefficients into Theorem 4.1. If $A = p^a$, use the period- A origin delta for the first word and an inverse- A dilate of the length- B alternating Newton word for the second. Their first meeting is AB . After prime-field descent, in equal characteristics the product has degree at most $A + B - 2$ and is nonzero at zero, so consecutive Newton interpolation gives the stronger bound.

For scaling, restriction to multiples of a characteristic removes its low digits; for the reverse direction insert a low-digit origin delta and inversely dilate the other word. These are identities for maxima, not divisibility assertions about arbitrary witness gaps. The necessity in cap equality uses digit-reversal flag equivalence and a rank-one moment matrix after removing characteristic parts. Two independent single-root exchanges would force $(t-s)(w-v) = 0$ for distinct exchanged roots, a contradiction. The equality characterization requires a separate proof; the cap alone does not imply it. [Detailed proof: Supplement P4.]

Let $U(X)$ be the greatest prime power at most X , and set

$$C(A, B) = \min_{\substack{P > A, Q > B \\ P, Q \text{ prime powers of distinct bases}}} (P - A)(Q - B).$$

Theorem 5.2 (Exact all-characteristic envelope). *For $A, B \geq 2$,*

$$h(A, B) = \max\{BU(A), AU(B), AB - C(A, B)\}. \quad (5.1)$$

Also $h(A, 1) = A$ and $h(1, B) = B$. Allowing equal characteristics does not increase these maxima.

Proof outline. Origin deltas give the first two branches. The opposite-period CRT rectangle gives the third whenever positive. For the upper bound choose a smallest pair of budgets attaining a putative larger value, so both are literally essential, and divide out characteristic parts. The retained relative-rigidity lemma says that reduced essential budgets $d, e > 1$, prime to their own characteristics, with $1 \leq de - H \leq de/49$ must have the rectangle value $de - (R - d)(S - e)$ for the least enclosing characteristic powers R, S . Monotonicity of this affine expression then contradicts the proposed excess over (5.1).

An explicit Dusart estimate [10] supplies the required near-cap guard at large budgets. Exact scalar enumeration below 400000 leaves 118 exceptional individual budgets, ending at 1360, and only seven unordered budget pairs. Their values are

(A, B)	(6, 6)	(12, 12)	(14, 14)	(20, 20)	(20, 21)	(33, 33)	(33, 34)
$h(A, B)$	34	140	190	385	410	1057	1098.

The upper bounds for these pairs import 690 oriented rank-certificate records. The full-support argument retains the weak endpoint $H \geq E + 1$; replacing it by a strict inequality would lose a boundary case. The imported records are not replaced by small brute-force corroboration. [Detailed proof: Supplement P4.] $\square$

At (200, 200) the deficit is 200 whereas $C(200, 200) = 253$, so the delta branches are essential. For $A, B \geq 2$, equality $h = AB$ occurs exactly when at least one budget is a prime power.

For the full-layer two-word unit family F_2 , replace $dQ_S(d) - 1$ in Theorem 2.1 by $h(d, d) - 1$. This gives an exact finite optimizer, retaining only rows $R_d = \min\{h(d, d) - 1, \lfloor M_d/2 \rfloor\} \geq d$, as well as the baseline. Equality with the original family persists at tuned prime-power targets despite

unbounded gains elsewhere. Bounded-deficit inverse results, characteristic-base restrictions and nonunit equality statements inherit their own hypotheses; they do not give arbitrary below-cap nonunit maxima or repeated-prime equality classifications.

6. FIXED-CHARACTERISTIC PROFILES AND CLASSICAL STRUCTURE

Fix $p \neq q$, first budget d , and let P be the least p -power at least d . Write $H(e) = H_{p,q}(d, e)$ and $\rho = \inf_{e \geq 1} H(e)/e$.

Theorem 6.1 (Continuous profile and density). *For $e \geq 2$, $0 \leq H(e) - H(e-1) \leq P$. There is a unique nondecreasing P -Lipschitz function F with*

$$F(0) = 0, \quad F(e/q^n) = q^{-n}H(e), \quad F(qx) = qF(x).$$

The ratio $F(x)/x$ on $[1, q]$ has image $[\rho, d]$, also the full cluster set of integer ratios. If w is the minimum nonzero weight of the period- P Newton code, then

$$H(e) \geq \left\lceil \frac{Pe}{w} \right\rceil, \quad w = \min_{0 \leq i < d} \prod_j (p - i_j), \quad i = \sum_j i_j p^j.$$

When $q > P$, $H(w) = P$ and $\rho = P/w$.

Main mechanisms. If an optimal gap $G > P$, the first word is nonzero at P , so $g(P) = 0$. The normalized difference $g(n) - g(n+P)$ has budget at most $e-1$ and a gap at least $G-P$, giving the jump bound. Scaling in Theorem 5.1 defines F on the dense q -adic grid, and the bound extends it continuously. A minimum-support first word, averaged over translations, leaves a sufficiently short list of constraints; interpolation and backward origin repair give the density bound. The minimum-weight formula is classical repeated-root coding theory [16, 17]; minimum weight alone does not characterize every possible support. $\square$

For ordinary $1 \leq d \leq p$, set $s = p - d + 1$ and write $e - 1 = ks + r$, $0 \leq r < s$. The stronger staircase bound is $H(e) \geq pk + d + r \geq pe/s$. If $2 \leq d < p < q$, equality in $H(e)/e = p/s$ holds exactly when $e = st$ and every balanced base- q digit of t belongs to $[-\lfloor q/p \rfloor, \lfloor q/p \rfloor]$. This uses p -unit Newton determinants and residue balance, rather than dividing high factorials in a finite field. Integer attainment and the full real minimizing locus are distinct questions.

6.1. Uniform rows and the complete ordinary binary solution. For $1 \leq d < p$ and N the greatest q -power at most e , the two-branch formula

$$H_{p,q}(d, e) = \max\{dN, pe - (p-d)qN\} \tag{6.1}$$

holds for every e precisely when $p > q(p-d)$ for odd q . For $q = 2$ the criterion is that d is even or $2d > p$. Codimension-one rows with $p > q$ have separate normalized-pair uniqueness and essentiality theorems; the value formula alone does not establish them.

All ordinary binary rows admit an exact recursion. For odd p , let $H_d(e) = H_{p,2}(d, e)$ and $N \leq e < 2N$ be a power of two. Then $H_1(e) = e$; even d satisfy (6.1); and for odd $d \geq 3$, $0 < r < N$,

$$H_d(N) = dN, \quad H_d(N+r) = \max\{dN, (d-1)N + H_3(r), 2dN - p(N-r)\}.$$

If $u = M_0 + \dots + M_k$ lists its nonzero binary digits in decreasing order and $S_i = \sum_{j < i} M_j$, then for $p \geq 5$,

$$H_3(u) = \max_i \{2S_i + 3M_i, pu - (p-2)S_i - 2(p-3)M_i\}.$$

Thus every row has a finite-knot piecewise-linear normalized profile. For $2 \leq d < p$, the ratio maximum d occurs exactly at powers of two, and its rational infimum is never attained at an integer budget. The row $d = 1$ is the explicit exception: its ratio is identically one.

The period numerator gives the following relation. For a characteristic- q word of period $Q = q^a$ and numerator $V(z) = \sum_{r < Q} g(r)z^r$,

$$\text{actual dimension} = Q - \text{ord}_{z=1} V(z).$$

This is classical periodic-sequence linear complexity [18, 19, 20]. Its binary folding and support intervals determine optimizer dimensions: the first ordinary budget is forced, odd rows have one possible second actual dimension, and even rows can have gaps in the dimension spectrum. Unique dimensions do not in general imply unique coefficients.

6.2. Two supported residues and the Han formula. Consider any integer period $P \geq 2$, coprime to q , and first supports containing zero with at least two residues. A larger support contains a two-set and cannot improve the maximum. Its exact profile is

$$F(x) = \frac{Px}{2} + P \sup_{j \in \mathbb{Z}} q^{-j} \left(\text{dist}(q^j x/2, \mathbb{Z}) - \frac{1}{P} \right)_+. \quad (6.2)$$

For integer e , let $D_Q(e) = \text{dist}(e, 2Q\mathbb{Z})$. Then

$$H(e) = \frac{Pe + \max(\{0\} \cup \{PD_Q(e) - 2Q : Q = q^k, 2Q \leq Pe\})}{2}.$$

There are $O(\log(Pe))$ arithmetic steps. For a genuine first code of budget $P - 1$, $P = p^a$, Pascal inversion makes its value vectors exactly the hyperplane $\sum_{r < P} f(r) = 0$. Every two-set is realized by values $1, -1$. This justifies the transfer to $H_{p,q}(P - 1, e)$; it does not say every budget- $P - 1$ word has two supported residues.

Here the exact Newton rank problem becomes a local Toeplitz block, then the syzygy gap of $X^e, Y^b, (X + Y)^k$. The tent envelope is an application of the classical Han theorem in the inspected Teixeira formulation [4], with related Lefschetz and Jordan-rank antecedents [2, 3, 5]. It is not an independent discovery of that scalar syzygy theorem. Source-version and original-thesis access limitations are recorded in Section 9.

Equality with the density line occurs when $\text{dist}(q^j x/2, \mathbb{Z}) \leq 1/P$ for every integer j , including endpoints. If $q < P$, the full normalized profile simplifies to $F(x) = \max\{P - 1, Px - q\}$ on $[1, q]$; the density line need not be attained. For $P \geq 3$, off its density locus the general profile is locally piecewise affine with slopes zero and P ; the density locus itself is nondifferentiable. These statements exclude the linear endpoint $P = 2$.

6.3. Beyond two supports. For uniform supports of size at least s , $2 \leq s \leq P - 2$, density equality is characterized by $\text{dist}(q^j x/s, \mathbb{Z}) \leq 1/P$ for all $j \in \mathbb{Z}$. At $q > P$ the slope P/s is attained at $e = s$; at $q = P - 1$ only the normalized point sq/P attains it, and no integer budget does; at $q \leq P - 2$ the actual infimum is larger. Missing-one supports require an additional rational P -grid. Their minimizing set combines a balanced-digit self-similar set with isolated non- q -adic points. These results include the complete first-budget-two row and explain why integer minima can miss part of the real locus.

For all supported blocks, retain every block in the local map

$$\bigoplus_i K[t]_{< L_i} \longrightarrow K[t]/(t^e), \quad (A_i) \longmapsto \sum_i (1 + t)^{k_i} A_i.$$

Homogenization yields displayed generators X^e and $Y^{K-k_i+\epsilon_i}(X + Y)^{k_i}$, where $K = \max k_i$, $L = \max L_i$ and $\epsilon_i = L - L_i \in \{0, 1\}$. If d_j are the syzygy shifts of this displayed module, retaining redundant generators, then

$$\text{rank } A = \sum_i L_i - \sum_j (K + L - d_j)_+.$$

This gives an exact implicit maximal-syzygy profile with certified convergent rounding bounds. For more than two blocks it is not a closed Han-style lattice formula; deleting intermediate blocks changes the rank.

A general fixed-row digit/rank recursion computes exact values, and a finite automaton makes the row graph q -recognizable. Rational-word extremum results [21] give rational, effectively computable ratio infima and decisions about integer attainment; the state spaces can be enormous. A zero-slack graph describes the full minimizing locus and its dimension $\log_q \lambda$. For the ordinary quadratic row $p > q \geq 3$, the exact ratio infimum is $pq/((p-1)(q-1))$, at the unique normalized scalar minimizer $x = (q-1)(p-1)/p$, never at an integer budget. Higher codimension and binary optimizer spectra also fit this framework. The adjacent-mask shortcut is false, for instance $H_{13,11}(3, 75) = 97$, not 96. Effective evaluability does not guarantee a practical uniform algorithm in all parameters.

7. POWER-PLUS-ONE PROFILES AND ACTUAL-DEGREE ALGORITHMS

The earlier isolated rows at budgets five, nine, ten and seventeen are instances of one complete upper-budget theorem. The scalar result is distinct from its equality and actual-degree refinements.

Theorem 7.1 (Universal power-plus-one row). *Let $p \neq q$ be primes and $P = p^a \geq 3$, $a \geq 1$. Put $T = pP$, $C = T - P - 1$, and let N be the largest q -power at most e , with $Q = qN$. Then, over arbitrary independent coefficient fields,*

$$H_{p,q}(P+1, e) = \max\{(P+1)N, Pe, Te - CQ\}. \quad (7.1)$$

Its minimum ratio is P , precisely on the normalized interval

$$\frac{P+1}{P} \leq x = \frac{e}{N} \leq \frac{Cq}{P(p-1)}.$$

At every integer budget in this interval, the normalized maximizing pair is unique. It is the period- P origin delta and

$$g(n) = \sum_{j=0}^{e-1} (-1)^j \binom{un}{j}, \quad uP \equiv 1 \pmod{Q},$$

with actual dimensions P, e . For $e > 1$, both literal upper budgets are essential exactly when

$$e = N \quad \text{or} \quad P(p-1)e > CQ.$$

The plateau, density line and upper affine branch have separate Newton witnesses. Lucas coordinates reduce a nonzero top coefficient to rooted fibres with one root per fibre; interval spans, cyclic-difference graphs and full-rank Pascal/Vandermonde windows exclude larger gaps. The arbitrary-extension argument uses support domination or rank invariance rather than a finite-field coefficient scan. At integer ratio minima, the endpoint inequalities are strict by coprimality. Strictly excluding every nonzero top coefficient forces the period- P first delta, and unit-dilated interpolation fixes the second word. Scalar equality alone would not prove uniqueness.

The theorem covers both odd and binary first characteristics, all powers $P \geq 3$, and every distinct second prime. It supersedes earlier sufficient ranges such as $q > pP$, and the previous interval-only statements. Classical Han reductions already give substantial scalar subfamilies and terminal branches; the anchored dictionary, complete composition and optimizer claims must be distinguished from those antecedents. The primitive binary power $P = 2$ is excluded from (7.1). For its ordinary-second slice $1 \leq B < q$ the exact formula is instead

$$H_{2,q}(3, B) = \max\{2B + (B \bmod 2), 4B - q\}, \quad H_{2,q}(2, B) = 2B. \quad (7.2)$$

For powers $P = p^a, R = q^b$ of distinct prime bases, $a, b \geq 1$ (including a power equal to two), the upper-budget value is $H(P + 1, R + 1) = PR + \max(P, R)$ in the corresponding characteristic order. Orienting $P > R$, its normalized pair is unique, with actual dimensions $P, R + 1$.

7.1. Actual first dimension is a different optimization. Here let $P = p^a \geq 2, a \geq 1$, and let $H_{p,q}^{\text{top}}(P + 1, e)$ require the first degree to be exactly P . It scales exactly under $e \mapsto qe$, but need not equal the upper-budget maximum. For odd p and binary second characteristic,

$$H_{p,2}^{\text{top}}(P + 1, e) = \max\{(P + 1)N, pPe - 2(pP - P - 1)N\}.$$

The middle Pe branch of (7.1) can use a smaller first dimension and disappears here. For ordinary linear second words, q odd,

$$H_{p,q}^{\text{top}}(P + 1, 2) = \begin{cases} P + 3, & P \equiv \pm 1 \pmod{q}, \\ P + 2, & \text{otherwise.} \end{cases}$$

There are explicit affine families of maximizing first words. Outside resonance there are two disjoint families, each of dimension $\lfloor P/q \rfloor$ over the first field. At positive or negative resonance the dimension is respectively $(P - 1)/q - 1$ or $(P + 1)/q - 1$. The unique-pair cases are exactly odd primes $q = P \pm 1$, including $(P, q) = (2, 3)$. Actual-top optimizers are therefore not unique in general. For quadratic second words and $q \geq 5$, the exact value is

$$H_{p,q}^{\text{top}}(P + 1, 3) = \begin{cases} P + 5, & P \equiv \pm 1, \pm 2 \pmod{q}, \\ P + 3, & \text{otherwise.} \end{cases}$$

It includes $P = 2$ but asserts nothing for $q = 2, 3$. The linear and quadratic scalar rows extend to second budgets $2q^b$ and $3q^b$ by the actual-top scaling identity; scaled optimizer classifications do not automatically follow.

The following graph gives an exact low-degree reduction. Assume $P > 2k, q > 2k$, and $k \geq 1$. With $\alpha = P \bmod q$, let Γ_s have edges $\{i, i + \alpha\}$ in $\mathbb{F}_q$ for $1 \leq i \leq s \leq 2k$. Write τ_0 for a minimum vertex cover avoiding zero. Then

$$H_{p,q}^{\text{top}}(P + 1, k + 1) = P + 1 + \max\{0 \leq s \leq 2k : \tau_0(\Gamma_s) \leq k\}. \quad (7.3)$$

Each component is a path, zero is isolated or an endpoint, and $\tau_0 = \sum_C \lceil \ell_C/2 \rceil$. Under the additional $q \geq 3k$, if $c = \min(\alpha, q - \alpha)$, the maximum s is $2k - (k \bmod c)$.

Why the graph is exact. Lucas gives $f(r + Pt) = A_r + ct$ with a common $c \neq 0$. Two positions $i, i + P$ cannot both be first zeros, so the ordinary second root set must hit each edge. Its k roots cover at most $2k$ labelled edges. Conversely, below $2P$ these paired positions are the only repeated low fibres. A cover therefore leaves at most one required root in each fibre, realized by choosing A_r with common top coefficient -1 ; the origin fibre is rooted at P . This constructs normalized actual-top words and a degree- k second polynomial. The restriction $P > 2k$ is essential to this two-block sufficiency argument. $\square$

Without that restriction the pair graph can overpredict: at $(p, P, q, k) = (3, 3, 7, 3)$ it predicts ten while the actual maximum is seven. Dropping $q > 2k$ also fails. These counterexamples show why both restrictions are needed in (7.3).

7.2. Single-leading-digit hypergraphs and penultimate rows. For actual first degree $D = dP, P = p^a$ with $a \geq 0, 1 \leq d < p$, and ordinary second degree at most $k < q$, Lucas coordinates make each low fibre a degree- d polynomial in the top digit, with one common nonzero leading coefficient. Exact prefix feasibility is a hitting-set problem on second residues: retain every projected hyperedge from $d + 1$ same-low-fibre positions with distinct top digits, and the mandatory constraints from positive multiples of pP . The root set avoids zero and has at most k elements. Projected smaller edges, even singleton zero, cannot be discarded.

Hitting all these constraints leaves at most d required roots in a fibre. Complete the origin roots to d nonzero digits; normalization fixes the common leading coefficient. Complete the other fibres with that same coefficient and invert the Pascal matrix. This proves prime-field realization, arbitrary-extension validity and equality of upper/actual second maxima. Finite event compression and bounded-rank branching give fixed-parameter tractability in the combined parameter (d, k) , with binary- P input cost. For $d = 1$ this becomes a graph with loops and mandatory vertices and is fixed-parameter tractable in k . A compact fibre oracle does not list the full Newton vector: outputting all coefficients costs $\Omega(dP)$.

The endpoint maxima are $D + 1$ at $k = 0$ and $(D + 1)q$ at $k = q - 1$. The short bound $D + (d + 1)k + 1$ requires $q > (d + 1)k$. Only the proved q -power-scaled ordinary-base rows follow by insertion; arbitrary multidigit second budgets are not solved by this reduction.

For fixed distinct primes p, q , $1 \leq d < p$ and $2 \leq B < q$, with $P = p^a$, $a \geq 1$, there is also an exact phase-affine reduction. For $P > q$ in a reachable nonzero residue phase $s = P \bmod q$, the actual-top maximum has the form $T_d(P; B - 1) = C_s P + R_s$, with $0 \leq R_s \leq q$, and

$$H_{p,q}(dP + 1, B) = \max\{PH_{p,q}(d, B), T_d(P; B - 1)\}.$$

Powers in one phase form a geometric sequence, so finitely many exact linear inequalities decide the all-exponent cases for those fixed parameters; $P \leq q$ is checked separately. The earlier rectangle certificate for 66 prime pairs is a bounded characteristic-grid result. Its strict-sixth conclusions now have a symbolic all-prime proof in Section 8; the finite grid alone gives no conclusion about all primes.

For actual first degree P and second upper dimension $q - 1$, the value is $(P + 1)q - pP$ exactly when $q \geq 2p - 1$, and is larger below that threshold. In range the second word is unique and actual; $q > 2p$ gives full unscaled uniqueness, while $q = 2p - 1$ can leave free first fibres. For the higher digit $2 \leq d < p < q$, the full upper-budget problem has value

$$H_{p,q}(D + 1, q - 1) = V = (D + 1)q - pP,$$

a unique normalized pair with both degrees actual, and first predecessor $V - q$. The first roots are $q, 2q, \dots, Dq$ modulo pP ; the second support is $\{0, -pP\}$ modulo q , with values $1, -1$. This scalar follows from the Han dictionary. The equality classification also requires the root-forcing proof; a unique maximizing Han scale is insufficient.

8. COPRIME SUPPORT GEOMETRY, SHARP RIGIDITY AND ISOLATION

The support problem also makes sense without polynomial coefficients. Let $M, N \geq 2$ be coprime, and let supports $U \subseteq \mathbb{Z}/M\mathbb{Z}$, $V \subseteq \mathbb{Z}/N\mathbb{Z}$ contain zero. Suppose they miss at most $\alpha < M$ and $\beta < N$ residues. Set

$$G = \min\{n \geq 1 : n \bmod M \in U, n \bmod N \in V\}, \quad S = M + N, \quad R = M(\beta + 1) - N(M - \alpha - 1).$$

The canonical zero sets are

$$Z_U^* = \{jN : 1 \leq j \leq \alpha\} \pmod{M}, \quad Z_V^* = \{jM : 1 \leq j \leq \beta\} \pmod{N}. \quad (8.1)$$

They give gap R whenever $R > 0$. The CRT rectangle construction and a large-gap uniqueness idea appear explicitly in Meyerowitz's 2013 discussion of this same support problem [24]. The dense-support branch $G \leq M + N - 1$ when $|U|/M + |V|/N \geq 1$ has an inequality-form Fine–Wilf antecedent [25, Theorem 2.2.2]. The quantitative results below are stated without a historical-priority claim.

Lemma 8.1 (Refined component inequality). *In an orientation with $G > M$, let $u = |U|$, $b = |V^c|$ and let c be the number of connected components of the forced second-zero intervals. Then*

$$cG \leq R + (c - 1)(S - c) - M(\beta - b) - N(u - (M - \alpha)), \quad (8.2)$$

with $1 \leq c \leq \min(u, b, \lfloor N/2 \rfloor)$.

Proof. Label first supports by $j = -vN^{-1} \bmod M$, $0 \leq j < M$. After multiplication by M^{-1} modulo N , the forced second zeros form intervals

$$[1, \lfloor (G-1)/M \rfloor] \quad (j=0), \quad [\lfloor jN/M \rfloor, \lfloor (G-1+jN)/M \rfloor] \quad (j>0).$$

They lie in $[1, N-1]$; their union has at most b positions. For component i , let j_i^-, j_i^+ be its extreme labels, w_i its length, and $s_i = j_i^- N \bmod M$, $e_i = (G-1+j_i^+ N) \bmod M$. Exact endpoints, including the origin component, give

$$Mw_i = (G-1) + (j_i^+ - j_i^-)N + s_i - e_i.$$

The starts s_i are distinct, as are the ends e_i . Hence $\sum s_i \geq c(c-1)/2$, $\sum e_i \leq c(M-1) - c(c-1)/2$, and $\sum (j_i^+ - j_i^-) \geq u - c$. Sum these inequalities and use $\sum w_i \leq b$ to obtain (8.2). Counting labels, occupied residues and separating gaps gives the stated range for c . $\square$

Theorem 8.2 (Support isolation and the short branch). *If $R > S$, every noncanonical support pair satisfies*

$$G \leq \max \left\{ \left\lfloor \frac{R+S-2}{2} \right\rfloor, R - \min(M, N) \right\} = R - \delta, \quad \delta = \min \left\{ \left\lceil \frac{R-S+2}{2} \right\rceil, M, N \right\}.$$

If $R < S$, writing $K = S - R$ and $m = \min(M, N)$ gives

$$G \leq \max\{m, S+1 - \lceil \sqrt{4K} \rceil\}.$$

The ceiling is evaluated with integer square roots. Neither bound is asserted to equal the exact second-best value in every profile.

Proof. In the strict affine branch the unrefined component bound gives $G \leq R$. For $c \geq 2$, (8.2) gives $G \leq S+1 - c + (R-S)/c \leq (R+S-2)/2$. If $c=1$, put $t = R - G$ and $u_0 = M - \alpha$. If $j_{\max} = u_0 - 1 + \ell$, the exact final endpoint is $\beta + \lfloor (M-1-t+\ell N)/M \rfloor \leq \beta$, so $\ell N \leq t$. Thus $t < N$ forces the canonical first support; $t < M$ forces the canonical second support. Together they force $G = R$. More generally each directional first-support difference from the canonical pair is at most $\lfloor t/N \rfloor$, and each second-support difference at most $\lfloor t/M \rfloor$.

For $R < S$, orient the smaller period first. Either $G \leq m$, or (8.2) yields $G \leq S+1 - c - K/c \leq S+1 - 2\sqrt{K}$; integer rounding proves the claim. $\square$

Both isolation terms are attained: $(M, N, \alpha, \beta) = (5, 7, 2, 5)$ with $U = \{0, 3, 4\}$, $V = \{0, 6\}$ has $G = 13$, $R = 16$; $(5, 7, 3, 5)$ with $U = \{0, 3\}$, $V = \{0, 4\}$ has $G = 18$, $R = 23$. The first is disconnected and the second connected. Arbitrary nonzero periodic values remain free even when the supports are unique.

8.1. The strict-sixth theorem and its Newton transfer.

Theorem 8.3 (Ordinary coprime-support rigidity). *For coprime $M, N \geq 3$, $2 \leq A < M$, $2 \leq B < N$, and supports missing at most $A-1, B-1$ residues,*

$$G \leq \max\{MB + NA - MN, 5AB/6\}.$$

If the maximum H exceeds $5AB/6$, it equals the affine value and both literal dimensions are essential. The $1/6$ deficit threshold is sharp in this support class: $M=3$, prime $N \geq 5$, $A=2$, $B=3$ gives $H=5$ with both predecessors three and affine value $9 - N < 5$.

The proof supplements the component bound with the exact B th first-support position, reversed overlap and an integer-cap inequality. Ordinary Newton dimensions below distinct prime characteristics realize exactly these root sets, so no essentiality assumption is needed there. Composite coprime periods in this auxiliary theorem do not authorize arbitrary integer Frankl–Wilson moduli. [Detailed proof: Supplement P7.]

Now return to primes $p \neq q$, $P = p^a$, $a \geq 1$, $1 \leq d < p$, and

$$A = dP + 1, \quad 2 \leq B < q, \quad H = H_{p,q}(A, B), \quad \Delta = AB - H, \quad R = pPB + qA - pPq. \quad (8.3)$$

Assume *first literal essentiality*. It forces every optimizer to have actual first degree dP . Each of its P Lucas fibres is a degree- d polynomial with common nonzero leading coefficient, so it has at most d zeros per fibre, at most $A - 1$ in period pP . Theorem 8.3 therefore applies and gives

$$0 < \Delta < AB/6 \implies H = R.$$

The integral-binomial CRT witness attains R . The second literal budget is automatically essential in this window: otherwise its predecessor would inherit first essentiality and contradict the same support bound at $B - 1$; $B = 2$ uses $H(A, 1) = A$.

This transfer requires first essentiality. An upper-budget word of lower actual degree can have more than $A - 1$ zeros in period pP ; the ordinary root cap is unavailable for arbitrary proper-power Newton words. The general arbitrary-budget strict-sixth rigidity question remains open.

8.2. Sharper thresholds, no-wrap families and equality.

Theorem 8.4 (Sharp nonrectangle thresholds and uniqueness). *In (8.3), under first literal essentiality, every numerical nonrectangle $H \neq R$ has*

$$\frac{\Delta}{AB} \geq \frac{2}{9}.$$

Equality occurs exactly at $(p, P, d, B) = (2, 2, 1, 3)$ with prime $q \geq 7$. Outside $(p, P, d) = (2, 2, 1)$, the sharp lower bound is $5/22$, with equality exactly at $(p, P, d, q, B) = (5, 5, 2, 3, 2)$.

Every normalized maximizing coefficient pair is unique when

$$0 < \frac{\Delta}{AB} \leq \frac{1}{3}.$$

This endpoint is included; the sufficient uniqueness window is not claimed optimal. Every first-essential point of the binary row $(p, P, d) = (2, 2, 1)$ is unique without a deficit restriction.

The numerical rigidity windows are strict: the equality cases are genuine nonrectangles. Uniqueness has a different boundary and includes nonrectangles. From (7.2), genuine first-essential binary nonrectangles have odd $B \geq 3$, $q \geq 2B + 1$, gap $2B + 1$ and relative deficit $(B - 1)/(3B)$. Their unique first word is $1 + \binom{n}{1} + \binom{n}{2}$; if $B = 2m + 1$, their second word is $\prod_{j=1}^m (1 - n/(4j - 1))(1 - n/(4j))$ in characteristic q .

The nonbinary endpoint belongs to a complete no-wrap family. Put $O = H_{p,q}(d, B)$. When $O < p$, first essentiality holds exactly for

$$d = m(q - B + 1), \quad mq < p \quad (m \geq 1).$$

Then $H = mqP + B$, the first predecessor is mqP , the second is smaller, and the normalized pair is unique. Its deficit is $mP(B - 1)(q - B)$, giving the sharp $5/22$ minimum above. The ordinary test is $O < p$ exactly when $d + (B - 1)\lceil d/(q - B + 1) \rceil < p$, with that expression equal to O . This classification uses the staircase and balanced-node rank mechanisms already described, and the penultimate slices also have Han antecedents.

In the wrap regime, support components and fibre caps force the affine value in the near-cap range; a symbolic large/small-power argument gives the stated sharp constants. The uniqueness proof separately forces saturated root sets, handles the binary corners and uses integral unit dilation plus Pascal interpolation. At a first-essential global maximum, the origin fibre has exactly d forced nonzero roots and fixes the common leading coefficient. An underforced other fibre has at least two saturated prime-field completions. Consequently uniqueness is equivalent to one saturated maximizing first support and exactly $B - 1$ forced second roots. This is an exact

finite criterion; uniqueness is not inferred from essentiality alone. [Detailed proof: Supplement P8.]

8.3. Quantitative coefficient isolation and its two guards. For actual first degree $D = dP$ and ordinary second degree at most $k < q$, take $M = pP$, $N = q$, $\alpha = D$, $\beta = k$ in Theorem 8.2. If $R > M + q$, canonical support forces the unique normalized coefficients

$$f(n) = \binom{D - \ell n}{D} \pmod{p}, \quad \ell q \equiv 1 \pmod{pP}, \quad g(n) = \prod_{j=1}^k \left(1 - \frac{n}{jM}\right) \pmod{q}.$$

Every other pair in this *actual-degree stratum* has gap at most $R - \delta$. In the full first-essential upper-budget problem, lower-degree competitors must be retained:

$$G_{\text{noncanonical}} \leq \max\{H_-, R - \delta\}, \quad H_- = H_{p,q}(D, k + 1). \quad (8.4)$$

For $p = 2, P = 4, d = 1, q = 5, k = 3, R = 17, \delta = 3$ and $H_- = 16 > 14$, proving that the predecessor term cannot be dropped.

Nor do prime-field second-best searches always control extensions. At $p = 2, P = 4, d = 1, q = 3, k = 2$, the noncanonical actual-degree maximum is nine over the prime fields, but an $\mathbb{F}_4/\mathbb{F}_3$ pair attains twelve, the isolation cutoff. Here $B = k + 1 = q$, outside the stricter $B < q$ domain of Theorem 8.4. Saturating a noncanonical competitor can turn it into the excluded canonical optimizer, so maximum-value descent does not preserve the second-best question.

9. EVIDENCE, PROVENANCE AND THE REMAINING PROBLEMS

This report refers to the supplementary texts for the longer proofs. The accompanying `PROOF_MAP.md` resolves Supplements P1–P8 to the supplied proof texts and identifies external dependencies. These supplementary source texts are outside this PDF’s twenty-page limit; they do not constitute a second PDF.

Proof group	Coverage and the principal dependency
P1 Original certificates	Finite optimizer, all-target isolation, deletion and tuning; harmonic bounds and bridge; rank, ordinary and projective comparisons
P2 Euclidean geometry	Exact crossing and prime comparison; finite record and frequency horizons; product law, signed phase refinement and exact counterexample
P3–P4 Fourier/envelope	All-degree sparse lift, nonunit residue-ring cover, characteristic scaling, cap equality, 1/49 interface and seven finite residual pairs
P5 Fixed profiles	Density, balanced digits, uniform/binary rows, dimensions, Han two-support dictionary, automata, uniform supports and all-block syzygies
P6 Power/actual degree	Universal power-plus-one composition and minimizing pairs; linear/quadratic top rows; edge/hypergraph algorithms and penultimate equality
P7–P8 Support/rigidity	Coprime components, strict-sixth theorem, no-wrap family, sharp 2/9 and 5/22 thresholds, inclusive 1/3 uniqueness and isolation

9.1. What the computations certify. The all-target theorem includes a guarded exact initial-block verification for $2 \leq Q \leq 49$; rational logarithm enclosures certify floors and all final comparisons use integers. The finite Euclidean record imports the large proper-power enumeration and the external prime-gap table, whereas the frequency scan has its separate, smaller horizon. The global envelope imports the 690 oriented rank records. These computations are identified as proof dependencies, without any claim that this condensed source archive independently regenerates every large certificate.

Independent implementations of small Newton rank and literal-coefficient searches corroborate profile and equality formulas within specified finite ranges. Later support audits include complete mask enumerations through bounded periods, original-coordinate Newton checks, and extension-field examples. Rerunning an existing checker tests reproducibility without providing an additional

independent implementation. Finite agreement does not replace the unbounded arguments. Exploratory scripts using floating-point floors, bounded scans, or an argmax without tie counting do not certify exact floors or unique maxima; assertion-dependent scripts must be run with assertions enabled. No claim is made that every historical diagnostic is a guarded exact verifier.

9.2. Attribution and assistance. Classical antecedents and external results are cited where used in the text; the literature search is not intended as an exhaustive determination of priority. AI assistance was used to develop, debug, and review Python scripts for finite scans and computational checks. Certified finite comparisons use exact arithmetic, and computational checks apply only within their stated ranges and do not replace proofs of infinite claims.

9.3. Open questions.

- (1) Does $G_E(m) \leq 12597/8815$ hold for every ground-set size m ? Can the product $W(q)$ at actual proper prime powers be controlled strongly enough to prove this exact bound or convergence to one? The signed local refinement alone does neither.
- (2) What is the full fixed-characteristic Newton profile for arbitrary first budgets, and what closed finite-state or syzygy descriptions remain practical beyond two supports? Actual single-leading-digit algorithms do not solve arbitrary multidigit second budgets.
- (3) Does strict-sixth rectangle rigidity extend to arbitrary reduced, literally essential Newton budgets? Can normalized optimizers be classified beyond the sufficient one-third single-digit window? No unrestricted first-essential classification is asserted here.
- (4) Can the selected-gap square–cube collision term be made smaller than the fourth-power count? What are the leading constants in the weighted Euclidean/Ramsey bridge and the precise harmonic growth?
- (5) What happens for below-cap nonunit covers, repeated-prime equality, genuinely selected-layer rank certificates, and the unresolved degree-five selected problem? Improvements in those models require their own certificates.
- (6) Which specialized statements have earlier formulations in the literature, and which proofs survive external mathematical review or formal verification?

REFERENCES

- [1] P. Gopalan, *Constructing Ramsey graphs from Boolean function representations*, *Combinatorica* **34** (2014), 173–206. Definition 2, pp. 175–176; Lemma 14 and Proposition 15, pp. 197–199. doi:10.1007/s00493-014-2367-1. The model predates the journal appearance; the ECCC report is TR05-143 (2005).
- [2] A. R. Kustin and A. Vraciu, *The Weak Lefschetz Property for monomial complete intersections*, arXiv:1110.2822v1 (2011), Definition 3.2 and Propositions 3.3–3.4, pp. 9–10. Preprint.
- [3] K.-i. Iima and R. Iwamatsu, *On the Jordan decomposition of tensored matrices of Jordan canonical forms*, *Math. J. Okayama Univ.* **51** (2009), 133–148; especially pp. 136–137. Journal PDF.
- [4] P. Teixeira, *Syzygy gap fractals—I. Some structural results and an upper bound*, *J. Algebra* **350** (2012), 132–162. Inspected text: arXiv:1008.0583v1 (2010), Section 2, pp. 4–6; equation (2), p. 7; Example 3.8, p. 9; Theorem 3.9 (Han) and Figure 2, p. 10. Preprint. Page locators here refer to that preprint, not journal pagination.
- [5] L. Nicklasson, *The strong Lefschetz property of monomial complete intersections in two variables*, *Collect. Math.* **69** (2018), 359–375. Inspected text: arXiv:1703.07970v1 (2017), Theorem 4.11 and proof, p. 19. Preprint.
- [6] P. Frankl and R. M. Wilson, *Intersection theorems with geometric consequences*, *Combinatorica* **1** (1981), 357–368. Theorems 2 and 8. DOI.
- [7] L. Babai, P. Frankl, S. Kutin and D. Štefankovič, *Set systems with restricted intersections modulo prime powers*, *J. Combin. Theory Ser. A* **95** (2001), 39–73. Definition 2.1, p. 44; Lemma 6.1, p. 50. DOI.
- [8] Z.-W. Sun, *Groups and combinatorial number theory*, arXiv:math/0411289v1 (2004), Theorem 1.2 and its proof, pp. 4–5. Version-pinned preprint.

- [9] V. Grolmusz, *Superpolynomial size set-systems with restricted intersections mod 6 and explicit Ramsey graphs*, *Combinatorica* **20** (2000), 71–86 (publisher pagination). Section 6, printed pp. 83–84 in the inspected typeset copy. DOI.
- [10] P. Dusart, *Estimates of some functions over primes without $R.H.$* , arXiv:1002.0442 (2010). The statement used is that for $x \geq 396\,738$ the interval $[x, x(1 + 1/(25 \log^2 x))]$ contains a prime. Preprint.
- [11] D. A. M. Barrington, R. Beigel and S. Rudich, *Representing Boolean functions as polynomials modulo composite numbers*, *Comput. Complexity* **4** (1994), 367–382. DOI.
- [12] R. C. Baker, G. Harman and J. Pintz, *The difference between consecutive primes, II*, *Proc. London Math. Soc.* **83** (2001), 532–562. Theorem 1 is imported for the $21/40$ exponent. Inspected primary scan.
- [13] T. Oliveira e Silva, S. Herzog and S. Pardi, *Empirical verification of the even Goldbach conjecture and computation of prime gaps up to $4 \cdot 10^{18}$* , *Math. Comp.* **83** (2014), 2033–2060. Right-endpoint bound, printed p. 2050. DOI.
- [14] O. Järviniemi, *On large differences between consecutive primes*, arXiv:2212.10965v2 (2022). The version-specific short-interval and prime-gap-tail inputs are used with the separate fixed-fraction threshold transfer; no journal-version identity is assumed. Primary version record.
- [15] J. Maynard, *Dense clusters of primes in subsets*, *Compositio Math.* **152** (2016), 1517–1554. Theorem 3.4, p. 1520. Primary text.
- [16] J. L. Massey, D. J. Costello, Jr., and J. Justesen, *Polynomial weights and code constructions*, *IEEE Trans. Inform. Theory* **19** (1973), 101–110. Lemma 1, p. 102; Theorem 6.1, p. 109. Primary author archive.
- [17] G. Castagnoli, J. L. Massey, P. A. Schoeller and N. von Seemann, *On repeated-root cyclic codes*, *IEEE Trans. Inform. Theory* **37** (1991), 337–342. DOI.
- [18] J. L. Massey and S. Serconek, *Linear complexity of periodic sequences: A general theory*, *CRYPTO 1996*. Inspected corrected author postprint: Proposition 2, p. 3; Theorem 9, p. 8; same-field products, Section 5.2. Corrected postprint.
- [19] R. A. Games and A. H. Chan, *A fast algorithm for determining the complexity of a binary sequence with period 2^n* , *IEEE Trans. Inform. Theory* **29** (1983), 144–146. The identity used here was checked through Lauder–Paterson’s Lemma 1, rather than a fresh inspection of this original full text.
- [20] A. G. B. Lauder and K. G. Paterson, *Computing the error linear complexity spectrum of a binary sequence of period 2^n* , HPL-1999-128(R.1), 9 August 2000. Sections II.A and II.C; Lemmas 1 and 3; Theorem 7. Inspected technical report. Later journal version: *IEEE Trans. Inform. Theory* **49** (2003), 273–280; theorem numbering here refers to the report.
- [21] L. Schaeffer and J. Shallit, *The critical exponent is computable for automatic sequences*, arXiv:1104.2303v3 (2011), Theorem 12 and Corollary 13. Section 8 explains a flaw in an earlier broader limsup claim; that claim is not used. Version-pinned preprint.
- [22] F. Ihringer and S. Mattheus, *An improved algebraic construction for Ramsey numbers*, arXiv:2608.21769v1. Preprint, August 2026; used only as an external family comparator. arXiv record.
- [23] N. D. Elkies, *Answer to Long enough interval of integers to solve a simultaneous congruence*, *MathOverflow*, 5 March 2013, edited 6 March 2013. Verified answer permalink. Roots-of-unity proof of the arbitrary-factor product bound.
- [24] A. Meyerowitz, *Answer to Long enough interval of integers to solve a simultaneous congruence*, *MathOverflow*, 4 March 2013, edited 5 March 2013. *MathOverflow* discussion. CRT rectangles and a large-gap optimality/uniqueness discussion; not cited as a proof of the complete quantitative theorems here.
- [25] M.-w. Wang, *Periodicity and Repetition in Combinatorics on Words*, Ph.D. thesis, University of Waterloo, 2004. Theorem 2.2.2, pp. 9–11, the inequality form of Fine–Wilf. Thesis.
- [26] T. Xylouris, *Über die Linniksche Konstante*, arXiv:0906.2749v1, 15 June 2009. Theorem 1.1, printed p. 9, gives exponent 5.2 with an effectively computable uniform constant; the least-prime maximum is defined on printed p. 6. Version-pinned preprint. This conservative exponent is not asserted to be current-best.