FLOORS FOR DETERMINISTIC INTEGER FACTORING: COVERING MODELS, PARTIAL INFORMATION, THE ADDITIVE STRUCTURE OF THE LEHMAN–HARVEY START SET, AND THREE FURTHER FLOORS

SEAN WARD AND MAESTRO

Abstract. The fastest proven deterministic factoring algorithm for a balanced semiprime $N = pq$ runs in $N^{1/5+o(1)}$ operations, and Harvey's question whether the search over Lehman's $N^{1/3}$ candidates can reach $N^{1/6}$ is open. This report defines a covering model that contains Harvey's algorithm exactly and, up to factors $N^{o(1)}$, those of Lehman and Harvey–Hittmeir, and proves its floors: $N^{1/3}$ candidates, $N^{1/5}$ with materialised giants, $N^{1/6}$ with free giants. The exponents persist for $p \geq N^{2/5}$, the first two against fixed-degree forms and adaptive cell queries respectively, and thinning by a residue hypothesis modulo $\mathsf{q} \leq N^{1/20}$ buys at most a power of $\log\log N$. Given an interval of length $L = N^{\lambda}$ containing p, a localised search costs $N^{o(1)}L^{3/5}N^{-1/10}$, matching the model's floor for $\lambda \geq 3/8$; for $N^{1/6} \ll L \leq N^{3/8}$ the in-model cost is classified by the interval's Diophantine depth; a known residue $p \bmod M$, $M \leq N^{1/7}$, gives $N^{1/5+o(1)}M^{-2/5}$. In the explicit-difference model an $N^{1/6}$ search needs a difference cover, and the additive structure of the window starts governs every such cover: proved lower bounds $r^{1/3}$ (modulus-free) and $N^{1/12}$ (at $r = N^{1/3}$; for the squarefree shell, for almost all N) against upper bounds $2.1\,r^{2/3}$ and $6N^{2/9}$, with the envelope problem open. Further, a common factor $g = N^{\gamma}$ of $p-1$ and $q-1$ lets N be factored deterministically from N alone below $N^{1/5}$ for every $\gamma > 1/20$, by composing McKee–Pinch's progression search, made deterministic, with Pollard–Strassen on $N-1$ and Harvey–Hittmeir's order selection; a fixed list of elliptic curves factors every product of two primes from a range once its exposure labels separate, with computed separating lists for three dyadic ranges below 2^{23}; and NFS polynomial selection has a random-root floor $N^{1/(d+1)-o(1)}$, a rigorous resultant floor and, under a stated counting heuristic, no room to move the GNFS constant. Every rigorous floor is a theorem inside a stated model, the heuristic optima are labelled as such, and nothing here is a lower bound on factoring.

Contents

Date: 2 September 2026.

Introduction and reading guide

The question. Let $N = pq$ be a balanced semiprime. The fastest deterministic factoring algorithm whose running time is proved runs in $N^{1/5+o(1)}$ operations (Harvey, 2021); it is a babystep–giantstep collision search over Lehman's $N^{1/3}$ candidate cells $aq + bp$, and its author asks whether the full square root of that candidate count, $N^{1/6}$, is attainable. Nothing in the literature bounds the cost of any deterministic factoring method from below; the known lower bounds in cryptography live in representation-oblivious models [75, 19, 2] that do not see the bits of N on which Lehman-type algorithms operate. This report asks a narrower question that can be answered: *inside a model that contains the known Lehman-type algorithms, what is the least possible cost, what would an algorithm below it have to do, and which of those things are actually possible?* The method of Parts 1–3 is the same throughout – define a model containing the published algorithms, prove floors inside it, identify the combinatorial statistic that governs the remaining gap, and attach certified or exact computation to every measured claim. Parts 4–6 apply pieces of it to three further questions: a deterministic algorithm with a proved cost for one class of moduli, a separation criterion with computed certificates and heuristic asymptotics for a fixed list of elliptic curves, and rigorous coefficient inequalities beside a heuristic counting optimum for polynomial selection.

What the report establishes. Tables 1 and 2 list the principal results with their status and their location; every entry is a reference into the text. Five status labels are used: *proved* means a theorem with a complete proof; *proved in the model* means a theorem whose hypotheses and conclusion are those of the stated model, so that it bounds algorithms inside the model and nothing outside it; *conditional* means proved under a conjecture stated in full; *heuristic* means a derivation whose hypothesis is a counting or distributional assumption, labelled as such where it occurs; *computed* means an exact or certified finite computation whose scope is stated. A row with several clauses names the status of each. In the table, r is the radius of the cell family (cells (a, b) with $ab \le r$), L and M are the length of the interval and the modulus of the residue known about p, $D_{\max}$ is the start statistic of Lemma 4.3, φ is Euler's function, $\|f\|_\infty$ is the largest absolute value of a coefficient of f, and c_θ is the constant of Theorem 20.11.

Table 1: The principal results of Parts 1–3, with status and location.

result	where	status
A covering model containing Harvey's algorithm exactly and, up to $N^{o(1)}$, those of Lehman and Harvey–Hittmeir; Hittmeir's residue-thinned search and Hart's one-line method are members for which the floors are not asserted	Definition 2.1, Proposition 2.9	proved
Floors: $N^{1/3}$ candidates, $N^{1/5}$ with materialised giants, $N^{1/6}$ with free giants	Corollary 3.3, Theorems 3.4, 3.6	proved in the model
Robustness: exponents kept for $p \ge N^{2/5}$; materialised-giant floor for fixed degree; candidate floor against adaptive queries; residue-hypothesis thinning modulo $\mathsf{q} \le N^{1/20}$ buys at most $(\log\log N)^{3/5}$	Proposition 3.7, Theorem 9.4, Proposition 9.5, Theorem 2.8	proved in the model

continued on the next page

result	where	status
In the explicit-difference representation model an $N^{1/6}$ search requires a difference cover (sufficient under order and decoding conditions); every explicit cover is bounded through the start statistic $D_{\max}$; a uniform envelope $D_{\max} \le x^{\eta}$ on the squarefree $a = 1$ shells, $0 \le \eta < 1/2$, gives the floor $N^{(1-\eta)/(5-4\eta)}$	Proposition 4.1, Lemma 4.3, Theorem 4.4	proved in the model
Pollard–Strassen's $1/4$ is optimal for divisibility tests by polynomially bounded integers in the difference-cover model	Proposition 4.5	proved in the model
Localised search in $N^{o(1)}L^{3/5}N^{-1/10}$ given an interval of length $L \ge N^{3/8}$, matching the model's floor	Theorem 13.6, Corollary 13.8, Theorem 13.2	algorithm proved (unconditional); floor proved in the model
For $c_6 N^{1/6} \le L \le N^{3/8}$ the in-model cost of an interval is classified by its Diophantine depth; a Farey-fraction family covers every interval at the typical cost, which is optimal for all but a vanishing fraction of the centres when $L/N^{1/6} \to \infty$	Theorem 13.15, Corollary 13.16	proved in the model
Residue information $p \bmod M$, $M \le N^{1/7}$: $N^{1/5+o(1)}M^{-2/5}$; the two-axis landscape against Coppersmith subdivision	Proposition 14.2, Corollary 14.3	algorithm proved (unconditional)
The start statistic on the balanced restriction: $\gg \sqrt{r}$ along rays and equal products; $\gg rN^{-1/4}$ along half-offset lines in the primitive family and in every family retaining the constructed cells	Lemmas 19.1, 19.2, 19.3	proved
Conditional in-model floors $N^{5/28}$ and $N^{2/11}$ for families satisfying the hypotheses of Proposition 19.5	Proposition 19.5, Corollary 19.7	conditional (Conjectures 19.4, 19.6)
The modulus-free statistic: lower bound $c_\theta r^{1/3}$ with the sharp constant of the symmetric mechanism; upper bound $2.1\, r^{2/3}$; one drift-free family holds at most $7\, r^{1/3}$	Theorems 20.11, 21.14, 20.16	proved
The planar regime: resonant clusters $\gg N^{1/12}$ at $r = N^{1/3}$ (for the squarefree shell, for almost all N); upper bound $6N^{2/9}$	Theorems 21.2, 21.4, 21.3, 21.13	proved

Table 2: The further results of Parts 4–6, with status and location.

result	where	status
A common factor $g = N^\gamma$ of $p-1$ and $q-1$: for $g \le (p-1)^{3/4}/(\log N)^3$, deterministic factoring from N alone in $N^{1/4-\gamma/2+o(1)}$ (McKee–Pinch's expected-time exponent made deterministic; beyond that bound a Burgess completion gives $N^{1/(8\sqrt{\mathrm{e}})+\varepsilon}$), and for arbitrary $g = N^\gamma$ the piecewise curve $E(\gamma)$ composing the progression search with Pollard–Strassen on $N-1$ and order selection, below $N^{1/5}$ for every $\gamma > 1/20$	Lemma 26.2, Corollary 27.1, Proposition 28.2	proved
A fixed list of curves factors pq iff the exposure labels of p and q differ; separating lists for the primes of $[2^{18}, 2^{19})$, $[2^{20}, 2^{21})$, $[2^{22}, 2^{23})$	Proposition 32.4, Table 12	criterion proved; lists computed
Pollard's $p-1$ succeeds on a positive proportion of nearly balanced semiprimes in $N^{\theta/2}$ for every $\theta > 1/(2\sqrt{\mathrm{e}})$	Proposition 35.3	proved
NFS polynomial selection: random-root floor $N^{1/(d+1)-o(1)}$ and resultant floor $\|f\|_\infty^e \|g\|_\infty^d \gg N$; joint floor $N^{(2d-1)/(d^2+d-1)}$ and the GNFS constant as existence optimum; small-degree enumerations	Propositions 38.1, 40.1; Propositions 39.1, 42.1; Section 37	proved; heuristic; computed

How the parts fit together. Part 1 is the base: it defines the covering model, proves that the published algorithms belong to it, proves the three floors and their robustness, reduces the $N^{1/6}$ question to a difference-cover question, and lists what the model forbids and what is known about each forbidden operation. Part 2 asks what changes when partial information about a factor is available – an interval, a residue, or both – and shows that the model's floors localise, that a localised form of Harvey's search attains them, and that below $N^{3/8}$ the cost of an interval depends on its position through a single Diophantine quantity. Part 3 studies the combinatorial statistic on which the difference-cover floor depends – the additive structure of the window starts – and is the longest part because that structure is rich: exact and approximate coincidences, their growth laws, their upper bounds, and the envelope problem that stands between the conditional floors and a theorem. Part 4 treats the one class of moduli that Part 1's bisection procedure leaves aside, and disposes of it from N alone. Parts 5 and 6 are independent of the covering model and of each other: Part 5 proves a separation criterion for a fixed list of elliptic curves, computes the lists that separate the primes of three dyadic ranges, and states its asymptotic consequences as a conjecture and a heuristic; Part 6 proves two coefficient inequalities for polynomial selection and places a heuristic counting optimum beside them, checked by exact enumeration at small sizes.

A reader who wants the main line should read Sections 1–4 of Part 1, Section 13 of Part 2, and Sections 19–21 of Part 3; Section 24 summarises what is proved, conditional, computed and open across the first three parts. Each part opens with a summary and an introductory section. Part 2 presupposes Definition 2.1 and Lemmas 3.1, 3.2; Part 3 presupposes those and the explicit-difference representation model with Lemma 4.3 and Theorem 4.4; Part 4 refers to Proposition 7.3 for context only; Parts 5 and 6 are self-contained. Sections and results are numbered consecutively through the report and every cross-reference names its target; the appendices of Parts 1 and 2 are the last sections of those parts.

What is not claimed. No statement in this report is a lower bound on the cost of factoring. Every rigorous floor is a theorem inside a stated model, and each part names the operations its model forbids; a faster algorithm would have to use one of them. The heuristic optima of Part 6 are labelled as such. Numerical statements are finite computations on stated populations, exact or certified where the text says so, and they are evidence for the constructions and the implementations, not for the asymptotics. Every measured or certified statement is regenerable from the code and archives in the public repository [23], whose reproduction guide names, for each statement, the function that produced it, the archive that holds it and the command that regenerates it.

Part 1. Covering floors for Lehman-type deterministic factoring, and what an $N^{1/6}$ algorithm would need

Summary. Harvey's deterministic factoring algorithm runs in $N^{1/5+o(1)}$ operations, and Harvey asks whether the babystep–giantstep search over Lehman's $N^{1/3}$ candidates admits a full square-root speed-up to $N^{1/6+o(1)}$. We define a coverage model – an oblivious family of integer-linear cells $aq+bp$ whose tested integers cover the real range of the smaller factor, searched by one collision step – which contains Harvey's algorithm exactly and, up to factors $N^{o(1)}$, those of Lehman and Harvey–Hittmeir, with Hittmeir's residue-thinned search and Hart's one-line method as further members, and we prove its floors: at least $cN^{1/3}$ candidates, and a cost of at least $cN^{1/5}$ when the giant elements are materialised and at least $cN^{1/6}$ when they are free. All three floors keep their exponents for unbalanced factors $p \geq N^{2/5}$; the materialised-giant floor holds for nonconstant forms of every fixed degree and the candidate floor against adaptive cell queries; and a family thinned by an enumerated residue hypothesis modulo $\mathsf{q} \leq N^{1/20}$ – the mechanism of the Harvey–Hittmeir refinement – gains at most a power of $\log\log N$, which is what that refinement obtains. In the explicit-difference representation model an $N^{1/6+o(1)}$ search requires an efficiently computable difference cover of the candidate set, and a lemma bounds every such cover through an approximate-Sidon statistic of the window starts. The statistic has provable power growth in every sub-family for which it has been determined, so a bounded statistic is impossible there, and under two stated envelope conjectures the explicit-difference floors for the controlled families are $N^{5/28}$ and $N^{2/11}$. None of this is a lower bound on factoring: every floor is a theorem inside a stated model, and the part closes with the shapes an algorithm below $N^{1/5}$ must take.

1. Introduction

Let $N = pq$ with primes $p < q$ and, for a fixed constant $C > 1$, $p \in J := [\sqrt{N/C}, \sqrt{N}]$. Lehman's algorithm [49] finds p in $N^{1/3+o(1)}$ operations by Dirichlet approximation of p/q: for a parameter $r \geq C$ there are coprime a, b with $ab \leq r$ and $0 \leq aq + bp - 2\sqrt{abN} < \sqrt{N}/(4r\sqrt{ab})$, and $(aq+bp)^2 - 4abN = (aq-bp)^2$ is then a square, detected by trying the few integers in the window. Hittmeir [40] and Harvey [31] replaced the square tests by collision searches; in Harvey's form, for a unit α modulo N, the unknown integer $U = aq+bp$ satisfies $\alpha^U \equiv \alpha^{aN+b} \pmod{p}$, so the $r \log r$ targets $\alpha^{aN+b-\lceil 2\sqrt{abN}\rceil}$ and the window offsets can be matched by babystep–giantstep with m baby steps and a product tree, and the cost $N^{1/2}/(r^{1/2}m) + r + m$ is minimised at $r = m = N^{1/5}$. Harvey asks (Remark 3.4 of [31]) whether the full square root of Lehman's candidate count, $N^{1/6}$, can be attained. The term that prevents it is the construction of the r targets themselves, one per cell.

This part studies the question inside a model and determines what the model allows. The model (Definition 2.1) contains Harvey's search exactly and, up to factors $N^{o(1)}$, Lehman's algorithm and the hypothesis-thinned search of Harvey–Hittmeir [32], and it has Hittmeir's residue-thinned search [39] and Hart's one-line factoring [30] as further members, the latter

where its family covers the range (Proposition 2.9); it does not encompass smoothness methods, and its lower bounds do not apply to algorithms that exclude parts of J using information computed from N.

1.1. **Results.** Write $c_C := (1 - C^{-1/2})^2/(80(C+1))$ and $c'_C := 2^{4/5}(c_C/2)^{2/5}$. The results are stated in the notation of Sections 2–4: Σ_W is the number of tested integers of a family, T its cost in the models M1 and M2, Z the multiset of its shifted tested exponents and $D_{\max}$ the statistic of Lemma 4.3.

- *Floors* (Section 3). Every oblivious covering family has $\Sigma_W \geq (c_C/2)^{2/3}N^{1/3}$ tested integers; the cost is $\geq c'_C N^{1/5}$ with materialised giants (Theorem 3.4) and $\geq 2(c_C/2)^{1/3}N^{1/6}$ with free giants (Theorem 3.6). Lehman's choice $r \asymp N^{1/3}$ is optimal for every oblivious family, and $1/5$ is optimal with materialised giants. All three floors keep their exponents for unbalanced factors $p \approx N^{\beta}$, $\beta \geq 2/5$ (Proposition 3.7); the materialised-giant floor $N^{1/5}$ holds for oblivious families of nonconstant integer forms of every fixed total degree, with a weaker candidate floor $N^{1/4}$ (Theorem 9.4); and the candidate floor $N^{1/3}$ holds against adaptive cell-membership queries (Proposition 9.5). A family that enumerates a hypothesis on p modulo q and tests one residue class modulo q^2 per form – the Harvey–Hittmeir mechanism – has, for $\mathsf{q} \leq N^{1/20}$, cost at least $c(\varphi(\mathsf{q})/\mathsf{q})^{3/5}N^{1/5}$ (Theorem 2.8), which their algorithm attains up to logarithms: hypothesis thinning buys at most a power of $\log\log N$.
- *Reduction* (Section 4). If $\operatorname{supp}(Z) \subseteq B - G$ for explicitly computable sets B, G of $O(\log N)$-bit integers with $|B| + |G| = N^{1/6+o(1)}$ satisfying an order condition and a decoding condition, there is an $N^{1/6+o(1)}$ factoring search (Proposition 4.1). In the explicit-difference representation model such a cover is required by definition, and Lemma 4.3 bounds its size below in terms of the approximate-Sidon statistic $D_{\max}$ of the window starts of any disjoint sub-family; Theorem 4.4 turns a uniform envelope $D_{\max} \leq r^{\eta}$ into the floor $cN^{(1-\eta)/(5-4\eta)}$. The same counting gives Pollard–Strassen's exponent $1/4$ as the floor of the difference-cover model for divisibility tests by polynomially bounded integers (Proposition 4.5).
- *The start set* (Section 5). The starts are the integers $\lfloor(\sqrt{aN} - \sqrt{b})^2\rfloor$. In the balanced family $a \leq b \leq Ca$ the statistic is $\gg \sqrt{r}$ along rays (Lemma 19.1) and along pairs of equal product (Lemma 19.2); after both are removed, half-offset lines force $D_{\max} \geq rN^{-1/4}/15 - 1$ for $47N^{1/4} \leq r \leq N^{1/3}$ and $C \geq 2$ (Lemma 19.3), so a bounded statistic is impossible for these families. Under an envelope hypothesis the conditional floor is $N^{5/28}$ (Proposition 19.5), and under a thinned-statistic hypothesis $N^{2/11}$ (Corollary 19.7).
- *Escapes* (Section 7). What the model forbids is listed with its status: the product modality (the tested product as one algebraic object) is formalised with its separation problem and a bisection that reduces that problem to a width bound and an explicit exceptional class of moduli (Propositions 7.2, 7.3); a strategy that evaluates pieces of low harmonic mass separately, at a cost at least the square root of each piece's length, pays $N^{1/4-o(1)}$ (Theorem 10.1).

1.2. **What is not claimed.** Nothing here is a computational lower bound. Theorems 3.4–4.4 hold within the model of Definition 2.1 and, for 4.4, the explicit-difference representation model of Section 4.1. The conditional floors are conditional on stated conjectures. Numerical statements are finite computations on the moduli listed in Section 6.

1.3. **Related work.** The closest lower-bound and barrier results are in representation-oblivious frameworks, and none is a lower bound for factoring. In a generic cyclic group with a prime-order component of order ℓ, the discrete logarithm and specified related problems require $\Omega(\sqrt{\ell})$ generic group operations [75], and root extraction in a generic group of unknown order has a generic lower bound of the same kind [19]; in the generic ring model, breaking RSA is *equivalent* to factoring [2], an equivalence and not a complexity bound. Neither framework sees the bit-level representation of N, which is exactly what Lehman-type algorithms exploit, and we have not found a published model that formalises coverage-based divisor searches and proves a cost floor within it. On the upper-bound side the exponent record is unchanged since [31]: the refinement [32] improves the polylogarithmic factor, and the element of large order that every Harvey-type search needs is now supplied for every bound D by Harvey and Hittmeir [33], so that its selection is no bottleneck at any exponent (Section 16 gives a balanced-semiprime proof with a different recovery step). An unrefereed note of Carella [14] claims a deterministic $O(N^{1/6+\varepsilon})$ algorithm; we have not been able to verify its argument, and the present part's floors say only what such an algorithm would have to look like (Section 7). Difference covers at the square-root scale are classical in other algorithmic settings – suffix-array construction [48], and discrete logarithms in an interval [29, 8] – but we know of no prior use of difference covers, or of the approximate-Sidon statistic of Lemma 4.3, in the design or analysis of factoring algorithms. Partial information about the factors – an interval containing p, or $p \bmod M$ – is treated in Part 2, where the materialised-giant floor of Section 3 is localised and matched by an algorithm, and the candidate and free-giant floors are localised.

2. The model and its members

Write $J_C := [\sqrt{N/C}, \sqrt{N}]$, so that $J = J_C$. Throughout, log is the natural logarithm and lg the binary one, $\pi(x)$ the number of primes up to x, φ Euler's function, $\omega(n)$ the number of distinct prime factors of n, $d(n)$ its number of divisors, v_p the p-adic valuation, $\Psi(x, y)$ the number of y-smooth integers in $[1, x]$, and $\tilde{O}$ hides factors $(\log N)^{O(1)}$.

Definition 2.1 (oblivious Lehman-cell BSGS algorithm)**.** The algorithm computes from N alone: (1) a finite set of pairwise distinct *forms* $u_i(p, q) = a_i q + b_i p$, integers $a_i, b_i \geq 1$, $i = 1, \ldots, P$; (2) for each form a finite set $T_i \subset \mathbb{Z}$ of *tested integers*, $W_i := |T_i| \geq 1$, with *candidate set* $S_i := \bigcup_{n \in T_i} (n-1, n+1)$ of measure $w_i := |S_i| \in [W_i, 2W_i]$; write $\Sigma_W := \sum_i W_i$ and $\Sigma_w := \sum_i w_i$; (3) *oblivious coverage*: for every real $p \in J$ there is an i with $u_i(p, N/p) \in S_i$; (4) a collision search between baby steps and giant elements, charged by the two cost models below: when T_i is an interval of integers, as in the standard family, the tested integers are split into blocks of m consecutive integers with one giant element per block, which is Harvey's realisation; for a general tested set the cost models are stipulated as accounting axioms – a relaxation, since a block of arbitrary integers need not be a translate of one common set of m baby offsets. A hit is the equality of a tested integer with the true $u_i(p, q)$.

Two cost models are used, with an integer block parameter $m \geq 1$. **M1 (materialised giants):** $T \geq m + \sum_i \lceil W_i/m \rceil$. **M2 (free giants):** $T \geq m + \Sigma_W/m$, the cost of a hypothetical full square-root speed-up.

Oblivious coverage is an assumption: a non-oblivious algorithm would have to use polylogarithmic information about N to exclude real sub-intervals of J containing no prime divisor of N. No such exclusion is known: the known-bits methods [17, 42] assume the excluded region given rather than computing it from N. A lower bound against genuinely non-oblivious algorithms would be a computational lower bound and is out of reach.

The integer convention matters. With $x_{ab} = 2\sqrt{abN}$ and $\delta_{ab} = \sqrt{N}/(4r\sqrt{ab})$ Harvey tests $T_{ab} = \{\lceil x_{ab} \rceil + j : 0 \leq j < \delta_{ab}\}$, so $W_{ab} = \lceil \delta_{ab} \rceil \geq 1$; at $r = N^{1/3}$ most real windows have width

below one, and the inequality $\Sigma_W \ge P$ would be false for the real widths. The *standard family* at parameter r is $\{(a,b) : a,b \ge 1,\ ab \le r\}$ with these windows.

2.1. **Thinned families.** Two published algorithms violate condition (3) as stated: Hittmeir's babystep–giantstep method [39] tests, for the Fermat form $p+q$, only the integers in the residue classes modulo 2^kQ, $Q = \prod_{2<\ell\le B} \ell$, compatible with $pq \equiv N$ (his sets $L_{N,m}$; his Theorem 5.5 gives $|L_{N,Q}| = O(Q\,2^{-\pi(B)} \log B)$), and the refinement [32] of Harvey's algorithm tests each of its cells on a single residue class modulo the square of a primorial. Both discard tested integers on the strength of arithmetic modulo small primes, not on the size of N. We make the mechanism explicit and examine what it can change.

Definition 2.2 (thinned family)**.** A *thinned family* consists of forms as in Definition 2.1, each carrying a *window* $\mathcal{W}_i$ of $\widetilde{W}_i$ consecutive integers with candidate set $\widetilde{S}_i := \bigcup_{n\in\mathcal{W}_i}(n-1, n+1)$ of measure $\widetilde{w}_i \in [\widetilde{W}_i, 2\widetilde{W}_i]$, the candidate sets satisfying the coverage condition (3), and testing $T_i = \mathcal{W}_i \cap \Lambda_i$ for a union Λ_i of residue classes modulo a common modulus $\mathfrak{q}$, of density $\theta_i := |\Lambda_i \bmod \mathfrak{q}|/\mathfrak{q}$; put $W_i := |T_i|$, $\theta := \min_i \theta_i$, $\Sigma_{\widetilde{W}} := \sum_i \widetilde{W}_i$ and $\Sigma_{\widetilde{w}} := \sum_i \widetilde{w}_i$. Since the tested integers need not be consecutive, clause (4) is replaced by the cost models directly: with a block parameter $h \ge 1$, $T \ge h + \sum_i \lceil W_i/h \rceil$ in M1 and $T \ge h + \Sigma_W/h$ in M2. If $\mathfrak{q} \le \widetilde{W}_i/2$ for every i, a window of $\widetilde{W}_i$ consecutive integers contains at least $\theta_i(\widetilde{W}_i - \mathfrak{q}) \ge \theta\widetilde{W}_i/2 \ge 1$ members of Λ_i, so $\Sigma_W \ge \theta\Sigma_{\widetilde{W}}/2$ and $\Sigma_W \ge P$.

Lemma 2.3 (thinning by the classes compatible with N)**.** *Let $\mathfrak{q}$ be squarefree and coprime to $N = pq$, let (a,b) be a form with $\gcd(a,b,\mathfrak{q}) = 1$, and let $\Lambda_{ab} \subseteq \mathbb{Z}/\mathfrak{q}\mathbb{Z}$ be the set of residues of $aq' + bp'$ over the pairs of units (p',q') modulo $\mathfrak{q}$ with $p'q' \equiv N$; it contains $aq + bp \bmod \mathfrak{q}$. Then $|\Lambda_{ab}| \ge \varphi(\mathfrak{q})/2^{\omega(\mathfrak{q})}$, so a thinned family of such forms whose classes are the compatible ones has density $\theta \ge \varphi(\mathfrak{q})/(\mathfrak{q}\,2^{\omega(\mathfrak{q})}) \ge \mathfrak{q}^{-c/\log\log\mathfrak{q}}$ for an absolute constant c and $\mathfrak{q} \ge 3$. For every thinned family with $\mathfrak{q}$ at most half the shortest window, of density θ, the floors of Corollary 3.3 and Theorems 3.4–3.6 below hold in the form $\Sigma_W \ge (\theta c_C/4)^{2/3} N^{1/3}$, $T \ge (\theta c_C)^{2/5} N^{1/5}$ in M1 and $T \ge 2(\theta c_C/4)^{1/3} N^{1/6}$ in M2; when the classes are the compatible ones and $\mathfrak{q} \le N$, $\theta \ge N^{-o(1)}$ and the floors hold up to a factor $N^{o(1)}$.*

Proof. For a prime $\ell \mid \mathfrak{q}$ the compatible pairs are (x, Nx^{-1}), $x \in \mathbb{F}_\ell^\times$. If $\ell \nmid ab$, the map $x \mapsto aNx^{-1} + bx$ takes each value v at most twice, at the roots of $bx^2 - vx + aN$, so its image has at least $(\ell-1)/2$ elements; if ℓ divides exactly one of a, b the map is aN/x or bx, a bijection of $\mathbb{F}_\ell^\times$; $\ell \mid \gcd(a,b)$ is excluded by hypothesis. The Chinese remainder theorem makes the image modulo $\mathfrak{q}$ the product of the local images, giving $|\Lambda_{ab}| \ge \prod_{\ell\mid\mathfrak{q}}(\ell-1)/2$. Since $\omega(\mathfrak{q}) \le (1+o(1))\log\mathfrak{q}/\log\log\mathfrak{q}$ and $\varphi(\mathfrak{q})/\mathfrak{q} \gg 1/\log\log\mathfrak{q}$, the density bound follows. For the floors, Lemma 3.2 below applied to the windows' candidate sets gives $\Sigma_{\widetilde{w}} \ge c_C\sqrt{N}/\sqrt{P}$, so $\Sigma_{\widetilde{W}} \ge c_C\sqrt{N}/(2\sqrt{P})$ and, by Definition 2.2, $\Sigma_W \ge \theta c_C\sqrt{N}/(4\sqrt{P})$; with $\Sigma_W \ge P$ this gives $\Sigma_W^{3/2} \ge \theta c_C\sqrt{N}/4$. In M1 every ceiling is at least one, so $T \ge P$, and $T^2 \ge 4\Sigma_W \ge \theta c_C\sqrt{N}/\sqrt{P} \ge \theta c_C\sqrt{N}/\sqrt{T}$, whence $T^{5/2} \ge \theta c_C\sqrt{N}$; in M2, $T \ge 2\sqrt{\Sigma_W}$. □

In the standard family the primitivity hypothesis costs nothing: if $(a,b) = d(a',b')$ covers a point then so does the reduced cell (a',b'), whose window is wider, so imprimitive cells may be discarded. Hittmeir's modulus 2^kQ has the even part 2^k, to which the lemma does not apply; it costs only a logarithmic factor.

Lemma 2.4 (the 2-adic part)**.** *For odd N and $k \ge 4$, the set $L_{N,2^k} := \{x + Nx^{-1} \bmod 2^k : x \text{ odd}\}$ has at least $2^{k-4}/(k+1)$ elements.*

Proof. Let $f(x) := x + Nx^{-1}$ on the 2^{k-1} odd residues and let E be the number of ordered pairs (x, y) with $f(x) = f(y)$; by Cauchy–Schwarz $|L_{N,2^k}| \geq 4^{k-1}/E$. Now $f(x) = f(y)$ iff $(x-y)(xy-N) \equiv 0 \pmod{2^k}$. The diagonal contributes 2^{k-1} pairs. For $x \neq y$ let $2^j \parallel x - y$, $1 \leq j \leq k-1$, so that $2^{k-j} \mid xy - N$, and write $x = y + 2^j t$ with t odd. If $j \geq k/2$ the condition reads $y^2 \equiv N \pmod{2^{k-j}}$, which holds for at most $4 \cdot 2^j$ odd residues y modulo 2^k, and then x ranges over at most 2^{k-j-1} values: at most 2^{k+1} pairs. If $j < k/2$ the condition reduces modulo 2^j to $y^2 \equiv N \pmod{2^j}$, satisfied by at most $4 \cdot 2^{k-j}$ odd y (all 2^{k-1} if $j \leq 2$), and then t is determined modulo 2^{k-2j} by $2^j ty \equiv N - y^2 \pmod{2^{k-j}}$, giving at most 2^j values of x: at most 2^{k+2} pairs. Hence $E \leq 2^{k-1} + (k-1)2^{k+2} \leq (k+1)2^{k+2}$ and $|L_{N,2^k}| \geq 4^{k-1}/((k+1)2^{k+2}) = 2^{k-4}/(k+1)$. □

With $\gcd(N, Q) = 1$, the Chinese remainder theorem and Lemma 2.3 for the Fermat form give $|L_{N,2^k Q}| \geq (2^{k-4}/(k+1))\,\varphi(Q)/2^{\pi(B)-1}$, so the 2-power costs at most a factor $16(k+1) = O(\log N)$ and the density remains $N^{-o(1)}$. An exponent gain attributable to the factor $2^{-\omega(\mathsf{q})}$ would need $2^{\omega(\mathsf{q})} \geq N^{\varepsilon}$, i.e. $\log \mathsf{q} \gg_\varepsilon \log N \log\log N$, a modulus far larger than the windows; scanning every window integer with a membership test then costs $\Omega(\Sigma_{\widetilde{W}})$, and materialising the class set by Chinese-remainder enumeration, as in [39, Algorithm 4.1], costs at least $|\Lambda| \geq \mathsf{q}^{1-o(1)} > N$, so neither known implementation yields an exponent gain, while an output-sensitive enumeration of $\mathcal{W}_i \cap \Lambda_i$ would escape this argument and is not known to us. Thinning by classes that are *not* the compatible ones – a single class per form, the form being chosen to make that class compatible with a hypothesis on p – is what [32] does; it is a member of the thinned model with $\theta = 1/\mathsf{q}$, for which the displayed floors are valid but weak, and Theorem 2.8 below gives its exact floor: $(\varphi(\mathsf{q})/\mathsf{q})^{3/5} N^{1/5}$, which their own analysis shows they attain up to logarithms (Proposition 2.9(iv)).

2.2. **Hypothesis thinning.** The Harvey–Hittmeir refinement does not thin by the classes compatible with N: for each unit class σ of p modulo a primorial q it chooses its forms so that the value $aq + bp$ falls into one predictable class modulo q^2 whenever $p \equiv \sigma$, and tests that class alone. Its saving is $(\varphi(\mathsf{q})/\mathsf{q})^{3/5}$. We show that this is the most the mechanism can buy inside the model: the requirement that one class serve every p of the hypothesis forces the forms into a sublattice of index q, and the covering sum on a sublattice is larger by exactly the factor that compensates the thinning.

Definition 2.5 (hypothesis-thinned family)**.** Let $\mathsf{q} \geq 2$ with $\gcd(\mathsf{q}, N) = 1$. A *hypothesis-thinned family* consists, for each unit class $\sigma \in (\mathbb{Z}/\mathsf{q}\mathbb{Z})^\times$, of a finite set F_σ of pairwise distinct forms, each form i carrying a window $\mathcal{W}_i$ of $\widetilde{W}_i \geq 1$ consecutive integers with candidate set $\widetilde{S}_i = (n_0 - 1, n_0 + \widetilde{W}_i)$ and covered set $\widetilde{S}(i) := \{p \in J_C : u_i(p, N/p) \in \widetilde{S}_i\}$, such that (1) for every σ the sets $\widetilde{S}(i)$, $i \in F_\sigma$, cover J_C; (2) each form tests $T_i := \mathcal{W}_i \cap \Lambda_i$ for a single residue class Λ_i modulo q^2, and $T_i \neq \emptyset$; (3) for every integer $p' \in \widetilde{S}(i)$ with $p' \equiv \sigma \pmod{\mathsf{q}}$ the residue $a_i N p'^{-1} + b_i p' \bmod \mathsf{q}^2$ lies in Λ_i. The cost is $T \geq h + \sum_\sigma \sum_{i \in F_\sigma} \lceil W_i/h \rceil$ with $W_i := |T_i|$ and a shared block parameter $h \geq 1$.

The family factors N: if $\sigma := p \bmod \mathsf{q}$, some $i \in F_\sigma$ has $aq + bp \in \widetilde{S}_i$ by (1), hence $aq + bp \in \mathcal{W}_i$ as it is an integer; p is a unit modulo q, so $q \equiv Np^{-1} \pmod{\mathsf{q}^2}$ and (3) puts $aq + bp$ in Λ_i, i.e. in T_i. Condition (2) excludes forms that cover reals but test nothing, without which the window coverage would be vacuous. Algorithm 4.3 of [32] is such a family: its forms satisfy $b \equiv aN\sigma^{-2} \pmod{\mathsf{q}}$ by construction (their (3.4)), which makes the residue of (3) constant along $p' \equiv \sigma$, its windows cover every real p of their interval (their (3.3)), and its babysteps $\beta^{\mathsf{q}^2 i}$ are shared by all forms; the accounting of Definition 2.5 charges one giant per form, as their Algorithm 4.3 does.

Lemma 2.6 (alignment)**.** *If $\widetilde{S}(i)$ contains two integers p' and $p'+\mathfrak{q}$ both congruent to σ modulo $\mathfrak{q}$, then $b_i \equiv a_i N\sigma^{-2} \pmod{\mathfrak{q}}$. A form that is not* aligned *in this sense has $|\widetilde{S}(i)| \le 4\mathfrak{q}$.*

Proof. Write $p' = \sigma + \mathfrak{q}t$ with σ a unit modulo $\mathfrak{q}^2$. Since $(\sigma+\mathfrak{q}t)(\sigma^{-1}-\mathfrak{q}t\sigma^{-2}) \equiv 1 \pmod{\mathfrak{q}^2}$, one has $a_i N p'^{-1} + b_i p' \equiv (a_i N\sigma^{-1} + b_i\sigma) + \mathfrak{q}t(b_i - a_i N\sigma^{-2}) \pmod{\mathfrak{q}^2}$, and likewise with $t+1$; both residues lie in the single class Λ_i, so $\mathfrak{q}(b_i - a_i N\sigma^{-2}) \equiv 0 \pmod{\mathfrak{q}^2}$. For the second statement, $\widetilde{S}(i)$ is the preimage of an interval under the strictly convex g_i, hence has at most two connected components; a component of length $> 2\mathfrak{q}$ contains two integers $p', p'+\mathfrak{q}$ of the class σ and would make the form aligned, so each component has length $\le 2\mathfrak{q}$. □

Lemma 2.7 (covering sum on a sublattice)**.** *Let γ be a residue modulo $\mathfrak{q}$, and let P pairwise distinct forms (a_i, b_i) with $b_i \equiv \gamma a_i \pmod{\mathfrak{q}}$, with measurable candidate sets of total measure Σ_w, cover a measurable $E \subseteq J_C$ of measure $\lambda|J_C|$, $0 < \lambda \le 1$. Then*

$$\Sigma_w \;\ge\; \frac{\lambda^2 |J_C|^2}{16\sqrt{N}\, A_{P,\mathfrak{q}}}, \qquad A_{P,\mathfrak{q}} := 2\sqrt{(C-1)P/\mathfrak{q}} + 3\ln(4P\mathfrak{q}) + C + 5.$$

Proof. As in Lemma 3.2, with $L := |J_C|$: the monotone forms cover at most w_i each and the wide ones at most $2\sqrt{w_i\sqrt{N}/a_i}$, so either $\Sigma_w \ge \lambda L/2$, which exceeds the claimed bound because $L \le \sqrt{N}$ and $A_{P,\mathfrak{q}} \ge 8$, or $\sum_{\text{wide}} \sqrt{w_i/a_i} \ge \lambda L N^{-1/4}/4$ and Cauchy–Schwarz gives $\Sigma_w \ge \lambda^2 L^2 N^{-1/2}/(16\sum_{\text{wide}} 1/a_i)$. For fixed a the wide forms of the family have b among the $\lceil Ca\rceil - a + 1 \le (C-1)a + 2$ integers of $[a, \lceil Ca\rceil]$ lying in one class modulo $\mathfrak{q}$, so at most $(C-1)a/\mathfrak{q}+3$ of them; hence for every integer $A \ge 1$, $\sum_{\text{wide}} 1/a_i \le (C-1)A/\mathfrak{q} + 3(1+\ln A) + P/A$. If $(C-1)P \le \mathfrak{q}$ take $A := P$: the sum is at most $5 + 3\ln P \le A_{P,\mathfrak{q}}$. Otherwise $\sqrt{P\mathfrak{q}/(C-1)} < P$; take $A := \lceil\sqrt{P\mathfrak{q}/(C-1)}\rceil \le P$, so that $(C-1)A/\mathfrak{q} \le \sqrt{(C-1)P/\mathfrak{q}} + (C-1)/\mathfrak{q}$, $P/A \le \sqrt{(C-1)P/\mathfrak{q}}$ and $\ln A \le \ln P$, whence $\sum_{\text{wide}} 1/a_i \le A_{P,\mathfrak{q}}$. □

Primitivity is not assumed and must not be arranged: dividing an aligned form by a common factor of its coordinates that is not a unit modulo $\mathfrak{q}$ can destroy the alignment, and the lemma only uses distinctness.

Theorem 2.8 (hypothesis thinning)**.** *Fix $C > 1$. There are $c_H = c_H(C) > 0$ and $N_0(C)$ such that for $N \ge N_0(C)$ and every $\mathfrak{q}$ with $2 \le \mathfrak{q} \le N^{1/20}$ and $\gcd(\mathfrak{q}, N) = 1$, every hypothesis-thinned family has*

$$T \;\ge\; c_H\Big(\frac{\varphi(\mathfrak{q})}{\mathfrak{q}}\Big)^{3/5} N^{1/5}.$$

Since $\varphi(\mathfrak{q})/\mathfrak{q} \gg 1/\log\log(3\mathfrak{q})$, hypothesis thinning saves at most a factor $O((\log\log N)^{3/5})$ over Theorem 3.4; for the primorials $\mathfrak{q} \in (N^{1/21}, N^{1/20})$ of [32], their algorithm attains the floor up to a factor $(\log N)^{O(1)}$.

Proof. Put $P_\sigma := |F_\sigma|$, $P_{\text{tot}} := \sum_\sigma P_\sigma$, $\varphi := \varphi(\mathfrak{q})$, $L := |J_C|$, $c_3 := (1-C^{-1/2})^2/(512\sqrt{C-1})$ and $c_H := c_3^{2/5}$, and let $T^* := c_H(\varphi/\mathfrak{q})^{3/5}N^{1/5}$. Since $T \ge P_{\text{tot}}$ we may assume $P_{\text{tot}} < T^* \le c_H N^{1/5}$. Fix σ. By Lemma 2.6 the non-aligned forms of F_σ cover at most $4\mathfrak{q}P_\sigma \le 4c_H N^{1/4} \le L/2$ for $N \ge N_0(C)$, so the aligned forms of F_σ – distinct forms in the sublattice $b \equiv N\sigma^{-2}a \pmod{\mathfrak{q}}$ – cover a set of measure at least $L/2$, and Lemma 2.7 with $\lambda = 1/2$ gives $\sum_{i\in F_\sigma} \widetilde{w}_i \ge L^2/(64\sqrt{N}A_{P_\sigma,\mathfrak{q}})$, where $\widetilde{w}_i = \widetilde{W}_i + 1$ is the measure of the window's candidate set. A window of $\widetilde{W}_i$ consecutive integers contains at least $\lfloor \widetilde{W}_i/\mathfrak{q}^2 \rfloor$ members of any class modulo $\mathfrak{q}^2$, and $W_i \ge 1$ by (2), so $W_i \ge \max(1, \widetilde{W}_i/\mathfrak{q}^2 - 1) \ge \widetilde{W}_i/(2\mathfrak{q}^2) \ge \widetilde{w}_i/(4\mathfrak{q}^2)$. Hence

$$\Sigma_W \;\ge\; \frac{L^2}{256\,\mathfrak{q}^2\sqrt{N}} \sum_\sigma \frac{1}{A_{P_\sigma,\mathfrak{q}}} \;\ge\; \frac{L^2\,\varphi}{256\,\mathfrak{q}^2\sqrt{N}\,A_{P_{\text{tot}}/\varphi,\mathfrak{q}}},$$

by Jensen, $x \mapsto 1/A_{x,\mathsf{q}}$ being convex (A is positive, increasing and concave). With $P_{\mathrm{tot}} < T^*$ the square-root term of $A_{P_{\mathrm{tot}}/\varphi,\mathsf{q}}$ is at most $X := 2\sqrt{C-1}\,c_H^{1/2}\varphi^{-1/5}\mathsf{q}^{-4/5}N^{1/10} \geq 2\sqrt{C-1}\,c_H^{1/2}N^{1/20}$ (as $\varphi^{1/5}\mathsf{q}^{4/5} \leq \mathsf{q} \leq N^{1/20}$), while its remaining terms are $3\ln(4P_{\mathrm{tot}}\mathsf{q}/\varphi) + C + 5 = O_C(\log N)$; so $A_{P_{\mathrm{tot}}/\varphi,\mathsf{q}} \leq (1+o(1))X$ and

$$\Sigma_W \;\geq\; (1-o(1))\,\frac{(1-C^{-1/2})^2 N\,\varphi}{256\,\mathsf{q}^2\sqrt{N}\,X} \;=\; (1-o(1))\,c_3\,c_H^{-1/2}\Big(\frac{\varphi}{\mathsf{q}}\Big)^{6/5}N^{2/5},$$

using $L^2 = (1-C^{-1/2})^2 N$. Therefore $T \geq 2\sqrt{\Sigma_W} \geq (1-o(1))\,2c_3^{1/2}c_H^{-1/4}(\varphi/\mathsf{q})^{3/5}N^{1/5}$, and $2c_3^{1/2}c_H^{-1/4} = 2c_3^{1/2-1/10} = 2c_H$, so $T \geq T^*$ for $N \geq N_0(C)$. □

The theorem prices the mechanism, not the enumeration: the $\varphi(\mathsf{q})$ hypotheses enter only through the form count P_{tot}, which is what [32] pay. Thinning by arbitrary classes of density q^{-2}, without the hypothesis condition, is bounded by Lemma 2.3 only by $(\mathsf{q}^{-2})^{2/5}N^{1/5}$; the sublattice forced by Lemma 2.6 is what raises the covering sum by $\sqrt{\mathsf{q}}$ and returns the exponent of the saving to $(\varphi(\mathsf{q})/\mathsf{q})^{3/5}$.

2.3. **Members.**

Proposition 2.9 (members of the model)**.** *(i) Lehman's algorithm* [49] *is simulated by an oblivious Lehman-cell BSGS algorithm with $m = 1$ whose candidate count exceeds Lehman's by at most the factor $\max_{k\leq r} d(k) = r^{o(1)}$: Lehman tests, for each product $k \leq r$, the integers U of the standard window of k for the square property of $U^2 - 4kN$; assigning each such test to every divisor pair (a,b) of k gives the standard family at $r = \lceil N^{1/3}\rceil$, which covers J_C by Harvey's covering lemma* [31, Lemma 3.3]*, proved for every real $\xi = p/q$ with $(N/r)^{1/2} \leq p < N^{1/2}$; the square test succeeds in particular at $U = aq + bp$, which is the hit of clause (4), and its other successes yield divisors of $4kN$ that Lehman's algorithm tests by a gcd. Hart's one-line factoring* [30] *tests the single integer $2\lceil\sqrt{kN}\rceil$ per product k and is the sub-family with that tested integer, assigned to the divisor pairs of k; Hart proves coverage for a restricted class of N, so it is a member of the model to the extent that its family covers J_C. (ii) Harvey's Algorithm 4.2* [31] *with parameters (r,m) is an oblivious Lehman-cell BSGS algorithm with materialised giants: its Step (2) forms one giant per block of m consecutive tested integers of each standard cell, its Steps (1), (3) and (4) realise the collision search, and its bit cost is the M1 cost up to $(\log N)^{O(1)}$* [31, Proposition 4.2]*; at $r = m = N^{1/5}$ it attains Theorem 3.4 up to logarithms. (iii) Hittmeir's Algorithm 8.1* [39, Theorems 7.1 and 8.3] *is a thinned family in the sense of Definition 2.2 with one form, the Fermat form $p+q$, whose window is the interval $[2\sqrt{N}, N^{1/2}+N/\Delta]$ of candidates for $p+q$ once his Strassen step has excluded the factors $p \leq \Delta$ (for his parameter $\Delta \in [N^{2/5}, N^{1/2}]$; the interval has length below N/Δ) and whose tested set is the intersection of that window with the classes $L_{N,2^kQ}$ compatible with N, and for $\Delta < \sqrt{N/C}$ that window covers J_C. Here $Q = \prod_{2<\ell\leq B}\ell$ with $B = \frac{1}{2}\log T$ for his bound T on $p+q$, so that $Q = T^{1/2+o(1)}$ and $2^{\omega(Q)} = N^{o(1)}$, and k is the least integer with 2^kQ at least the window length, so that the modulus lies between the window length and twice it and each class has at most one representative in the window. By Lemmas 2.3 and 2.4 the compatible classes have density $N^{-o(1)}$ modulo 2^kQ; but the hypothesis $\mathsf{q} \leq \widetilde{W}/2$ under which Definition 2.2 converts that density into a count of tested integers is not met, so the floors of Lemma 2.3 are not asserted for this member. Its cost is $N^{2/9+o(1)}$* [39, Theorem 8.3]*, above every floor proved here. (iv) The Harvey–Hittmeir refinement, Algorithm 4.3 of* [32] *with its Propositions 3.3 and 4.4, is a thinned family with materialised giants: for a primorial m coprime to N, an integer m_0, each unit class σ modulo m and each interval $[\sigma_0, (1+1/m_0)\sigma_0)$ of a geometric progression covering*

$[1, N^{1/2})$, *their Proposition 3.3 computes from* N *alone, by two-dimensional lattice reduction, a form* (a, b) *such that* $aq + bp$ *lies within* $4N^{1/2}\mathsf{m}^{1/2}/m_0^{3/2}$ *of* $aN/\sigma_0 + b\sigma_0$ *for* every *real* p *of the interval (their (3.3), which does not use the class hypothesis), and such that* $aq + bp$ *lies in one prescribed class modulo* m^2 *whenever* $p \equiv \sigma \pmod{\mathsf{m}}$ *(their (3.4)); the window of the form is the interval of* $2\lambda\mathsf{m}^2$ *consecutive integers around* $aN/\sigma_0 + b\sigma_0$, $\lambda = \lceil 4N^{1/2}/(\mathsf{m}m_0)^{3/2} \rceil$, *the tested set is its intersection with that single class* ($\theta_i = \mathsf{m}^{-2}$, $W_i = 2\lambda + 1$), *the* $2\lambda + 1$ *babysteps* $\beta^{\mathsf{m}^2 i}$ *are shared by all forms and one giant* β^{aN+b-j} *is formed per form, so that the block parameter is* $h = 2\lambda + 1$ *and the family has* $P = \varphi(\mathsf{m}) \cdot O(m_0 \log N)$ *forms; coverage of* J_C *holds because for every real* p *the* $\varphi(\mathsf{m})$ *windows of* p*'s interval all contain* $aq + bp$*; and the bit cost is the M1 cost* $h + P$ *up to* $(\log N)^{O(1)}$ *(their Proposition 4.4, whose bound is* $O(\mathsf{m} \lg N (\lg \lg N)^2 + \varphi(\mathsf{m}) m_0 \lg^4 N + N^{1/2} \lg^2 N/(\mathsf{m}m_0)^{3/2})$*). The thinning is by an enumerated hypothesis on* p *modulo* m*, not by the classes compatible with* N*, so Lemma 2.3's density bound does not describe it; the floors of Lemma 2.3 hold with* $\theta = \mathsf{m}^{-2}$*, and their own cost bound shows that the saving over Harvey's algorithm is the factor* $(\varphi(\mathsf{m})/\mathsf{m})^{3/5} \asymp (\log \log N)^{-3/5}$ *for every admissible choice of* m*: the exponent does not move, and Theorem 2.8 shows that no hypothesis-thinned family does better than this saving up to constants. (v) Hittmeir's time–space tradeoff* [40] *is not classified here. Pollard–Strassen* [68, 78] *is not a cell algorithm: it tests divisibility rather than a linear form; it belongs to the explicit difference-cover model of Section 4 through the additive mechanism, and Proposition 4.5 gives it the floor* $N^{1/4}$ *there.*

Proof. (i) and (ii) restate the algorithms: Lehman's window at r and Harvey's tested set T_{ab} are the standard windows; coverage of every real p in the range is the content of Harvey's Lemma 3.3 (Lehman's original lemma covers the same range); if $U = aq + bp$ then $U^2 - 4abN = (aq - bp)^2$; the multiplicity of the assignment is the divisor function of k, whose maximum over $k \le r$ is $r^{o(1)}$; Harvey's Step (2) computes $\lceil W_{ab}/m \rceil$ giants per cell and Step (1) the m babies, which is M1. (iii) Hittmeir's Algorithm 8.1 first excludes the factors $p \le \Delta$ by Strassen's method, so that $p + q < N^{1/2} + N/\Delta$, and then applies his Theorem 7.1 to the candidates s for $p + q$ in $[2\sqrt{N}, N^{1/2} + N/\Delta]$ that lie in $L_{N,2^kQ}$, testing each by his Lemma 2.1; the window is that interval, the tested set its intersection with the compatible classes, and the parameters $B = \frac{1}{2} \log T$ and $2^{k-1}Q < \alpha + 1 \le 2^kQ$, with $\alpha + 1$ the window length, are those of his proof of Theorem 7.1; $Q = e^{\vartheta(B)}/2 = T^{1/2+o(1)}$ by the prime number theorem, and $\omega(Q) = \pi(B) - 1 = O(\log N/\log \log N)$. (iv) is read off the cited statements as described; the coverage claim is their (3.3), which is proved from (3.1) alone. (v) is Proposition 4.5. □

Which step of the lower bound is not definitional. The cost models charge what the algorithms do: m babies and one giant per block in M1, which is Harvey's Step (2); the stipulated $m + \Sigma_W/m$ in M2. The content of the floors is elsewhere: Lemma 3.1 bounds the measure of real p covered by one cell in terms of the measure of its tested set, and the coverage clause (3) forces the tested sets to exhaust J_C; that geometry, not the accounting, yields the candidate floor $\Sigma_W \gg N^{1/3}$, and the two cost floors are its consequences.

3. The floors

Lemma 3.1 (coverage of one cell)**.** *Let* $g(p) = aN/p + bp$ *with* $a, b \ge 1$ *and let* $J' \subseteq (0, \sqrt{N}]$ *be an interval. Then* $g'' = 2aN/p^3 \ge \mu := 2a/\sqrt{N}$ *on* J'*, and for every measurable* $S \subset \mathbb{R}$ *of measure* w*,*

$$|\{p \in J' : g(p) \in S\}| \le 2\sqrt{2w/\mu} = 2\sqrt{w\sqrt{N}/a}.$$

Proof. Split J' at the critical point $p^* = \sqrt{aN/b}$ if it lies in J'; on each piece J_j the restriction $h = g|_{J_j}$ is strictly monotone onto an interval $h(J_j)$ with left endpoint u_0. The p-measure

of $\{h \in S\}$ is $\int_S \rho$ with $\rho(u) = 1/|h'(h^{-1}(u))|$ on $h(J_j)$ and 0 outside. Since $g'' > 0$, $|h'|$ increases along the direction of travel on each piece, so ρ is non-increasing in u on $[u_0, \infty)$, and $\int_S \rho \le \int_{[u_0,u_0+w)} \rho$, which is the length ℓ of $\{p \in J_j : u_0 \le h(p) < u_0 + w\}$, an interval starting at the endpoint x_0 of J_j where $h = u_0$. Taylor's theorem, with $h'' \ge \mu$ and $h'(x_0)$ pointing in the direction of travel, gives $w \ge |h(x_0 \pm \ell) - h(x_0)| \ge \mu\ell^2/2$, so $\ell \le \sqrt{2w/\mu}$; sum over the at most two pieces. □

The factor 2 is sharp: it is approached when the critical point $p^* = \sqrt{aN/b}$ lies inside J' and S is an interval at the minimum of g, both pieces then contributing equally. If $b \ge \lceil Ca \rceil + 1$ or $b \le a - 1$ then $|g'| \ge 1$ on J and the coverage is $\le w$ (the *monotone bound*). Both bounds hold for every measurable S, not only for unions of unit intervals around integers; this is used in Section 9.

Lemma 3.2 (covering sum)**.** *For an oblivious covering family of $P \ge 1$ distinct forms, $\Sigma_w \ge c_C\sqrt{N}/\sqrt{P}$. More generally, if the forms' measurable candidate sets S_i cover a measurable $E \subseteq J_C$ of measure $\lambda|J_C|$, $0 < \lambda \le 1$, then $\Sigma_w \ge \lambda^2 c_C\sqrt{N}/\sqrt{P}$.*

Proof. Call a form *wide* if $a \le b \le \lceil Ca \rceil$ and *monotone* otherwise; the monotone bound covers the latter and Lemma 3.1 every form. With $L = |J| = \sqrt{N}(1 - C^{-1/2})$, coverage gives $L \le \sum_{\text{mon}} w_i + \sum_{\text{wide}} 2\sqrt{w_i\sqrt{N}/a_i}$. If the first sum is $\ge L/2$ we are done. Otherwise $\sum_{\text{wide}} \sqrt{w_i/a_i} \ge LN^{-1/4}/4$, and by Cauchy–Schwarz this is $\le \sqrt{\Sigma_w}\sqrt{\sum_{\text{wide}} 1/a_i}$. For each a there are $\kappa_a := \lceil (C-1)a \rceil + 1 \le (C+1)a$ wide forms; with $n_a \le \kappa_a$ of them present, $\sum_a n_a \le P$ and any integer $A \ge 1$, $\sum_{\text{wide}} 1/a_i \le \sum_{a \le A} \kappa_a/a + P/A \le (C+1)A + P/A$. If $P \ge C+1$ take $A = \lceil \sqrt{P/(C+1)} \rceil$, so the sum is $\le 3\sqrt{(C+1)P}$ and $\Sigma_w \ge L^2/(48\sqrt{C+1}\sqrt{NP}) \ge c_C\sqrt{N}/\sqrt{P}$; if $P < C+1$ take $A = 1$, so the sum is $< 2(C+1)$ and $\Sigma_w \ge L^2/(32(C+1)\sqrt{N}) \ge c_C\sqrt{N} \ge c_C\sqrt{N}/\sqrt{P}$. The general statement is the same argument with L replaced by λL, every bound being quadratic in L. □

Discarding forms whose coverage misses J_C from a family covering a larger $J_{C'}$ leaves a covering family of J_C with no more forms or candidates, so the bounds below, proved for a fixed C, apply to every algorithm covering any larger range, in particular to Harvey's $[\sqrt{N/r}, \sqrt{N}]$ with $r \to \infty$.

Corollary 3.3 (candidate-count floor)**.** $\Sigma_W \ge \max(P, (c_C/2)\sqrt{N}/\sqrt{P}) \ge (c_C/2)^{2/3}N^{1/3}$.

Proof. $\Sigma_W \ge \Sigma_w/2$ and $\Sigma_W \ge P$; the maximum is minimised at $P^{3/2} = (c_C/2)\sqrt{N}$. □

This is the converse of Harvey's Remark 3.4: Lehman's parameter choice is optimal for every oblivious covering family.

Theorem 3.4 (M1)**.** *Every oblivious Lehman-cell BSGS algorithm with materialised giants has $T \ge c'_C N^{1/5}$ with $c'_C = 2^{4/5}(c_C/2)^{2/5}$.*

Proof. $T \ge \max(P, m + \Sigma_W/m) \ge \max(P, 2\sqrt{\Sigma_W}) \ge \max(P, 2\sqrt{c_C/2}\, N^{1/4}P^{-1/4})$; the two terms cross at $P^{5/4} = 2\sqrt{c_C/2}\, N^{1/4}$. □

Harvey's $r = m = N^{1/5}$ attains this up to logarithms: within the materialised-giant model, 1/5 is optimal. Chunked babystep–giantstep is covered, since $\sum_c 2\sqrt{W_c} \ge 2\sqrt{\Sigma_W}$ and the non-empty first block of every form still gives $T \ge P$.

Remark 3.5 (randomised coverage). If the family is sampled from a distribution such that every real $p \in J$ is covered with probability $\ge \delta$, then by Fubini the covered set has expected measure $\ge \delta L$, the event that it has measure $\ge \delta L/2$ has probability $\ge \delta/2$, and on that event the proofs

above go through with L replaced by $\delta L/2$ (Lemma 3.2 scales as δ^2). Hence the expected cost of the whole sampled family is $\geq c(\delta/2)^{9/5}N^{1/5}$ in M1 and $\geq c(\delta/2)^{5/3}N^{1/6}$ in M2: randomness inside the covering paradigm does not change the exponent. The bound charges the full family, not the stopping time of a sequential search.

Theorem 3.6 (M2, the $N^{1/6}$ floor)**.** *Even with free giants, $T \geq 2\sqrt{\Sigma_W} \geq 2(c_C/2)^{1/3}N^{1/6}$.*

Taking $r > N^{1/3}$ does not help, since then $\Sigma_W \geq P \gtrsim r$. The floor is exactly the square root of the minimum candidate count.

3.1. Unbalanced factors.

Proposition 3.7 (unbalanced factors)**.** *Fix $\beta \in (0, 1/2)$ and let $J_\beta := [x, 2x]$ with $x := N^\beta$, so that a factor $p \in J_\beta$ of $N = pq$ has $q/p \in [N^{1-2\beta}/4, N^{1-2\beta}]$. For every oblivious covering family of J_β with P forms, and N large,*

$$\Sigma_W \;\geq\; \max\Big(P,\; \frac{N^{1-\beta}}{512\,P}\Big) \quad \text{if } P < 3N^{1-2\beta}, \qquad \Sigma_W \;\geq\; \max\Big(P,\; \frac{N^{1/2}}{1340\sqrt{P}}\Big) \quad \text{if } P \geq 3N^{1-2\beta},$$

whenever the monotone forms cover less than half of J_β; if instead they cover at least half, then $\Sigma_W \geq \max(N^\beta/4,\; 3N^{1-\beta}/64)$. Consequently $\Sigma_W \gg N^{\max(1/3,\,(1-\beta)/2)}$; with materialised giants $T \gg N^{\max(1/5,\,(1-\beta)/3)}$; with free giants $T \gg N^{\max(1/6,\,(1-\beta)/4)}$. For $\beta \geq 2/5$ the three floors are those of Corollary 3.3 and Theorems 3.4, 3.6; the materialised-giant floor rises below $\beta = 2/5$, the candidate-count and free-giant floors below $\beta = 1/3$.

Proof. On J_β we have $g'' = 2aN/p^3 \geq aN/(4x^3)$, so the argument of Lemma 3.1 bounds the coverage of every form by $4\sqrt{2}\,\sqrt{wx^3/(aN)}$. Put $H := N/x^2 = N^{1-2\beta} > 1$. A form with critical point $p^* = \sqrt{aN/b} \notin [x/\sqrt{2}, 4x]$ is monotone on J_β with $|g'| \geq 1$: if $p^* < x/\sqrt{2}$ then $b > 2aH$ and $g'(p) \geq b - aH > aH \geq H > 1$; if $p^* > 4x$ then $b < aH/16$, which with $b \geq 1$ forces $aH > 16$, and $-g'(p) \geq aH/4 - b > 3aH/16 > 3$. So every monotone form has coverage at most both w and $16w/(3H)$ (the latter from $|g'| \geq 3aH/16 \geq 3H/16$ in the second case and $|g'| \geq H$ in the first), and if the monotone forms cover at least $x/2$ then $\Sigma_W \geq x/4$ and $\Sigma_W \geq 3xH/64 = 3N^{1-\beta}/64$, which exceeds every floor below. Otherwise the wide forms cover at least $x/2$; they have $b/a \in [H/16, 2H]$, at most $\kappa_a \leq 3aN^{1-2\beta}$ of them for each a, and $\sum_{\text{wide}} \sqrt{w_i/a_i} \geq \sqrt{N/x}/(8\sqrt{2}) = N^{(1-\beta)/2}/(8\sqrt{2})$; Cauchy–Schwarz gives $\Sigma_w \geq N^{1-\beta}/(128\sum_{\text{wide}} 1/a_i)$, with $\sum_{\text{wide}} 1/a_i \leq \min(P, 3N^{1-2\beta}A) + P/A$ for every integer $A \geq 1$: $A = 1$ gives $\leq 2P$, and for $P \geq 3N^{1-2\beta}$ the choice $A = \lceil\sqrt{P/(3N^{1-2\beta})}\rceil$ gives $\leq 3\sqrt{3N^{1-2\beta}P}$, whence $\Sigma_w \geq N^{1/2}/(666\sqrt{P})$. With $\Sigma_W \geq \max(P, \Sigma_w/2)$ this is the display. In the range $P < 3H$ the candidate count is least at $P = N^{(1-\beta)/2}/\sqrt{512}$, admissible for $\beta \leq 1/3$; if $\beta > 1/3$, every $P < 3H$ has $\Sigma_W > N^{1-\beta}/(1536H) = N^\beta/1536 \geq N^{1/3}/1536$. In the range $P \geq 3H$ the unconstrained balance $P = 1340^{-2/3}N^{1/3}$ is admissible for $\beta > 1/3$; if $\beta < 1/3$ then $P \geq 3H \geq 3N^{(1-\beta)/2}$. Hence $\Sigma_W \gg N^{\max(1/3,(1-\beta)/2)}$. For materialised giants, $T \geq \max(P, 2\sqrt{\Sigma_W})$: in the first range $T \geq \max(P, N^{(1-\beta)/2}P^{-1/2}/12)$, balanced at $P \asymp N^{(1-\beta)/3}$, admissible for $\beta \leq 2/5$, while for $\beta > 2/5$ and $P < 3H$ one has $T > N^{\beta/2}/21 \geq N^{1/5}/21$; in the second range $T \geq \max(P, 2N^{1/4}P^{-1/4}/\sqrt{1340})$, balanced at $P \asymp N^{1/5}$, admissible for $\beta \geq 2/5$, while for $\beta < 2/5$ one has $T \geq P \geq 3H \geq 3N^{(1-\beta)/3}$. Hence $T \gg N^{\max(1/5,(1-\beta)/3)}$. Free giants give $T \geq 2\sqrt{\Sigma_W} \gg N^{\max(1/6,(1-\beta)/4)}$. □

The balanced range itself is Section 3's J_C, and the proposition is stated for $\beta < 1/2$ because $[\sqrt{N}, 2\sqrt{N}]$ contains no smaller factor. Thus the exponent $1/5$ is not an artefact of balance: the floors are the same for every $\beta \geq 2/5$, which is exactly the range in which Harvey's algorithm is the record (for smaller β, Pollard–Strassen finds a factor $p \leq 2N^\beta$ in $N^{\beta/2+o(1)}$ operations, and

for the standard family a wide cell must have $b/a \asymp N^{1-2\beta}$, hence $r \gg N^{1-2\beta}$, so materialising the full hyperbolic family costs $N^{1-2\beta+o(1)}$, which Pollard–Strassen beats exactly when $\beta < 2/5$). In the gap $1/3 \le \beta < 2/5$ the model's floor $N^{(1-\beta)/3}$ lies below the cost $N^{1-2\beta}$ of the standard family, and the optimal exponent within the materialised-cell paradigm is not known there. Two further robustness statements are proved in Section 9: the M1 floor $N^{1/5}$ holds for oblivious families of nonconstant integer forms of every fixed total degree (Theorem 9.4), and the candidate floor holds against adaptive cell-membership queries (Proposition 9.5).

4. Difference covers

4.1. **The reduction.** Fix $\alpha \in (\mathbb{Z}/N\mathbb{Z})^\times$, and write $\mathrm{ord}_N(\alpha)$ for its multiplicative order modulo N. At $r = N^{1/3}$, $m = N^{1/6}$, let $c_{ab} = \lceil 2\sqrt{abN} \rceil$ and let $U \in T_{ab}$ be a tested integer. The true cell satisfies $\alpha^U \equiv \alpha^{aN+b} \pmod p$, i.e. $\alpha^z \equiv 1 \pmod p$ with $z := U - aN - b$. Let Z be the multiset $\{U - aN - b\}$ over all cells and tested integers, of size $\Sigma_W = N^{1/3+o(1)}$, and $\mathrm{supp}(Z)$ its support; its elements have $O(\log N)$ bits.

Proposition 4.1 (sufficiency)**.** *Suppose that* $\mathrm{supp}(Z) \subseteq B - G$ *for sets* $B, G \subset \mathbb{Z}$ *of integers of* $O(\log N)$ *bits, computable in* $\tilde{O}(|B| + |G|)$ *operations, that the elements of* B *are pairwise distinct modulo* $\mathrm{ord}_N(\alpha)$*, and that every integer* z *can be decoded in* $O(\log N)$ *operations into the* $O(1)$ *triples* (a, b, U) *with* $z = U - aN - b$*,* $U \in T_{ab}$*, that* z *may represent. Then there is a factoring search of cost* $\tilde{O}(|B| + |G|)$*.*

Proof. The modular distinctness of B is a hypothesis; in the usual situation in which B lies in an interval of length $< N^{1/3}$, it follows from an element of order $> N^{1/3}$, which the large-order routine of Harvey and Hittmeir [33, Theorem 1.1] supplies, or factors N, with $D = N^{1/3}$ at cost $N^{1/6+o(1)}$. Decoding holds for the Lehman set: writing $z = -aN + (U - b)$ with $0 < U - b < N/2$ determines $a = \lceil -z/N \rceil$ and $s = U - b$; for fixed a the map $b \mapsto c_{ab} - b$ is strictly increasing with increments $\ge \sqrt{N/(2r)} - 2$, larger than the window width $W_{ab} \le N^{1/6}/4 + 1$ for $N \ge 10^9$, so at most one b is compatible with s and is found by binary search. Compute the babies $\{\alpha^\beta\}_{\beta \in B}$ and the giants $\{\alpha^\gamma\}_{\gamma \in G}$, each by one modular exponentiation with an $O(\log N)$-bit exponent. Sort and match them: since the babies are pairwise distinct modulo N, each giant matches at most one baby, so there are at most $|G|$ exact matches modulo N; decode each as $z = \beta - \gamma$ and test the resulting triples by checking that $U^2 - 4abN$ is a perfect square. A matched giant may not simply be discarded: writing $x_\beta := \alpha^\beta$, a giant γ equal to x_{β_*} modulo N can still agree modulo p alone with a different baby x_β, and that agreement may be the one representing the true cell (for $N = 15$, $\alpha = 2$, $B = \{0, 2\}$, $G = \{0\}$ the giant 1 equals the baby x_0 and agrees with $x_2 = 4$ modulo 3). So, with $F(X) := \prod_{\beta \in B}(X - x_\beta)$, compute for every matched baby

$$h := \gcd\big(F'(x_{\beta_*}), N\big) = \gcd\Big(\prod_{\beta \ne \beta_*} (x_{\beta_*} - x_\beta),\ N\Big);$$

the values $F'(x_\beta)$ at all babies come from the product tree of F and one multipoint evaluation of F', in $\tilde{O}(|B|)$. If $1 < h < N$, return h. If $h = N$, scan the differences $x_{\beta_*} - x_\beta$, $\beta \ne \beta_*$: each is nonzero modulo N by distinctness, so none has gcd N with N, and since their product is divisible by N not all have gcd 1, so the scan returns a proper factor in $O(|B|)$ gcds. If $h = 1$, no baby other than x_{β_*} agrees with γ modulo p or modulo q, and the giant is removed. (The test is what a baby set certified distinct modulo N alone needs. In Harvey's Algorithm 4.1 the babies are the consecutive powers $\alpha^0, \dots, \alpha^{m-1}$ and its Step (1) either finds a factor or certifies $\mathrm{ord}_p(\alpha), \mathrm{ord}_q(\alpha) > m$ [31], so the babies are pairwise distinct modulo each prime and a second, one-sided agreement cannot occur; the order-selection lemma for explicit baby sets in Part 2 secures the same separation modulo both primes.) Run Harvey's collision algorithm [31,

Algorithm 4.1, Proposition 4.1] on the babies and the remaining giants: it forms the product tree of the babies, evaluates it at every giant, and, whenever a value has a nontrivial gcd with N, returns a factor – directly if the gcd is proper, and otherwise by scanning the babies once, since a giant that equals no baby modulo N but agrees with some baby modulo p has gcd exactly p with that baby. The total cost is $\tilde{O}(|B| + |G|)$. The pair representing $z_0 = (aq + bp) - aN - b$ for the true cell either matched exactly and was decoded, or its giant was matched exactly by another baby and the derivative test returned a factor, or its giant survived to the collision algorithm and agrees with its baby modulo p, yielding a gcd. □

Definition 4.2 (explicit-difference representation model)**.** The algorithm tests the candidates Z' of an oblivious covering family by materialising lists B, G and supplying, for every $z \in \operatorname{supp}(Z')$, a pair $(\beta_z, \gamma_z) \in B \times G$ with $z = \beta_z - \gamma_z$; correctness is certified through these represented differences, and the cost is $\geq |B| + |G|$.

In this model $\operatorname{supp}(Z') \subseteq B - G$ by definition. Minimal oblivious coverage alone does not imply the containment (a spurious collision $\beta - \gamma \equiv 0 \pmod{\operatorname{ord}_p \alpha}$ may certify a candidate without representing it), and Corollary 3.3 bounds the multiset size, not the support. Counting alone gives only $|B||G| \geq |\operatorname{supp}(Z')|$, which at $|B|, |G| \asymp \sqrt{|\operatorname{supp} Z'|}$ says nothing; the obstruction below uses the approximate differences among the window *starts*.

4.2. **The main lemma.** Let $Z' = \bigsqcup_{I \in \mathcal{F}} I$ be a disjoint union of R' integer windows $I = [s_I, s_I + W_I) \cap \mathbb{Z}$ with $1 \leq W_I \leq W$, $n' := \sum_I W_I$, and suppose $Z' \subseteq B - G$ with $|B| = K_1$, $|G| = K_2$. Define the *approximate-Sidon statistic*

$$D_{\max}(\mathcal{F}, W) := \max_{t \in \mathbb{Z} \setminus \{0\}} \#\{(I, I') \in \mathcal{F}^2 : I \neq I',\ |s_I - s_{I'} - t| < W\}.$$

Throughout, $|\mathcal{F}| \geq 2$; the windows being disjoint, their starts are distinct, so $D_{\max}(\mathcal{F}, W) \geq 1$. We call $D_{\max}(\mathcal{F}, 1)$, the largest multiplicity of a nonzero difference among the starts, the *exact statistic* of $\mathcal{F}$, and $D_{\max}(\mathcal{F}, W)$ for the family's own window length its *tolerance statistic*.

Lemma 4.3 (the cover lemma)**.** *If $K_1 \leq n'/(2W)$ then $K_1 K_2^2 \geq n'^2/(4W D_{\max})$; if $K_2 \leq n'/(2W)$ then $K_2 K_1^2 \geq n'^2/(4W D_{\max})$. Consequently*

$$|B| + |G| \ \geq\ \min\Big(\frac{n'}{2W},\ 2\Big(\frac{n'^2}{4W D_{\max}}\Big)^{1/3}\Big).$$

Proof. Choose for every $z \in Z'$ one representation $z = \beta_z - \gamma_z$ and form the bipartite graph H on $B \sqcup G$ with the n' edges $\{\beta_z, \gamma_z\}$ (H is simple since an edge (β, γ) determines the unique difference $z = \beta - \gamma$). Call a pair of distinct edges at a common baby β an *overlap*. By Jensen, $\sum_\beta \binom{d_\beta}{2} \geq n'^2/(2K_1) - n'/2$. An overlap whose edges lie in the same window I pairs two elements of I through one baby, so these number at most $\sum_I \binom{W_I}{2} \leq (W-1)n'/2$. An overlap with edges in $I \neq I'$ and giants $\gamma \neq \gamma'$ has $\beta \in (I + \gamma) \cap (I' + \gamma')$, which is non-empty only if $|s_I - s_{I'} - (\gamma' - \gamma)| < W$ and then has at most W elements; hence the cross-window overlaps number at most $W K_2^2 D_{\max}$. Altogether $n'^2/(2K_1) - Wn'/2 \leq W K_2^2 D_{\max}$, and if $K_1 \leq n'/(2W)$ the left side is $\geq n'^2/(4K_1)$. The second statement is the same count at the giants (two edges at a common giant have distinct babies, and $\gamma \in (\beta - I) \cap (\beta' - I')$); negating Z' does not reduce one case to the other, because the start differences change by $W_{I'} - W_I$ when the widths vary. If both $K_1, K_2 \leq n'/(2W)$, multiply the two inequalities and use $K_1 + K_2 \geq 2\sqrt{K_1 K_2}$. □

Harvey's cover (B an interval of length m, G the window starts plus multiples of m) satisfies the lemma with room; the bound bites when a cover uses fewer giants than windows, which forces approximate coincidences among the starts. We refer to the right side of the lemma as *the bound of Lemma 4.3* for the sub-family $\mathcal{F}$; its second branch is the *Sidon branch*.

4.3. **From an envelope to a floor.** For the cells $(1,k)$ write $s_k = \lceil 2\sqrt{kN}\rceil - N - k$ for the first tested exponent, $\mathcal{F}_1(r) = \{(1,k) : r/2 < k \le r\}$ for the *full $a=1$ shell*, $\mathcal{F}_1^{\square}(r)$ for its squarefree members, and $W(N,\rho) := \max_{\rho/2<k\le\rho} W_{1k} = \lceil W_{\mathrm{real}}(\rho)(1+O(1/\rho))\rceil$ for the lemma's window on the shell at ρ, with

$$W_{\mathrm{real}}(\rho) := \frac{\sqrt{N}}{2\sqrt{2}\,\rho^{3/2}};$$

we write W for $W(N,r)$ and W_{real} for $W_{\mathrm{real}}(r)$ when r is fixed.

Theorem 4.4 (envelope to floor)**.** *In the explicit-difference representation model, for the standard family at parameter r and all sufficiently large N, writing $D_\rho := D_{\max}(\mathcal{F}_1^{\square}(\rho), W(N,\rho))$:*

(1) *if $r \le N^{1/5}$: $|B|+|G| \ge 2^{-1/2}N^{1/4}r^{-1/4} \ge 2^{-1/2}N^{1/5}$;*
(2) *if $N^{1/5} \le r \le N^{1/3}$: $|B|+|G| \ge \min\big((3/(4\pi^2)-o(1))r,\ (0.253-o(1))N^{1/6}r^{1/6}D_r^{-1/3}\big)$;*
(3) *if $r > N^{1/3}$: $|B|+|G| \ge (0.569-o(1))u^{2/3}D_u^{-1/3}$ with $u = \lfloor N^{1/3}\rfloor$, hence $\gg N^{2/9}D_u^{-1/3}$.*

If $D_x \le x^\eta$ for all $N^{1/5} \le x \le N^{1/3}$ with a fixed $0 \le \eta < 1/2$, then $|B|+|G| \ge c\,N^{(1-\eta)/(5-4\eta)}$ for every r.

Proof. In the shell $\{(1,k) : u/2 < k \le u\}$ consecutive starts differ by $\approx \sqrt{N/k} \gg W$ for $u \le N^{1/3}$, so the windows are disjoint and $|B||G| \ge n'$ for each sub-family. (i) The shell at $u = r$ has $\ge r/2$ cells with $W_{1k} \ge \sqrt{N}/(4r\sqrt{r})$, so $n' \ge \sqrt{N}/(8\sqrt{r})$ and $|B|+|G| \ge 2\sqrt{n'}$. (ii) $\mathcal{F}_1^{\square}(r)$ has $R' = (3/\pi^2+o(1))r$ windows of integer widths $\ge W/2$, so $n' \ge R'W/2$; Lemma 4.3 gives $n'/(2W) \ge R'/4$ and $2(n'^2/(4WD_r))^{1/3} \ge 2(R'^2W/(16D_r))^{1/3} \ge (0.253+o(1))N^{1/6}r^{1/6}D_r^{-1/3}$ using $W \ge (1-O(1/r))W_{\mathrm{real}}$. (iii) The fixed shell at u is contained in the family and has all windows of width 1 at parameter $r \ge u$, i.e. $W(N,u) = 1$; its start differences do not depend on r; Lemma 4.3 with $W = 1$, $n' = R' = (3/\pi^2+o(1))u$ gives $2(R'^2/(4D_u))^{1/3} = (2^{1/3}(3/\pi^2)^{2/3}+o(1))u^{2/3}D_u^{-1/3}$, and this is the smaller branch of the minimum, since $D_u \ge 1$ gives $2(R'^2/4)^{1/3} < R'/2$ as soon as $R' > 16$. For the consequence, the counting bound of (i) holds for every $r \le N^{1/3}$. In (ii) the envelope gives the Sidon branch $\gg N^{1/6}r^{(1-2\eta)/6}$, and the other branch of the minimum, $\gg r$, is also $\gg N^{1/6}r^{(1-2\eta)/6}$ for $r \ge N^{1/5}$ and $\eta \ge 0$, so (ii) gives $|B|+|G| \gg N^{1/6}r^{(1-2\eta)/6}$. This and the counting bound cross at $r = N^{1/(5-4\eta)}$, where both have order $N^{(1-\eta)/(5-4\eta)}$, and case (iii), under the envelope at $x = u$, is $\ge N^{(2-\eta)/9}$, which exceeds that exponent by $(1-2\eta)^2/(9(5-4\eta)) \ge 0$. □

The exponent is $1/5$ at $\eta = 0$, $2/11$ at $\eta = 1/3$ and the floor $1/6$ at $\eta = 1/2$. The sub-family to which Lemma 4.3 is applied is a free choice: thinning by a factor f costs $f^{2/3}$ in the Sidon branch, so a polylogarithmic thinning costs $N^{o(1)}$, and the exponent obtained is a supremum over admissible sub-families.

4.4. **The additive mechanism.** The same incidence count applies to divisibility tests, and shows that Pollard–Strassen is optimal in the difference-cover model for them: the tested set is polynomially bounded, contains a multiple of every prime in J_C, is represented inside a difference set $Y - D$ as in Pollard–Strassen, and is charged cost at least $|Y|+|D|$.

Proposition 4.5 (the additive mechanism)**.** *Fix $C > 1$ and $K \ge 1$. Let $H \subset [1, N^K]$ be a set of integers such that every prime $p \in J_C$ divides some element of H, and suppose $H \subseteq Y - D$ for finite sets $Y, D \subset \mathbb{Z}$. Then*

$$|Y|+|D| \ \ge\ 2\sqrt{|H|}, \qquad |H| \ \ge\ \frac{(1-C^{-1/2}-o(1))\,2\sqrt{N}}{(2K+1)\log N},$$

so that $|Y|+|D| \gg_{C,K} N^{1/4}/\sqrt{\log N}$.

Proof. For $h \in H$ let $\Omega_{J_C}(h) := \sum_{p \in J_C \text{ prime}} v_p(h)$. Since every prime in J_C is at least $\sqrt{N/C}$, $(N/C)^{\Omega_{J_C}(h)/2} \le h \le N^K$, hence $\Omega_{J_C}(h) \le 2K/(1-\log C/\log N) \le 2K+1$ for all N sufficiently large in terms of C and K. Count the incidences $\mathcal{I} := \{(p,h) : p \in J_C \text{ prime},\ h \in H,\ p \mid h\}$: every prime of J_C occurs in at least one, and each h occurs with at most $2K+1$ distinct primes of J_C, so $\pi(J_C) \le |\mathcal{I}| \le (2K+1)|H|$. By the prime number theorem $\pi(J_C) = (1 - C^{-1/2} + o(1))\, 2\sqrt{N}/\log N$, which gives the bound on $|H|$. Finally $H \subseteq Y - D$ implies $|H| \le |Y|\,|D|$, and $|Y| + |D| \ge 2\sqrt{|Y|\,|D|} \ge 2\sqrt{|H|}$. □

Put $k_0 := \lfloor \sqrt{N} \rfloor$. Pollard–Strassen takes an integer $B \asymp N^{1/4}$, $H = \{1, \dots, k_0\}$, $D = \{0, \dots, B-1\}$ and $Y = \{jB : 1 \le j \le \lceil k_0/B \rceil\}$; then $H \subseteq Y - D$, and $\prod_{d \in D}(y-d) = P_B(y-B)$ with $P_B(X) := \prod_{i \le B}(X+i)$, so the $|Y|$ block products are the values of one fixed polynomial, computed by multipoint evaluation in $\tilde{O}(|Y| + |D|)$ ring operations. The proposition concerns divisibility by integers of polynomial size; a test of the form $p \mid \alpha^m - 1$, whose integer is exponentially large, is outside it (Section 7). The two mechanisms differ in their coverage: a divisibility test $p \mid h$ with $h \le N^K$ accounts for at most $2K+1$ primes of J_C, whereas Lemma 3.1 allows a cell with tested-set measure w to cover p-measure up to $2\sqrt{w\sqrt{N}/a}$; this square-root gain in coverage, after balancing the number of cells against the materialised-giant cost, is the source of the improvement from exponent $1/4$ to exponent $1/5$. In both, the realisation cost is a difference-cover cost, and the only unpriced lever is the algebraic evaluation of the product over the tested set as a single object (Section 7).

5. The start set

Since $aN + b$ is an integer,

$$aN + b - \lceil 2\sqrt{abN} \rceil = \left\lfloor aN + b - 2\sqrt{abN} \right\rfloor = \left\lfloor (\sqrt{aN} - \sqrt{b})^2 \right\rfloor :$$

the (negated) start exponents of the standard family are the floors of squared coordinate differences between the two square-root grids, the window offsets adding integers. Simultaneous negation of all starts does not change $D_{\max}$, since it sends the parameter t to $-t$. The statistic of Lemma 4.3 on the natural families is the subject of Part 3. Its conclusions, which the results list of Section 1 quotes, are these. In the balanced family $a \le b \le Ca$ the statistic is $\gg \sqrt{r}$ along rays (Lemma 19.1) and along pairs of equal product (Lemma 19.2); after both are excised by keeping one cell per product, half-offset lines force $D_{\max} \ge rN^{-1/4}/15 - 1$ for $47N^{1/4} \le r \le N^{1/3}$ and $C \ge 2$ in every family retaining the constructed cells (Lemma 19.3), so a bounded statistic is impossible for these families; under an envelope hypothesis the conditional floor is $N^{5/28}$ (Conjecture 19.4, Proposition 19.5) and under a thinned-statistic hypothesis $N^{2/11}$ (Conjecture 19.6, Corollary 19.7). For the $a = 1$ shell, whose starts are points moving linearly in $u = 2\sqrt{N}$, the statistic is bounded below by a multiple of $r^{1/3}$ in the modulus-free range $r \le N^{1/5-\varepsilon}$ and by a multiple of $N^{1/12}$ at $r = \lfloor N^{1/3} \rfloor$ (for the squarefree shell, for almost all N), and above by $2.1\, r^{2/3}$ and $6N^{2/9}$ (Sections 20 and 21); a return to $N^{1/5}$ within the method must come from a thinner sub-family, whose statistic is not controlled.

6. Measurements

The direct cover search of Section 23.5 and the balanced census of Section 23.6, both in Part 3, are exact integer computations on the moduli of Table 3: balanced semiprimes $N = pq$ with $q/p < \sqrt{2}$ generated by rejection sampling of two random primes from a seeded generator (seed 7 for the cover search, seed 5 for the balanced census). The product-modality sample of Section 7.1 uses 80 further moduli of 34 bits from seed 21. The exactness status of every other

computation of Part 3 – floating-point sweeps, boundary brackets and 200-bit recomputations among them – is stated where it is reported. The data and the code that regenerate every number are in the repository [23].

bits	N (Section 23.5)	bits	N (Section 23.6)
22	3 302 401	30	880 634 351; 964 728 493
24	10 142 569	32	3 722 008 519; 3 601 086 119
26	46 450 753	34	16 577 631 001; 14 048 619 007
28	164 541 959	36	46 551 225 037; 51 465 965 107
30	713 555 207	38	197 762 968 751; 187 452 546 833
32	3 032 841 269	40	876 190 896 151; 680 093 802 697
34	11 085 944 759	42	3 890 830 613 443; 3 072 789 281 321
36	42 281 311 687	44	11 427 345 574 369; 16 247 581 009 873
		46	55 862 770 399 391; 46 686 914 783 221
		48	176 552 314 063 291

TABLE 3. The moduli of the direct cover search and of the balanced census. Where two moduli are listed they are reported in that order throughout.

The cover search on the $a = 1$ shell and the census of the balanced sub-family are reported in Sections 23.5 and 23.6 of Part 3; the sample of Section 7.1 is reported there.

7. Escapes

None of the floors above is a lower bound on factoring; each is a theorem inside a stated model, and any faster algorithm must leave the model of whichever floor it evades. Table 4 lists the operations the models forbid, why, and what is known about each.

excluded operation	why the floor does not apply	status
non-oblivious exclusion of parts of J_C from N	violates coverage clause (3)	no such exclusion known; thinning by the classes compatible with N is worth $N^{o(1)}$ (Lemma 2.3), and thinning by an enumerated hypothesis on p modulo q at most $(\varphi(\mathsf{q})/\mathsf{q})^{3/5} = O((\log\log N)^{-3/5})$ (Theorem 2.8), which [32] attain; excluding an interval is worth exactly what knowing its complement's leading bits is worth (Part 2), and an arbitrary excluded set is not priced
non-linear or non-polynomial tests	the model has linear forms	fixed total degree: same M1 floor (Theorem 9.4); smoothness tests are outside the model
free giants without an explicit cover	M2 stipulates $m + \Sigma_W/m$; Lemma 4.3 prices explicit covers only	an implicit exact-product evaluation (Proposition 7.2) would leave both; no evaluator of exponent $< 3/5$ is known
shared computation across pieces	Theorem 10.1 charges pieces separately	open; separately evaluated chirp chains, rays and square arrays cost $N^{1/4-o(1)}$
Coppersmith-type lattice methods	outside the model	the univariate method [17, 42] runs in polynomial time once one factor is located in an interval of length $O(N^{1/4})$, i.e. once about $\frac{1}{4}\log_2 N$ of its leading bits are known; no necessity is proved
divisibility tests by integers of polynomial size	tested objects are divisors, not linear forms	Pollard–Strassen is optimal in the difference-cover model for them (Proposition 4.5)
smooth-order and elliptic-curve methods	test $p \mid \alpha^m - 1$, an exponentially large integer, or a group order	outside the covering and explicit-difference models of Part 1 (fixed-list elliptic curves are treated in Part 5); no floor is proved; the algebraic evaluation of $\prod_{i\le k}(X+i)$ below exponent $1/2$ is open

TABLE 4. What the models forbid, and what is known about each escape.

7.1. **The product modality.** Difference covers are not the only way to batch the cells. For a cell (a,b) put $c_{ab} := \lceil 2\sqrt{abN}\rceil$ and $T_{ab} := \{c_{ab}+i : 0 \le i < W_{ab}\}$. If the true value $U = aq + bp$ belongs to T_{ab} then, for every $\alpha \in (\mathbb{Z}/N\mathbb{Z})^\times$, $\alpha^U \equiv \alpha^{aN+b} \pmod p$. Hence the standard family at radius r gives the batched product

$$\Pi_r(\alpha) \;:=\; \prod_{ab\le r}\ \prod_{0\le i<W_{ab}} \left(\alpha^{c_{ab}+i} - \alpha^{aN+b}\right) \bmod N,$$

and $p \mid \Pi_r(\alpha)$ whenever the family covers p. If $\operatorname{supp}(\mathcal{T}) \subseteq B - G$ for the multiset $\mathcal{T}$ of shifted exponents, the Cartesian-product overproduct $Q_{B,G} := \prod_{b\in B, g\in G}(\alpha^{b-g}-1)$ is computable in $\tilde{O}(|B|+|G|)$ ring operations by a product tree and multipoint evaluation, and contains the collision factors; this overproduct realisation is what Lemma 4.3 and Theorem 4.4 concern. But $\Pi_r(\alpha)$ is an element of $\mathbb{Z}/N\mathbb{Z}$ with a definite algebraic description, and any straight-line program for it is a batching that is not a cover. Its evaluation faces a structural obstacle.

Lemma 7.1 (mirror symmetry)**.** *The standard family is invariant under $(a,b) \mapsto (b,a)$, and c_{ab}, W_{ab} depend only on ab. For every α, $\alpha^{ap+bq} \equiv \alpha^{aN+b} \pmod q$, and the window of (a,b) contains $ap+bq$ iff the cell (b,a) covers p. Consequently, if the full standard family covers p then $N \mid \Pi_r(\alpha)$; in particular, if $p \in J_C$, $p<q$ and $r \ge C$, then $\gcd(\Pi_r(\alpha), N) = N$.*

Proof. Fermat's congruence gives $\alpha^q \equiv \alpha \pmod q$, so $\alpha^{ap+bq} \equiv \alpha^{ap+b} \equiv \alpha^{aN+b} \pmod q$. The tested windows of (a,b) and (b,a) coincide, and $ap+bq$ is the value of the cell (b,a); if an off-diagonal cell (a,b) covers p, its factor vanishes modulo p and the corresponding factor of (b,a) vanishes modulo q, while a diagonal cell's own factor vanishes modulo both. For $p \in J_C$ and $r \ge C$, Harvey's Lemma 3.3 guarantees coverage. □

For the product modality we therefore use the *half-family* $a<b$. If a standard cell (a,b) covers p at $r = \lfloor N^{1/3} \rfloor \ge 2$, then with $m := |aq-bp|$ one has $m^2 < 6N/r$ and $|a/b - p/q| < 3/(b\sqrt{r})$ by Proposition 11.2(i); so a covering cell with $a \ge b$ forces $q-p < 3q/\sqrt{r}$, and the half-family covers every $p \in J_C$ with $q-p \ge 3q/\sqrt{r}$; the excluded moduli satisfy $(\sqrt{q}-\sqrt{p})^2 \le (\frac{9}{4}C^{3/2}+o(1))N^{1/6}$ and are found by $O_C(N^{1/6})$ Fermat square tests. Any remaining mod-q zero of the half-family product $\Pi_r^{<}(\alpha)$ is a non-mirrored coincidence. In a sample of 80 balanced 34-bit semiprimes ($q/p < \sqrt{2}$, seed 21) with $\alpha = 2$ and $r = \lfloor N^{1/3} \rfloor$, call a factor a *geometric hit* when its tested integer equals $aq+bp$, a *mirrored hit* when it equals $ap+bq$ (a zero modulo q by Lemma 7.1), and every other zero modulo p or q *coincidental*. The full family (mean 19 396 factors) had equal numbers of geometric and mirrored hits in every modulus (mean 2.53), on average 4.78 coincidental zeros modulo p and 2.55 modulo q, and $\gcd = N$ throughout. The half-family (mean 9 658 factors) had $\gcd = p$ in 53 moduli and $\gcd = N$ in 20 – there through its coincidental zeros modulo q, of which it has 1.24 per modulus on average over all 80, 4.85 on average over those 20, and at most 34 – and no zero modulo p in 7, the moduli with $q-p \in [626, 2070]$, which are inputs for the Fermat preprocessing.

Say that the half-family products have *uniform evaluation exponent at most* γ if, for each N, a deterministic generator given N streams a division-free straight-line program A_N, with one ring input X, the constants $0, \pm 1$ and the operations $+, -, \times$, such that $A_N(x) = \Pi^{<}_{\lfloor N^{1/3} \rfloor}(x)$ in $\mathbb{Z}/N\mathbb{Z}$ for every x, the program length is $\Sigma_W^{\gamma+o(1)}$, and the generator runs in that length times $(\log N)^{O(1)}$ bit operations.

Proposition 7.2 (the product dial)**.** *Fix $C>1$, let $N=pq$ with primes $p<q$ and $p \in J_C$, and put $r = \lfloor N^{1/3} \rfloor$. Suppose the half-family products have uniform evaluation exponent at most γ. Run $O_C(N^{1/6})$ Fermat square tests, which factor every modulus with $q-p < 3q/\sqrt{r}$. Suppose the input survives, and that an element α is given or computed at cost T_α with $1 < \gcd(\Pi_r^{<}(\alpha), N) < N$. Then N is factored with $N^{\gamma/3+o(1)}$ ring operations, plus T_α, the uniformity overhead and one gcd; in all $O_C(N^{1/6+o(1)} + N^{\gamma/3+o(1)} + T_\alpha)$. Subject to the separation hypothesis, $\gamma = 1/2$ gives an $N^{1/6+o(1)}$ algorithm and $\gamma = 1$ the Lehman exponent.*

Proof. If the preprocessing does not factor N, some cell with $a<b$ covers p, so $p \mid \Pi_r^{<}(\alpha)$; the half-family has $\Sigma_W = N^{1/3+o(1)}$ factors and the program evaluates its product in $N^{\gamma/3+o(1)}$ ring operations, so a proper gcd yields p. □

The separation hypothesis is the proposition's weakest point. It reduces to an evaluation hypothesis for sub-products, a width factor, and an explicit exceptional class of moduli.

Proposition 7.3 (isolation by bisection)**.** *Order the factors of $\Pi_r^{<}(X)$ with the cells lexicographic and the tested exponents increasing within a cell, and suppose that interval sub-products have uniform evaluation exponent at most γ: for every interval J of this order a*

deterministic generator streams a division-free straight-line program of length $|J|^{\gamma+o(1)}$ computing $\prod_{(a,b,U)\in J}(X^U - X^{aN+b})$. Let $N = pq$ with $p \in J_C$, $N \geq N_0(C)$, survive the Fermat preprocessing, and let (a_0, b_0) be a cell furnished by Harvey's Lemma 3.3, so that $a_0 < b_0$; put $\mathcal{H}_0 := \{x \in \mathbb{F}_q^\times : x^{(p-1)(a_0q-b_0)} = 1\}$, of order $\gcd((p-1)(a_0-b_0), q-1)$. For every unit $\alpha \notin \mathcal{H}_0$, breadth-first bisection of the interval sub-products factors N in $(1+Z_\alpha)N^{\gamma/3+o(1)}$ ring operations, where Z_α is the number of factors of $\Pi_r^{\leq}(\alpha)$ vanishing modulo p or modulo q. Unless

$$\gcd(p-1, q-1) \;\geq\; \frac{N^{5/24}}{2\sqrt{C}\,(\log N)^3},$$

some $\alpha \leq (\log N)^4$ lies outside $\mathcal{H}_0$, so trying $\alpha = 2, 3, \ldots$ factors N within $(\log N)^4$ candidates at total cost $(1+Z_*)N^{\gamma/3+o(1)}$, where $Z_* := \max Z_\alpha$ over the units $\alpha \leq (\log N)^4$ that are tried. The subgroup $\mathcal{H}_0$ depends on the unknown p, q, a_0, b_0 and is not available to the algorithm.

Proof. For an interval J write $\Pi_J := \prod_J(\alpha^U - \alpha^{aN+b})$; then $\gcd(\Pi_J, N)$ is p if J contains a factor vanishing modulo p and none vanishing modulo q, q symmetrically, N if both kinds occur, and 1 if neither. Explore the binary subdivision of the full order breadth first, computing at each node the gcd of its product, and stop as soon as a node has gcd p or q. Every node retained for subdivision has gcd N and hence contains a vanishing factor; retained nodes at one level are disjoint, so at most Z_α are retained and at most $2Z_\alpha$ children are evaluated at the next level; the depth is $O(\log N)$ and every interval has length at most $\Sigma_W = N^{1/3+o(1)}$, so all evaluations cost $(1+Z_\alpha)N^{\gamma/3+o(1)}$. If the exploration ends without a proper gcd, every vanishing factor vanishes modulo both primes; in particular the true factor $\alpha^{a_0q+b_0p} - \alpha^{a_0N+b_0}$ vanishes modulo q, i.e. $\alpha^{(p-1)(a_0q-b_0)} \equiv 1 \pmod q$, since $a_0N + b_0 - a_0q - b_0p = (p-1)(a_0q-b_0)$; that is $\alpha \in \mathcal{H}_0$, and $|\mathcal{H}_0| = \gcd((p-1)(a_0q-b_0), q-1) = \gcd((p-1)(a_0-b_0), q-1)$ because $a_0q \equiv a_0 \pmod{q-1}$. If every $\alpha \leq y := (\log N)^4$ lies in $\mathcal{H}_0$ then every y-smooth integer below q does, so $\Psi(q-1, y) \leq |\mathcal{H}_0|$, while $\Psi(q-1, y) \geq (q-1)^{3/4}/(\log N)^3$ for N large (the products of $u := \lfloor \log(q-1)/\log y \rfloor$ distinct primes $\leq y$ are distinct y-smooth integers below q, and $\binom{\pi(y)}{u} \geq (\pi(y)/u)^u \geq (\log N)^{3u} \geq (q-1)^{3/4}/(\log N)^3$); hence some $\alpha \leq y$ lies outside $\mathcal{H}_0$ unless $|\mathcal{H}_0| \geq (q-1)^{3/4}/(\log N)^3$. In that case $|\mathcal{H}_0| \leq \gcd(p-1, q-1)\,|a_0 - b_0| < \gcd(p-1, q-1)\,b_0$, and Proposition 11.2(ii) (Section 11), applied to the Lemma 3.3 cell, gives $b_0 < \sqrt{rq/p} + C^{1/4}/(2r^{1/4}) + 1/(8r) \leq \sqrt{Cr} + 1$ for N large; so with $r = \lfloor N^{1/3} \rfloor$ and $q \geq \sqrt{N}$,

$$\gcd(p-1, q-1) \;\geq\; \frac{(q-1)^{3/4}}{b_0(\log N)^3} \;\geq\; \frac{(\sqrt{N}-1)^{3/4}}{(\sqrt{C}\,N^{1/6}+1)(\log N)^3} \;\geq\; \frac{N^{5/24}}{2\sqrt{C}\,(\log N)^3}$$

for N large. □

In the sample above the half-family had on average about six zeros modulo p or q per modulus for $\alpha = 2$, the coincidental zeros modulo q numbering at most 34; nothing deterministic is proved about Z_α beyond $Z_\alpha \leq \Sigma_W$. The exceptional common-factor class is handled below $N^{1/5}$ from N alone by a separate argument: every balanced N with $\gcd(p-1, q-1) \geq N^{5/24}/(2\sqrt{C}(\log N)^3)$ is factored in $N^{7/48+o(1)}$ operations up to $\gcd(p-1, q-1) \leq (p-1)^{3/4}/(\log N)^3$, by a collision search on α^{N-1}, and in $N^{1/(8\sqrt{\mathrm{e}})+\varepsilon}$ beyond, by Burgess's bound on the least element outside a proper subgroup (Lemma 26.2 and Corollary 27.1 of Part 4). The conditional part of the product modality is thereby narrowed to the interval-evaluation hypothesis and the width bound $Z_* = N^{o(1)}$.

The same square-root pattern appears in Pollard–Strassen, where the batched object is the rising factorial $P_k(X) = \prod_{i\leq k}(X+i)$: a uniform division-free straight-line evaluation of P_k with exponent γ would factor a semiprime in $N^{\gamma/2+o(1)}$ ring operations by one evaluation at 0 and one

gcd, so that $\gamma = 1/3$ would give $N^{1/6+o(1)}$; the best known evaluator has exponent $1/2$ (Strassen; see [15] for the record), no evaluator of exponent $1/2 - \delta$ is known, and no unconditional lower bound above polylogarithmic size is known for the complexity of $k!$ or of P_k [52, 76, 12, 46]. The τ-conjecture of Shub and Smale, if it holds with exponent c, forces $\gamma \geq 1/c$, since P_k has k distinct integer roots; with $c = 2$ it would make Strassen's exponent optimal. We record the problem, not an approach to it.

7.2. **Piecewise evaluation.** One might hope to evaluate the tested product piece by piece along structured sub-families – chirp chains (u, vm^2), whose exponents are quadratic-plus-Beatty in m and admit sub-linear evaluators in the bit model, rays, or square arrays. Theorem 10.1 (Section 10) shows that a strategy which evaluates the pieces separately, at cost at least the square root of each piece's length, pays at least $(1 - C^{-1/2})N^{1/4}/(4\sqrt{H})$ where H is the largest harmonic mass $\sum_{\text{near}} 1/a_i$ of a piece; chirp chains, rays and square arrays all have $H \leq \sqrt{C}(1+\log N)+2$, so such decompositions cost $N^{1/4-o(1)}$, above Harvey's $N^{1/5}$. To come below $N^{1/5}$ some piece must have harmonic mass $\gg_C N^{1/10}$, hence contain $\gg_C N^{1/5}$ cells spanning $\gg_C N^{1/10}$ distinct values of a: a genuinely two-dimensional set whose exponents $\lfloor(\sqrt{aN} - \sqrt{b})^2\rfloor$ carry the square root in both parameters.

8. What an algorithm below $N^{1/5}$ would need

The picture inside the model is as follows. Lehman's $N^{1/3}$ candidates are forced by coverage (Corollary 3.3); Harvey's $N^{1/5}$ is optimal when giants are materialised (Theorem 3.4), also for unbalanced factors down to $p \approx N^{2/5}$ and for forms of every fixed degree, and the candidate floor holds against adaptive cell queries; an $N^{1/6}$ search follows from a difference cover of the candidate support, subject to the order and decoding conditions of Proposition 4.1, and in the explicit-difference representation model any such cover is bounded below in terms of the approximate-Sidon statistic of the window starts of a sub-family (Lemma 4.3, Theorem 4.4). The statistic has positive power growth in every sub-family for which it has been determined – $\sqrt{r}$ along rays and equal products, $rN^{-1/4}$ along half-offset lines after both are removed (for $C \geq 2$), and at least of order $r^{1/3}$ and $N^{1/12}$ in the $a = 1$ shell – so a bounded statistic is impossible for those sub-families, and the conditional floors for the unique-product family are $N^{5/28}$ under Conjecture 19.4 and $N^{2/11}$ under Conjecture 19.6. Under those conjectures and the family hypotheses of Proposition 19.5, the gain conceivable within the babystep–giantstep paradigm is therefore bounded: the materialised-giant cost is $N^{1/5}$, and the conditional explicit-cover floors are $N^{5/28}$ and $N^{2/11}$. A cover attaining the bound of Lemma 4.3 would need the start differences to be spread over $\asymp |B|$ values each of multiplicity $\asymp r^{1/3}$, a heuristic reading of the lemma's equality case, whereas in every computation of Section 6 the maximal cluster is carried by a single family and the direct search found no compression beyond that of a random set. The tested set is plausibly incompressible, in which case the true cover floor for the standard family is $N^{1/5}$; proving that would require an incompressibility theorem beyond Lemma 4.3, and disproving it a cover exploiting a structure the computations do not show.

Subject to the qualifications stated, each of the following would give a deterministic exponent below $1/5$ on every balanced semiprime, and none is excluded by the floors of this part, which price materialised giants (M1), the free-giant BSGS implementation (M2) and explicit difference representations (Theorem 4.4), and are not lower bounds for arbitrary exact-product identities or for factoring in general:

(i) a uniform division-free straight-line evaluation of the rising factorial P_k with exponent below $2/5$;

(ii) a uniform evaluation of every interval sub-product of the half-family Lehman product with exponent $\gamma < 3/5$, together with the width bound $Z_* = N^{o(1)}$ (Propositions 7.2 and 7.3): the bisection then costs $(1+Z_*)N^{\gamma/3+o(1)}$, the Fermat preprocessing $N^{1/6+o(1)}$, and the common-factor exception is handled below $N^{1/5}$ by Part 4;

(iii) an explicitly computable difference representation $\mathrm{supp}(Z) \subseteq B - G$ with $|B| + |G| = N^{1/5-\delta}$ satisfying the decoding condition of Proposition 4.1 and its order condition – which [33] supplies when B lies in an interval of length D at cost $D^{1/2+o(1)}$, and which Lemma 16.2 supplies for an explicitly given set B of $O(\log N)$-bit integers of spread $S \le N$ at cost $\tilde{O}(|B| + \sqrt{S})$, subordinate exactly when $\sqrt{S} \le N^{1/5-\delta}$; for the standard family Theorem 4.4 permits this only if the tolerance statistic of the squarefree $a = 1$ shell is large enough for the theorem's Sidon branch to fall below $N^{1/5-\delta}$, a quantitative condition on its growth, which Section 5 bounds from both sides;

(iv) a covering family of total tested multiplicity $N^{1/3+o(1)}$ whose tested product is a product of $N^{o(1)}$ pieces, each uniformly evaluable with exponent below $3/5$, together with a separation mechanism – a way to obtain a proper gcd from the pieces' values, which the bisection of Proposition 7.3 provides under its width hypothesis – for which no proposition is proved here: this is a direction, not a result; for any strategy with per-piece cost at least the square root of the piece's length, Theorem 10.1 forces some piece to have harmonic mass $\gg_C N^{1/10}$, so no decomposition entirely into chirp chains, rays or square arrays achieves it;

(v) a deterministic algorithm of exponent below $1/5$ outside the cell model – Coppersmith's method needs a quarter of the bits of one factor, thinning by the classes compatible with N modulo a squarefree modulus below half the windows changes only $N^{o(1)}$ factors (Lemma 2.3), and thinning by an enumerated residue hypothesis at most a power of $\log\log N$ (Theorem 2.8), so none of these yields such an algorithm on its own.

9. Appendix to Part 1: forms of bounded degree and adaptive queries

Definition 2.1 admits only the linear forms $aq+bp$. Any integer polynomial $u(p,q)$ is admissible for the Fermat identity $\alpha^{u(p,q)} \equiv \alpha^{u(1,N)} \pmod p$, and one may ask whether a covering family of non-linear forms escapes the floors. For nonconstant forms of every fixed total degree the answer is no in the materialised-giant model.

Lemma 9.1 (measure lemma)**.** *Let I be an interval, $k \ge 1$, $f \in C^k(I)$ with $|f^{(k)}| \ge \lambda > 0$ on I, and $S \subseteq \mathbb{R}$ measurable. Then $|\{t \in I : f(t) \in S\}| \le C_k\,(|S|/\lambda)^{1/k}$ with $C_1 = 1$ and $C_k = 2 + 2C_{k-1}$.*

Proof. For $k = 1$, f is strictly monotone and $|\{f \in S\}| = \int_{S\cap f(I)} |f'(f^{-1}(u))|^{-1}\,du \le |S|/\lambda$. For $k \ge 2$ and $\eta > 0$, $f^{(k-1)}$ is strictly monotone, so $E := \{t \in I : |f^{(k-1)}(t)| \le \eta\}$ is an interval of length at most $2\eta/\lambda$ by the case $k = 1$, and $I \setminus E$ consists of at most two intervals on which $|f^{(k-1)}| > \eta$; by induction each contributes at most $C_{k-1}(|S|/\eta)^{1/(k-1)}$. If $0 < |S| < \infty$, choosing $\eta = \lambda^{(k-1)/k}|S|^{1/k}$ makes both η/λ and $(|S|/\eta)^{1/(k-1)}$ equal to $(|S|/\lambda)^{1/k}$, giving $C_k = 2+2C_{k-1}$; if $|S| = 0$ let $\eta \downarrow 0$, and if $|S| = \infty$ there is nothing to prove. □

Every integer form $u(p,q)$ of total degree at most e restricts on $pq = N$ to $g(p) = \sum_{m=1}^{e}(\alpha_m p^m + \beta_m (N/p)^m) + \alpha_0$ with integer coefficients, since $p^i q^j = N^j p^{i-j}$ for $i > j$, $N^i q^{j-i}$ for $j > i$ and N^i for $i = j$; the restriction is nonconstant iff some α_m or β_m with $m \ge 1$ is nonzero, and $M := \max\{m \ge 1 : (\alpha_m, \beta_m) \ne (0,0)\}$ is its *effective degree*. Constant restrictions must be excluded: a single such form with one tested integer covers every real p, and its hit carries no information.

Lemma 9.2 (uniform derivative bound)**.** *Fix $C > 1$ and $M \geq 1$. There is $c_M = c_M(C) > 0$ such that for every $h(t) = \sum_{0<|m|\leq M} v_m t^m$ with real coefficients satisfying $\max(|v_M|, |v_{-M}|) \geq 1$, and every $t \in I_C = [C^{-1/2}, 1]$, $\max_{1\leq k\leq 2M} |h^{(k)}(t)| \geq c_M$.*

Proof. $(h'(t), \dots, h^{(2M)}(t))^{\mathsf{T}} = \mathsf{W}(t)\, v$ with $\mathsf{W}(t)_{k,m} = m(m-1)\cdots(m-k+1)\, t^{m-k}$. The rows are the derivatives of orders $0, \dots, 2M-1$ of the $2M$ functions $f_m(t) := mt^{m-1}$, nonzero multiples of powers with pairwise distinct exponents; the Wronskian of distinct powers $t^{e_1}, \dots, t^{e_n}$ equals $\prod_{i<j}(e_j - e_i)\, t^{\sum e_i - n(n-1)/2} \neq 0$ for $t > 0$, so $\det \mathsf{W}(t) \neq 0$ on I_C. The least singular value $\sigma(t)$ of $\mathsf{W}(t)$ is continuous and positive on the compact I_C; with $\sigma_0 := \min_{I_C} \sigma > 0$, $\max_k |h^{(k)}(t)| \geq \|\mathsf{W}(t)v\|_2/\sqrt{2M} \geq \sigma_0\|v\|_2/\sqrt{2M} \geq \sigma_0/\sqrt{2M} =: c_M$. □

Lemma 9.3 (coverage of a cell of effective degree M)**.** *Fix $C > 1$, $M \geq 1$. There is $K_M = K_M(C)$ such that for every nonconstant restriction g of effective degree M and every measurable S of measure $w \leq (c_M/2)N^{M/2}$, $|\{p \in J_C : g(p) \in S\}| \leq K_M\, w^{1/(2M)} N^{1/4}$.*

Proof. With $p = \sqrt{N}\, t$, $g(p) = N^{M/2}h(t) + \alpha_0$ where $h(t) = \sum_{0<|m|\leq M} v_m t^m$, $v_m = \alpha_m N^{(m-M)/2}$, $v_{-m} = \beta_m N^{(m-M)/2}$; thus $(v_M, v_{-M}) = (\alpha_M, \beta_M) \in \mathbb{Z}^2 \setminus \{0\}$ and Lemma 9.2 applies. Put $S' := (S - \alpha_0)/N^{M/2}$, $|S'| = wN^{-M/2} \leq c_M/2$, and $U_k := \{t \in I_C : |h^{(k)}(t)| > c_M/2\}$ for $1 \leq k \leq 2M$; by Lemma 9.2, $I_C = \bigcup_k U_k$. For either sign, $P_{k,\pm}(t) := t^{M+k}(h^{(k)}(t) \pm c_M/2)$ is a real polynomial with at most $2M+1$ nonzero monomials; if it is not identically zero it has at most $2M$ positive roots (Descartes), and if one of $P_{k,\pm}$ vanishes identically then $U_k = \emptyset$. So U_k has at most $4M+2$ components, on each of which $|h^{(k)}| \geq c_M/2$. By Lemma 9.1, since $2|S'|/c_M \leq 1$ and $C_k \leq C_{2M}$, each component contributes at most $C_{2M}(2|S'|/c_M)^{1/(2M)}$, so $|\{t \in I_C : h(t) \in S'\}| \leq 2M(4M+2)\, C_{2M}(2w/(c_M N^{M/2}))^{1/(2M)}$, and multiplying by $\sqrt{N}$ gives the claim with $K_M := 2M(4M+2)C_{2M}(2/c_M)^{1/(2M)}$. □

Theorem 9.4 (forms of bounded total degree)**.** *Fix $e \geq 1$ and $C > 1$. There are $c, c' > 0$ depending only on e and C such that every oblivious covering family of pairwise distinct integer forms of total degree at most e with nonconstant restrictions, covering J_C in the sense of Definition 2.1(3), satisfies $\Sigma_W \geq c\, N^{1/4}$ and $T \geq c'\, N^{1/5}$ in the model M1.*

Proof. Group the forms by effective degree $M \in \{1, \dots, e\}$. The set covered by each class is measurable (a finite union of preimages of measurable sets under continuous maps), and the e classes together cover J_C, so some class covers a measurable $E_M \subseteq J_C$ with $|E_M| \geq L/e$, $L = (1 - C^{-1/2})\sqrt{N}$. *Case $M = 1$.* Each restriction is $g_i(p) = \delta_i N/p + \varepsilon_i p + \kappa_i$ with integers $(\delta_i, \varepsilon_i) \neq (0,0)$. If $\delta_i\varepsilon_i \leq 0$ then $|g_i'| = |\varepsilon_i - \delta_i N/p^2| \geq 1$ on J_C, the two terms having the same sign or one of them vanishing, so the monotone bound gives coverage at most w_i. Otherwise write $(\delta_i, \varepsilon_i) = \sigma_i d_i (a_i, b_i)$ with $\sigma_i = \pm 1$, $d_i \geq 1$ and coprime $a_i, b_i \geq 1$; then $\{p : g_i(p) \in S_i\} = \{p : a_i N/p + b_i p \in S_i'\}$ with $S_i' := \sigma_i(S_i - \kappa_i)/d_i$, a measurable set of measure $w_i/d_i \leq w_i$. Group these forms by their normalised pair (a, b), and give the pair the measurable tested set $S'_{(a,b)} := \bigcup S_i'$ over the forms of the group, of measure at most the sum of their w_i. This is a family of $R \leq P_1$ distinct pairs with measurable tested sets which, together with the monotone forms, covers E_1; Lemma 3.1 and the monotone bound hold for measurable sets, so the proof of Lemma 3.2 applies verbatim and gives $\Sigma_w \geq e^{-2} c_C \sqrt{N}/\sqrt{R}$, the total measure of the transformed sets being at most Σ_w. Hence $\Sigma_W \geq \Sigma_w/2 \gg_{C,e} \sqrt{N}/\sqrt{R}$, while $\Sigma_W \geq P_1 \geq R$ and, in M1, $T \geq P_1 \geq R$ and $T \geq 2\sqrt{\Sigma_W}$; optimising in R as in Corollary 3.3 and Theorem 3.4 gives $\Sigma_W \gg N^{1/3}$ and $T \gg N^{1/5}$. *Case $M \geq 2$.* Let the class have P_M forms. If some cell has $w_i > (c_M/2)N^{M/2}$ then $\Sigma_w > (c_M/2)N^{M/2}$; otherwise Lemma 9.3 and Hölder's inequality $\sum w_i^{1/(2M)} \leq P_M^{1-1/(2M)} (\sum w_i)^{1/(2M)}$ give $L/e \leq K_M N^{1/4} P_M^{1-1/(2M)} \Sigma_w^{1/(2M)}$, i.e. $\Sigma_w \geq \kappa_M N^{M/2} P_M^{-(2M-1)}$ with $\kappa_M := ((1 - C^{-1/2})/(eK_M))^{2M} \leq c_M/2$, so the same bound

holds in both cases. Hence $\Sigma_W \geq \max\{P_M, \kappa_M N^{M/2}/(2P_M^{2M-1})\} \geq (\kappa_M/2)^{1/(2M)} N^{1/4}$, and in M1, $T \geq \max\{P_M, 2\sqrt{\Sigma_W}\} \geq \max\{P_M, \sqrt{2\kappa_M}\, N^{M/4} P_M^{-(2M-1)/2}\} \geq (2\kappa_M)^{1/(2M+1)} N^{M/(4M+2)}$, and $M/(4M+2) \geq 1/5$ for $M \geq 2$. □

For $M = 1$ the crude argument would give only $N^{1/6}$ in M1, which is why the linear class is routed through Lemma 3.2. The candidate floor $N^{1/4}$ is weaker than the $N^{1/3}$ of Corollary 3.3 because no bound on the number of forms with degenerate critical points in J_C is used; the materialised-giant floor does not move. This is a model-relative bound: it says nothing about algorithms that share work between algebraically related forms outside M1.

Proposition 9.5 (adaptive cell queries)**.** *Fix $C > 1$. Let a deterministic algorithm, knowing N, learn about a real parameter $p \in J_C$ only through an oracle which, on a query $(a, b, n) \in \mathbb{Z}^2_{\geq 1} \times \mathbb{Z}$, answers whether $|aN/p + bp - n| < 1$. Suppose that for every real $p \in J_C$ the execution against the oracle determined by $(p, N/p)$ eventually receives an affirmative answer. Then for every N there is a real $p \in J_C$ against which the algorithm makes at least $(c_C/2)^{2/3} N^{1/3}$ queries.*

Proof. Follow the all-negative transcript for as long as the algorithm issues queries, and denote its successive queries by $Q_j = (a_j, b_j, n_j)$. Let $H_j := \{p \in J_C : |a_j N/p + b_j p - n_j| < 1\}$ and $U_M := \bigcup_{j \leq M} H_j$; each H_j is relatively open in the compact interval J_C. Some finite U_M contains J_C: if the all-negative execution issues only finitely many queries, say L, and $U_L \neq J_C$, then any $p \in J_C \setminus U_L$ produces exactly that execution and never receives an affirmative answer, contrary to the hypothesis; if it issues infinitely many and no finite U_M contains J_C, then either $\bigcup_M U_M \neq J_C$ and a point outside the union is answered negatively forever, or $\bigcup_M U_M = J_C$ and compactness gives a finite subcover – a contradiction either way. Let M be least with $U_M \supseteq J_C$. Group $Q_1, \ldots, Q_M$ by the form (a, b) and put $T_{a,b} := \{n_j : j \leq M,\ (a_j, b_j) = (a, b)\}$. These data satisfy conditions (1)–(3) of Definition 2.1, which are all that Corollary 3.3 uses, with $\Sigma_W \leq M$ (repeated queries only reduce the count); hence $M \geq (c_C/2)^{2/3} N^{1/3}$. Since M is least, $J_C \setminus U_{M-1} \neq \emptyset$; against the oracle of any p in that set the first $M-1$ answers are negative, so determinism forces the algorithm to issue its M-th query. □

The proposition bounds the number of singleton cell-membership queries, i.e. the candidate count; it is not a bound in M1 or M2 for adaptive batched searches. If a query may name a finite block of integers and is answered affirmatively when any member hits, the number of oracle queries can be smaller; expanding the blocks into their tested integers still gives the candidate-count floor, and the cost is governed by M1 or M2 only when the blocks are implemented by the batching mechanisms those models specify. A test of the form "does p divide $\alpha^m - 1$" is answered by a single gcd and eliminates every p whose order divides m; that is the smooth-order paradigm of Pollard's $p - 1$ method [68], not a cell query, and we have no lower bound for it.

10. Appendix to Part 1: piecewise evaluation

Theorem 10.1 (piecewise evaluation)**.** *Let a family of cells (a_i, b_i), with tested sets T_i and candidate sets S_i as in Definition 2.1, cover J_C, and write $W_i := |T_i| \geq 1$ and $w_i := |S_i| \in [W_i, 2W_i]$. Partition the cells, with their tested integers, into nonempty pieces $I_1, \ldots, I_F$, and put $W_j := \sum_{i \in I_j} W_i$. Call a cell* near *if $a_i \leq b_i \leq \lceil C a_i \rceil$, and define the harmonic mass of a piece and its maximum by*

$$H_j := \sum_{i \in I_j \text{ near}} \frac{1}{a_i}, \qquad H := \max\{1, \max_j H_j\}.$$

Consider a strategy that evaluates the sub-products of the pieces separately, with additive charged operation counts, the sub-product of piece j costing at least W_j^γ ring operations for a fixed $\gamma \geq$

1/2. *Then its total charged cost is at least* $(1-C^{-1/2})N^{1/4}/(4\sqrt{H})$. *Moreover* $H_j \le (C+1)A_j$ *and* $H_j \le 2\sqrt{(C+1)P_j}$, *where* A_j *is the number of distinct values of* a_i *in piece* j *and* $P_j := |I_j|$. *Chirp chains* (u, vm^2) *and* (um^2, v), *rays* $m(a_0, b_0)$ *and square arrays* (s^2u, t^2v) *with* $a_i \le N$ *all satisfy* $H_j \le \sqrt{C}(1+\log N)+2$.

Proof. Put $x := 1-C^{-1/2}$, so $|J_C| = x\sqrt{N}$, and $\Sigma := \sum_j \sqrt{W_j}$. A cell that is not near has $b_i \le a_i - 1$ or $b_i \ge \lceil Ca_i \rceil + 1$, and the monotone bound gives coverage at most $w_i \le 2W_i$; a near cell covers, by Lemma 3.1, at most $2\sqrt{w_i\sqrt{N}/a_i} \le 2\sqrt{2}\,N^{1/4}\sqrt{W_i/a_i}$. As the covered sets exhaust J_C,

$$x\sqrt{N} \le 2\sum_j W_j + 2\sqrt{2}\,N^{1/4}\sum_j \sum_{i\in I_j \text{ near}} \sqrt{\frac{W_i}{a_i}} \le 2\Sigma^2 + 2\sqrt{2}\,N^{1/4}\sqrt{H}\,\Sigma,$$

by Cauchy–Schwarz within each piece and $\sum_j W_j \le \Sigma^2$. If $\Sigma < xN^{1/4}/(4\sqrt{H})$, the right side is less than $x\sqrt{N}(x/(8H) + 1/\sqrt{2}) < x\sqrt{N}$, a contradiction; so $\Sigma \ge xN^{1/4}/(4\sqrt{H})$. Every piece is nonempty, so $W_j^{\gamma} \ge W_j^{1/2}$, and the total charged cost is at least Σ. For fixed a the near cells have $b \in [a, \lceil Ca \rceil]$, at most $Ca+1$ values, contributing at most $C+1$ to H_j; hence $H_j \le (C+1)A_j$. For real $A_0 \ge 1$ the near cells with $a_i < A_0$ contribute at most $(C+1)A_0$ and those with $a_i \ge A_0$ at most P_j/A_0; if $P_j \ge C+1$, take $A_0 = \sqrt{P_j/(C+1)}$, and otherwise $H_j \le P_j \le 2\sqrt{(C+1)P_j}$. For a chain (u, vm^2) a near cell has $vm^2 \le Cu+1$, so $H_j \le \sqrt{C+1}$; for (um^2, v), $H_j \le \pi^2/6$; for a ray, $H_j \le 1+\log N$; for an array (s^2u, t^2v) a near cell has $t^2v \le Cs^2u+1$, so $H_j \le \sqrt{C}(1+\log S)+\pi^2/6$ with $S \le \sqrt{N}$. □

11. Appendix to Part 1: hits are convergents or intermediate fractions

The bisection of Proposition 7.3 needs a bound on the denominator of a covering cell. Every Lehman hit is a convergent or an intermediate fraction of p/q; the natural route – Fatou's theorem, which classifies a/b with $|\xi - a/b| < 1/b^2$ – falls short by a factor $1+O(1/r)$ at the top of the shell (hits with $b^2|\xi - a/b| \in (1, 1.0002)$ exist), and the right threshold is the Stern–Brocot one.

Lemma 11.1 (Stern–Brocot parent criterion)**.** *Let* $\xi > 0$ *and let* a/b *be in lowest terms with* $b \ge 2$. *If* $|\xi - a/b| < 1/(b(b-1))$ *then* a/b *is a convergent or an intermediate fraction of* ξ.

Proof. In the Stern–Brocot tree [28, §4.5] every reduced a/b with $b \ge 2$ is the mediant of its two parents $a_1/b_1 < a/b < a_2/b_2$, with $a = a_1 + a_2$, $b = b_1 + b_2$ and $a_2b_1 - a_1b_2 = 1$; hence $a/b - a_1/b_1 = 1/(bb_1)$ and $a_2/b_2 - a/b = 1/(bb_2)$, both at least $1/(b(b-1))$. So the hypothesis places ξ strictly between the parents of a/b, i.e. a/b lies on the Stern–Brocot path to ξ, and the fractions on that path are exactly the convergents and intermediate fractions [28, §4.5], [44, §§6, 15]. □

Proposition 11.2 (hits are convergents or intermediate fractions)**.** *Let* $N = pq$ *with primes* $p < q \le Cp$, $C > 1$, *let* $r \ge 16C$, *and let* (a, b) *with* $a, b \ge 1$, $ab \le r$ *be a hit:* $0 \le aq + bp - 2\sqrt{abN} < W := \sqrt{N}/(4r\sqrt{ab})$. *Put* $m := |aq - bp|$ *and* $\xi := p/q$. *Then*

$$\textit{(i)}\ m^2 < \frac{N}{r} + W^2, \qquad \textit{(ii)}\ b < \sqrt{rq/p} + \frac{C^{1/4}}{2r^{1/4}} + \frac{1}{8r}, \qquad \textit{(iii)}\ m(b-1) < q,$$

so $|\xi - a/b| < 1/(b(b-1))$ *for* $b \ge 2$, *and* a/b *is a convergent or an intermediate fraction of* p/q.

Proof. Write $x = aq$, $y = bp$, so the hit is $(\sqrt{x} - \sqrt{y})^2 < W$ and $m = |\sqrt{x} - \sqrt{y}|(\sqrt{x} + \sqrt{y})$. (i) $(\sqrt{x} + \sqrt{y})^2 = (\sqrt{x} - \sqrt{y})^2 + 4\sqrt{xy} < W + 4\sqrt{abN}$, and $W \cdot 4\sqrt{abN} = N/r$, so $m^2 < W^2 + N/r$.

(ii) From $|\sqrt{x}-\sqrt{y}|<\sqrt{W}$ and $bp\ge p\ge\sqrt{N/C}>W$ we get $\sqrt{x}>\sqrt{y}-\sqrt{W}\ge 0$, hence $aq>bp-2\sqrt{bpW}$, and with $ab\le r$, $b^2p-2b\sqrt{bpW}<rq$, i.e. $(b\sqrt{p}-\sqrt{bW})^2<rq+bW$, so $b\sqrt{p}<\sqrt{bW}+\sqrt{rq}+bW/(2\sqrt{rq})$. Since $a\ge 1$ and $b\le r$, $bW\le\sqrt{N}/(4\sqrt{r})$; dividing by $\sqrt{p}\ge(N/C)^{1/4}$ gives (ii). (iii) By (i), $m<\sqrt{N/r}\,(1+1/(32r))$; by (ii), $b-1<\sqrt{rq/p}\,(1-(1-\varepsilon)\sqrt{p/(rq)})$ with $\varepsilon:=C^{1/4}/(2r^{1/4})+1/(8r)<1/2$ for $r\ge 16C$; and $\sqrt{N/r}\sqrt{rq/p}=q$. Hence $m(b-1)<q\,(1+1/(32r))(1-\delta)$ with $\delta>1/(2\sqrt{Cr})$, and $(1+1/(32r))(1-\delta)<1$. For $b\ge 2$, $|\xi-a/b|=m/(bq)<1/(b(b-1))$; with $g=\gcd(a,b)$, $a'=a/g$, $b'=b/g$, if $b'\ge 2$ then $|\xi-a'/b'|<1/(b'(b'-1))$ and Lemma 11.1 applies, and if $b'=1$ the estimate $<1/2$ with $0<\xi<1$ forces $a'=1$ and $\xi>1/2$, so $1/1$ is the convergent h_1/k_1. For $b=1$: by (i), $m/q<\sqrt{p/(rq)}\,(1+1/(32r))<1/2$, so $a=1$ and again $1/1$ is the convergent h_1/k_1. □

Part 2. Partial information: sharp floors and matching algorithms

Summary. Harvey's deterministic factoring algorithm runs in $N^{1/5+o(1)}$ operations. We ask what changes when partial information about the smaller factor p of a balanced semiprime $N=pq$ is available – an interval of length $L=N^\lambda$ containing p, the residue $p \bmod M$ with $M=N^\mu$, or both – within the covering model of Part 1, which contains Harvey's algorithm. Given the interval, a localised version of Harvey's search factors N deterministically in $N^{o(1)}L^{3/5}N^{-1/10}$ operations for $L\ge N^{3/8}$, matching the model's floor there and lying below $N^{1/5}$ for every $L<N^{1/2}$; below $N^{3/8}$ the floor's exponent $\lambda-1/4$ is that of Coppersmith subdivision, which lies outside the model, while inside it, for $N^{1/6}\ll L\le N^{3/8}$, the cost of an interval is classified up to constants by one Diophantine quantity, its depth $\mathfrak{D}(I)$ (the effective denominator of the best single form): $T(I)\asymp\min\{\sqrt{\mathfrak{D}(I)}\,LN^{-1/4},L^{3/5}N^{-1/10}\}$, the second scale – attained for every interval by a family built from Farey fractions – being the cost of all but a vanishing fraction of the intervals. With the residue known the same localisation gives $N^{1/5+o(1)}M^{-2/5}$ for $M\le N^{1/7}$, and with both kinds of information the cells beat Coppersmith subdivision exactly on the region $8(1/2-\lambda)+12\mu<1$. The element of large order that the search needs is supplied for every bound by a recent theorem of Harvey and Hittmeir. The algorithms are deterministic and unconditional; the floors are statements about the covering model and are labelled as such.

12. Introduction

Let $N=pq$ with primes $p<q$ and, for a fixed $C>1$, $p\in J_C:=[\sqrt{N/C},\sqrt{N}]$. The best deterministic factoring algorithm known runs in $N^{1/5+o(1)}$ bit operations [31, 32]; it is a babystep–giantstep search over Lehman's cells $aq+bp$ [49], and Part 1 proves that within a model containing Harvey's algorithm and, up to factors $N^{o(1)}$, Lehman's (and Hart's [30] where its family covers the range), its exponent cannot be improved when giant elements are materialised. Here we ask how the picture changes when partial information about p is available: an interval $I\ni p$ of length $L=N^\lambda$ (the leading bits of p), or the residue $p \bmod M$ with $M=N^\mu$ (the trailing bits), or both. Coppersmith's method [17, 42] factors N in polynomial time once $L/M\le N^{1/4}$, i.e. once about half of the bits of p are known; between that threshold and no information, we have found in the literature only Coppersmith subdivision, of cost $N^{\lambda-\mu-1/4}$, and Hittmeir's Strassen-type search on a residue class [39], of cost $N^{1/4}/\sqrt{M}$.

Our results are as follows. The covering model and its floors are those of Part 1 (Definition 2.1, Lemmas 3.1 and 3.2, Corollary 3.3, Theorems 3.4 and 3.6); the covering-floor arguments below rest on the two lemmas. Section 13 proves the floors under localisation (Proposition 13.1 and the sharp Theorem 13.2), constructs the localised search (Lemmas 13.3, 13.4, 13.5 and Theorem 13.6), recalls the theorem of Harvey and Hittmeir that supplies the element of large order on which every Harvey-type search rests, for every bound D (Section 13.2; Section 16 gives a balanced-semiprime proof with a different recovery step, and its variant for explicit baby sets),

and assembles the unconditional landscape (Corollaries 13.8 and 13.9): $\min\{\lambda - 1/4, 3\lambda/5 - 1/10\}$, at the exponent level a floor in the model at every λ (Theorem 13.2 carries a logarithmic loss in its first branch, which Theorem 13.15 removes for $c_6 N^{1/6} \le L \le N^{3/8}$), attained inside the model for $\lambda \ge 3/8$ and, for $\lambda < 3/8$, by Coppersmith subdivision outside it. Section 13.4 then shows that below $N^{3/8}$ the in-model cost depends on the position of the interval and not only on its length (Theorem 13.15): it is determined up to constants by a single Diophantine quantity, the depth of the interval; the floor $LN^{-1/4}$ holds without logarithmic loss and is the cost of the intervals of bounded depth, and a family built from Farey fractions covers every interval at cost $O(L^{3/5}N^{-1/10})$, which is also the floor for all but a vanishing fraction of the centres. Section 14 treats arithmetic thinning of the tested set with genuine side information: starting from the thinned families of Part 1 (Definition 2.2, Lemmas 2.3 and 2.4), it gives the floors for thinned families (Proposition 14.1), constructs the residue-class variant of Harvey's search, of cost $N^{1/5+o(1)}M^{-2/5}$ for $M \le N^{1/7}$ (Proposition 14.2), and the two-axis landscape, on which the cells beat Coppersmith subdivision exactly when $8(1/2 - \lambda) + 12\mu < 1$ (Corollary 14.3), and records that two per-cell lattice methods do not improve the Coppersmith branch. Section 15 places the results against the literature and says what in them is more than running Harvey's search on a shorter range.

Every algorithm and theorem stated here is deterministic and unconditional, and every model-relative statement says so. The measurements reported are finite checks of implementations, not evidence for the asymptotics. The notation and the model are those of Part 1; throughout, $\omega(n)$ is the number of distinct prime factors of n, $\pi(x)$ the number of primes up to x, φ Euler's function, $d(n)$ the number of divisors, log the natural logarithm and $\tilde{O}$ hides factors $(\log N)^{O(1)}$. Harvey's algorithm is cited by its numbered steps [31, Algorithms 4.1–4.3], Hittmeir's by [39, Section 6]. Moduli whose factors satisfy a structural rather than a partial-information hypothesis – a large common factor of $p-1$ and $q-1$ – are treated in Part 4.

13. Floors under partial information

Nothing in the proof of Lemma 3.2 uses that the set to be covered is the whole of J_C: Lemma 3.1 bounds the coverage of any measurable set by that of the interval containing it, and the two cases of Lemma 3.2 only use the total measure L to be covered. This bounds from below, within the model, what a non-oblivious algorithm – one that uses polylogarithmic information about N to exclude parts of J_C – would gain.

In this section the localising set or interval is regarded as auxiliary input: the covering family may depend on N and on that localisation data, but not on the unknown point p within it, so "as in Definition 2.1" refers to the form, tested-set and cost conventions of that definition, with obliviousness relative to the remaining uncertainty in p.

Proposition 13.1 (floors under localisation)**.** *Let $S \subseteq J_C$ be measurable of measure L, and let a family as in Definition 2.1 satisfy the coverage condition (3) for every real $p \in S$ (rather than for every $p \in J_C$). Then*

$$\Sigma_w \;\ge\; \frac{L^2}{48\sqrt{C+1}\,\sqrt{NP}} \quad (P \ge C+1), \qquad \Sigma_w \;\ge\; \frac{L^2}{32(C+1)\sqrt{N}} \quad (P < C+1),$$

and consequently, for $N^{1/4} \le L \le |J_C|$ and with constants depending only on C,

$$\Sigma_W \;\ge\; c_1\Big(\frac{L}{N^{1/4}}\Big)^{4/3}, \qquad T \;\ge\; c_2\Big(\frac{L}{N^{1/4}}\Big)^{4/5} \text{ in M1}, \qquad T \;\ge\; c_3\Big(\frac{L}{N^{1/4}}\Big)^{2/3} \text{ in M2}.$$

At $L = |J_C| \asymp \sqrt{N}$ these are the floors $N^{1/3}$, $N^{1/5}$, $N^{1/6}$ of Corollary 3.3 and Theorems 3.4–3.6; at $L = N^{1/4}$ the three displayed lower-bound scales are constants depending on C (the proposition gives no upper bound there).

Proof. Run the proof of Lemma 3.2 with S in place of J_C. For each form, $|\{p \in S : g_i(p) \in S_i\}| \le |\{p \in J_C : g_i(p) \in S_i\}|$, which is $\le w_i$ for a monotone form ($|g_i'| \ge 1$ on $J_C \supseteq S$) and $\le 2\sqrt{w_i\sqrt{N}/a_i}$ for a wide one by Lemma 3.1; coverage of S gives $L \le \sum_{\text{mon}} w_i + \sum_{\text{wide}} 2\sqrt{w_i\sqrt{N}/a_i}$. If the first sum is $\ge L/2$ then $\Sigma_w \ge L/2$; otherwise the Cauchy–Schwarz step of Lemma 3.2 gives $\Sigma_w \ge L^2N^{-1/2}/(16\sum_{\text{wide}} 1/a_i)$, and the bound on $\sum_{\text{wide}} 1/a_i$ – which counts available wide forms and does not depend on the set covered – is unchanged, giving $\Sigma_w \ge L^2/(48\sqrt{C+1}\sqrt{NP})$ for $P \ge C+1$ and $\Sigma_w \ge L^2/(32(C+1)\sqrt{N})$ for $P < C+1$. Since $L \le |J_C| < \sqrt{N}$, each of these is smaller than $L/2$ (their ratios to $L/2$ are $L/(24\sqrt{C+1}\sqrt{NP}) < 1$ and $L/(16(C+1)\sqrt{N}) < 1$), so the monotone case only strengthens the bound and the two displayed inequalities hold in all cases.

For the consequences put $X := L/N^{1/4} \in [1, N^{1/4}]$, $K_C := 1/(96\sqrt{C+1})$ and $D_C := 1/(64(C+1))$. Using $\Sigma_W \ge \Sigma_w/2$ and $\Sigma_W \ge P$: for $P \ge C+1$, $\Sigma_W \ge \max\{P, K_C X^2 P^{-1/2}\} \ge K_C^{2/3}X^{4/3}$ (balance at $P^{3/2} = K_C X^2$); for $P < C+1$, $\Sigma_W \ge D_C X^2 \ge D_C X^{4/3}$ since $X \ge 1$. In M1, $T \ge m + \sum_i \lceil W_i/m \rceil \ge P$ (each ceiling is at least 1) and $T \ge m + \Sigma_W/m \ge 2\sqrt{\Sigma_W}$; hence for $P \ge C+1$, $T \ge \max\{P, 2K_C^{1/2}XP^{-1/4}\} \ge 2^{4/5}K_C^{2/5}X^{4/5}$ (balance at $P^{5/4} = 2K_C^{1/2}X$), and for $P < C+1$, $T \ge 2D_C^{1/2}X \ge 2D_C^{1/2}X^{4/5}$. In M2, $T \ge 2\sqrt{\Sigma_W} \ge 2c_1^{1/2}X^{2/3}$. One may take $c_1 = \min\{K_C^{2/3}, D_C\}$, $c_2 = \min\{2^{4/5}K_C^{2/5}, 2D_C^{1/2}\}$ and $c_3 = \min\{2K_C^{1/3}, 2D_C^{1/2}\}$. □

The proposition charges every form the stationary coverage of Lemma 3.1, but a form is stationary inside I only if its ratio a/b lies close to the ratio interval $\{p^2/N : p \in I\}$; elsewhere the slope $|g'|$ is large and the coverage small. Accounting for this sharpens all three floors to the exponents actually achieved.

Theorem 13.2 (sharp floors under localisation)**.** *Let $I = [p_0, p_1] \subseteq J_C$ be an interval of length L with $N^{1/4} \le L \le \sqrt{N}/(4C)$, let $N \ge 4$, and let a family as in Definition 2.1 (depending on I, not on $p \in I$) satisfy the coverage condition (3) for every real $p \in I$. Then*

$$T \;\ge\; c^{(1)}_{C,\mathrm{loc}} \min\Big\{L^{3/5}N^{-1/10},\; \frac{LN^{-1/4}}{\sqrt{\log_2 N}}\Big\} \text{ in M1},$$

$$T \;\ge\; c^{(2)}_{C,\mathrm{loc}} \min\Big\{L^{1/2}N^{-1/12},\; \frac{LN^{-1/4}}{\sqrt{\log_2 N}}\Big\} \text{ in M2},$$

$$\Sigma_W \;\ge\; (256C)^{-1} \min\Big\{LN^{-1/6},\; \frac{L^2N^{-1/2}}{\log_2 N}\Big\},$$

where $c^{(1)}_{C,\mathrm{loc}} := \min\{(453C)^{-2/5}, (64C)^{-1/2}\}$ and $c^{(2)}_{C,\mathrm{loc}} := \min\{(250C)^{-1/3}, (64C)^{-1/2}\}$. For $L = N^{\lambda+o(1)}$, ignoring logarithmic factors, the M1 floor has exponent $\lambda - 1/4$ on $[1/4, 3/8]$ and $3\lambda/5 - 1/10$ on $[3/8, 1/2]$, and the M2 floor $\lambda - 1/4$ on $[1/4, 1/3]$ and $\lambda/2 - 1/12$ on $[1/3, 1/2]$.

Proof. Write $g_i(p) = a_iN/p + b_ip$, $\rho_i := a_i/b_i$, $X := [p_0^2/N, p_1^2/N] \subset [1/C, 1]$, $\ell := |X| = L(p_0+p_1)/N$, so that $2L/\sqrt{CN} \le \ell \le 2L/\sqrt{N} \le 1/(2C)$, and $\delta_i := \operatorname{dist}(\rho_i, X)$. For $p \in I$ put $x = p^2/N \in X$; then $g_i'(p) = b_i - a_i/x = (b_i/x)(x - \rho_i)$, so $|g_i'| \ge b_i|x - \rho_i| \ge b_i\delta_i$ on I. Hence, besides Lemma 3.1, every form with $\delta_i > 0$ is strictly monotone on I, its inverse is $1/(b_i\delta_i)$-Lipschitz, and it covers a set of measure at most $w_i/(b_i\delta_i)$ in I for any measurable tested set of measure w_i. Call a form *near* if $\delta_i \le \ell$ and *far* otherwise. A near form has

$\rho_i \ge 1/C - \ell \ge 1/(2C)$, i.e. $a_i \ge b_i/(2C)$, so Lemma 3.1 gives coverage $\le 2\sqrt{2Cw_i\sqrt{N}/b_i}$; and for each b the near forms with $b_i = b$ have a_i in an interval of length $3\ell b$, so at most $3\ell b + 1$ of them, whence at most $6\ell 4^k + 2^k$ near forms have $b_i \in [2^k, 2^{k+1})$. Coverage of I gives

$$L \;\le\; \frac{1}{\ell}\sum_{\text{far}} w_i \;+\; 2\sqrt{2C}\,N^{1/4}\sum_{\text{near}}\sqrt{w_i/b_i} \;\le\; \frac{\Sigma_w}{\ell} \;+\; 2\sqrt{2C}\,N^{1/4}\sqrt{\Sigma_w}\,H^{1/2}, \qquad H := \sum_{\text{near}}\frac{1}{b_i}.$$

Throughout, $\Sigma_w \le 2\Sigma_W$ (Definition 2.1) and, in both models, $T \ge m + \Sigma_W/m \ge 2\sqrt{\Sigma_W}$, so $\Sigma_w \le T^2/2$.

Far case ($\Sigma_w/\ell \ge L/2$): $\Sigma_w \ge L\ell/2 \ge L^2/\sqrt{CN}$, so $T \ge \sqrt{2}\,C^{-1/4}LN^{-1/4}$ and $\Sigma_W \ge L^2/(2\sqrt{CN})$; both exceed the corresponding logarithmic terms below.

Near case ($\Sigma_w H \ge L^2/(32C\sqrt{N})$). Let P_k be the number of near forms with $2^k \le b_i < 2^{k+1}$ and put $z := 6\ell$. Since $P_k \le z4^k + 2^k$, write $P_k = Q_k + R_k$ with $0 \le Q_k \le z4^k$ and $0 \le R_k \le 2^k$; then $\sum_k Q_k, \sum_k R_k \le \sum_k P_k \le P$. Splitting at $2^k = \sqrt{P/z}$, $\sum_k 2^{-k}Q_k \le \sum_k 2^{-k}\min\{P, z4^k\} \le 4\sqrt{zP}$, the two geometric sums being each at most $2\sqrt{zP}$; and if $n = \lfloor \log_2 P \rfloor$ and $f = \log_2 P - n$, then $\sum_k 2^{-k}R_k \le n + 2^{-n}\sum_{k\ge n} R_k \le n + P2^{-n} = n + 2^f \le \log_2 P + 1$. Consequently $H \le 4\sqrt{6\ell P} + \log_2 P + 1 \le 10\sqrt{\ell P} + \log_2 P + 1$.

In M1, if $10\sqrt{\ell P} \ge \log_2 P + 1$ then $20\Sigma_w\sqrt{\ell P} \ge L^2/(32C\sqrt{N})$, and $\ell \le 2L/\sqrt{N}$ gives $\Sigma_w\sqrt{P} \ge L^{3/2}N^{-1/4}/(906C)$; with $\Sigma_w \le T^2/2$ and $P \le T$ this is $T^{5/2} \ge L^{3/2}N^{-1/4}/(453C)$. Otherwise $2\Sigma_w(\log_2 P + 1) \ge L^2/(32C\sqrt{N})$; we may assume $T \le N$ (else the asserted bound is trivial), so $P \le N$, $\log_2 P + 1 \le 2\log_2 N$, and $T^2 \ge L^2N^{-1/2}/(64C\log_2 N)$. Together with the far case this is the M1 bound.

For the candidate bound argue separately. If $\Sigma_W > N$ the asserted bound is immediate, so assume $\Sigma_W \le N$; then $P \le \Sigma_W \le N$ and $\log_2 P + 1 \le 2\log_2 N$. In the first near sub-case $2\Sigma_W^{3/2} \ge L^{3/2}N^{-1/4}/(906C)$, whence $\Sigma_W \ge (1812C)^{-2/3}LN^{-1/6} \ge (256C)^{-1}LN^{-1/6}$; in the logarithmic sub-case $L^2/(32C\sqrt{N}) \le 2\Sigma_w(\log_2 P + 1) \le 8\Sigma_W\log_2 N$, so $\Sigma_W \ge L^2N^{-1/2}/(256C\log_2 N)$; the far case gives the stronger $\Sigma_W \ge L^2/(2\sqrt{CN})$.

In M2, if $T > N^{1/6}$ the asserted bound is immediate, so assume $T \le N^{1/6}$. Since every form has $W_i \ge 1$, $P \le \Sigma_W \le T^2/4 \le N^{1/3}/4$. Put $K := \max\{0, \lceil \log_2 \sqrt{P/(6\ell)} \rceil\}$. For every distribution of the P_k,

$$H \le \sum_{k\le K}(6\ell 2^k + 1) + 2^{-K-1}P \le \tfrac{9}{2}\sqrt{6\ell P} + K + 1,$$

and $K + 1 \le \frac{1}{2}\log_2 P + \frac{1}{2}\log_2(1/(6\ell)) + 2 \le \log_2 N + 2$, since $1/(6\ell) \le \sqrt{CN}/(12L) \le \sqrt{N}$ (the hypotheses on L give $4CL \le \sqrt{N} \le L^2$). With $\sqrt{P} \le T/2$ this is $H \le A + B$, $A := (9\sqrt{6}/4)\sqrt{\ell}\,T$, $B := \log_2 N + 2$. If $A \ge B$, then $L^2/(32C\sqrt{N}) \le 2A\Sigma_w \le (9\sqrt{6}/4)T^3\sqrt{\ell}$, and $\ell \le 2L/\sqrt{N}$ gives $T^3 \ge L^{3/2}N^{-1/4}/(72\sqrt{12}\,C) \ge L^{3/2}N^{-1/4}/(250C)$. If $A < B$ then $H < 2B \le 4\log_2 N$ for $N \ge 4$, and $T^2 \ge L^2N^{-1/2}/(64C\log_2 N)$. Together with the far case this is the M2 bound. The exponent statements follow from $LN^{-1/4} = L^{3/5}N^{-1/10}$ at $L = N^{3/8}$ and $LN^{-1/4} = L^{1/2}N^{-1/12}$ at $L = N^{1/3}$. □

The M1 floor $L^{3/5}N^{-1/10}$ for $L \ge N^{3/8}(\log_2 N)^{5/4}$ is the cost of the localised search of Corollary 13.8 below: within the materialised-giant model, the localised Harvey search is optimal up to logarithmic factors for every such L. At the power-law level, if $L = \sqrt{N}/2^k$ in this regime, the common M1 scale is $N^{1/5}2^{-3k/5}$, so one additional known leading bit of p improves it by a factor $2^{3/5}$, up to C-dependent constants and logarithmic factors. Below $N^{3/8}$ the floor's exponent $\lambda - 1/4$ is that of Coppersmith subdivision, which lies outside the model; inside it, the single Fermat form $(1, 1)$ attains that exponent near $\sqrt{N}$: if $s = \sqrt{N}$ and $I = [s - 2L, s - L]$, then

$N/(s-d)+(s-d) = 2s+d^2/(s-d)$ for $L \le d \le 2L$, so $g(I)$ has length $O_C(L^2/\sqrt{N})$, consecutive tested integers covering it give $W = O_C(L^2/\sqrt{N})$, and the optimised M1 cost is $O_C(\sqrt{W}) = O_C(LN^{-1/4})$. Section 13.4 shows that the logarithm in the first branch can be removed, and that for all but a vanishing fraction of the centres the in-model cost below $N^{3/8}$ is the larger $\Theta_C(L^{3/5}N^{-1/10})$. In M2 the floor $L^{1/2}N^{-1/12}$ is not known to be attained by any family; Harvey's search [31] attains the M1 floor. For $\sqrt{N}/(4C) < L \le |J_C|$, Proposition 13.1 applied to I gives $T \ge c_2(L/N^{1/4})^{4/5} \ge c_2(4C)^{-4/5}N^{1/5}$ in M1, while $L < \sqrt{N}$ gives $L^{3/5}N^{-1/10} \le N^{1/5}$; so, after adjusting the C-dependent constant, the M1 bound $T \gg_C L^{3/5}N^{-1/10}$ holds for every $N^{1/4} \le L \le |J_C|$.

The scale $L = N^{1/4}$ at which the three lower-bound scales become constant is the standard known-prefix threshold for balanced factorisation. If $p, q = \Theta_C(N^{1/2})$ and one factor is known to lie in an interval of length $O_C(N^{1/4})$, then, after $O_C(1)$ subdivisions if necessary, N is factored in $(\log N)^{O(1)}$ time by Coppersmith's univariate small-root method in Howgrave-Graham's formulation [17, 42]; in binary terms this is, up to $O_C(1)$ bits and endpoint ambiguities, knowledge of the leading $\frac{1}{4}\log_2 N$ bits of p, about half of its bits, and conversely an interval of that length is covered by $O_C(1)$ aligned prefix intervals. This is a computational reduction, not an information-theoretic identity: factoring reveals the prefix trivially, and an algorithm recovering that prescribed prefix from N can be followed by Coppersmith's method; nothing is asserted for an arbitrary choice of $\frac{1}{4}\log_2 N$ bit positions, and a measurable set S of measure $N^{1/4}$ need not supply Coppersmith information at all, since it may be disconnected or lack an efficient interval description. Proposition 13.1 identifies the same localisation scale numerically, within the covering model only.

Within that model the proposition places non-oblivious coverage precisely. Information about N that excludes a fixed positive proportion of J_C changes the constants and not the exponents; to change an exponent through interval localisation one must reduce the uncertainty in p from $N^{1/2}$ to $N^{1/2-\varepsilon}$, i.e. determine, up to $O_C(1)$ terms, $\varepsilon \log_2 N$ additional leading bits of p. We are not aware of a general worst-case $(\log N)^{O(1)}$-time algorithm that, from a balanced semiprime N alone, recovers any additional prescribed leading-prefix bit of its smaller factor beyond what the size of N and the balance bound C already force – the known-bits literature [72, 17, 36] assumes the bits given and does not recover them – nor of a general reduction showing that a leading prefix shorter than $\frac{1}{4}\log_2 N + O_C(1)$ bits suffices to factor; this statement excludes results that assume auxiliary leakage, special bit positions, average-case distributions or additional algebraic relations. The floors of this section are the $\varepsilon = 0$ end of that interpolation, and the search for a non-oblivious exclusion is, inside the covering paradigm, the search for a partial-information factoring algorithm. The floors are attained, up to constants and logarithms, at $L = |J_C|$ by Harvey's algorithm. For shorter intervals a localised version of his search gives an upper bound with a different exponent, which we now prove, first given an element of large order and then, with that element supplied by the theorem of Harvey and Hittmeir recalled in Section 13.2, unconditionally; it rests on the observation that a cell covering p has its ratio a/b close to $p/q = p^2/N$.

Lemma 13.3 (proximity)**.** *Let $N = pq$ with $0 < p \le q$, let $2 \le r \le N^{1/3}/4$, and let the cell (a, b) of the standard family (Definition 2.1: $a, b \ge 1$, $ab \le r$, tested integers $c_{ab}, \dots, c_{ab} + W_{ab}^{\mathrm{int}} - 1$ with $c_{ab} = \lceil 2\sqrt{abN} \rceil$, $W_{ab}^{\mathrm{int}} = \lceil \sqrt{N}/(4r\sqrt{ab}) \rceil$) cover p: $|aq + bp - n| < 1$ for some tested n. Then $|a/b - p/q| < 2/(b\sqrt{r})$.*

Proof. Put $W := \sqrt{N}/(4r\sqrt{ab})$ and $W' := W + 2$. Since $aq + bp \ge 2\sqrt{abN}$ and $aq + bp < n + 1 \le c_{ab} + W_{ab}^{\mathrm{int}} < 2\sqrt{abN} + W + 2$, we have $(\sqrt{aq} - \sqrt{bp})^2 < W'$. With $m := |aq - bp|$, $m^2 = (\sqrt{aq} - \sqrt{bp})^2(\sqrt{aq} + \sqrt{bp})^2 < W'(W' + 4\sqrt{abN}) = W'^2 + N/r + 8\sqrt{abN}$, using $4W\sqrt{abN} =$

N/r. Now $8\sqrt{abN} \le 8\sqrt{rN} \le N/r$ for $r \le N^{1/3}/4$. The range $2 \le r \le N^{1/3}/4$ is nonempty only for $N \ge 512$, and then $\sqrt{N}/(4r) \ge N^{1/6} \ge 2$, so $W' \le \sqrt{N}/(4r) + 2 \le \sqrt{N}/(2r)$ and $W'^2 \le N/(4r^2) \le N/(8r)$. Hence $m^2 < 3N/r$, and $|a/b - p/q| = m/(bq) < \sqrt{3}\sqrt{N/r}/(bq) = \sqrt{3}\sqrt{p/(rq)}/b < 2/(b\sqrt{r})$ since $p \le q$. □

13.1. **The localised search.** For an interval $I \subseteq J_C$ put $R(I) := \{x^2/N : x \in I\}$ and define the *localised family*

$$\mathcal{F}(I, r) := \{(a, b) : a, b \ge 1,\ ab \le r,\ \operatorname{dist}(a/b, R(I)) < 2/(b\sqrt{r})\},$$

with the standard tested integers. If $|I| = L$ then $|R(I)| = L(2\inf I + L)/N \le 2L/\sqrt{N}$. The family is enumerated by listing $b \le \sqrt{2Cr}$ and, for each b, the integers a in an explicit interval, in $O(|\mathcal{F}(I, r)| + \sqrt{r})$ arithmetic operations on $O(\log N)$-bit numbers when the endpoints of I are $O(\log N)$-bit rationals.

Lemma 13.4 (localised covering and count)**.** *Let $16C^2 \le r \le N^{1/3}/4$ and $I \subseteq J_C$ of length L. (i) Every $p \in I$ is covered by a cell of $\mathcal{F}(I, r)$, and that cell satisfies the conclusion of Harvey's covering lemma* [31, Lemma 3.3]*. (ii) $P := |\mathcal{F}(I, r)| \le 2CLr/\sqrt{N} + 3\sqrt{2Cr}$. (iii) $\Sigma_W(\mathcal{F}(I, r)) \le CL/\sqrt{r} + \sqrt{2CN}\,\log(6Cr)/(2r) + P$.*

Proof. (i) Since $r \ge 16C^2 \ge C$, every $p \in J_C$ satisfies $(N/r)^{1/2} \le p < N^{1/2}$, and Harvey's Lemma 3.3 gives a cell (a, b) with $ab \le r$ and $0 \le aq + bp - \lceil 2\sqrt{abN}\rceil < \sqrt{N}/(4r\sqrt{ab})$; its tested integer $aq + bp$ lies in the standard window, so it covers p in the sense of Definition 2.1, and by Lemma 13.3 $\operatorname{dist}(a/b, R(I)) \le |a/b - p/q| < 2/(b\sqrt{r})$, as $p/q = p^2/N \in R(I)$. (ii) A cell of $\mathcal{F}(I, r)$ has $a/b \ge \inf R(I) - 2/\sqrt{r} \ge 1/C - 2/\sqrt{r} \ge 1/(2C)$, so $b^2 \le 2Cab \le 2Cr$ and $b \le B := \sqrt{2Cr}$. For fixed b the admissible a lie in an interval of length $b(|R(I)| + 4/(b\sqrt{r})) \le b\ell + 1$, $\ell := 2L/\sqrt{N}$, so there are at most $b\ell + 2$ of them; summing, $P \le \ell(B^2 + B)/2 + 2B \le C\ell r + 3\sqrt{2Cr}$ as $\ell \le 2$. (iii) $W_{ab}^{\mathrm{int}} \le \sqrt{N}/(4r\sqrt{ab}) + 1 \le \sqrt{2CN}/(4rb) + 1$ since $\sqrt{ab} \ge b/\sqrt{2C}$; hence $\Sigma_W \le \sum_{b \le B}(b\ell + 2)(\sqrt{2CN}/(4rb) + 1) \le B\ell\sqrt{2CN}/(4r) + (\sqrt{2CN}/(2r))(1 + \log B) + P$, and $B\ell\sqrt{2CN}/(4r) = CL/\sqrt{r}$, $1 + \log B \le \log(6Cr)$. □

Lemma 13.5 (list-restricted Harvey search)**.** *Let N be either prime or a semiprime pq. Let $\mathcal{L}$ be an explicitly enumerated list of forms (a, b) with $a, b \ge 1$, each carrying a window $\{n_{ab}, n_{ab} + 1, \dots, n_{ab} + W_{ab} - 1\}$ of consecutive integers – the standard window $n_{ab} = \lceil 2\sqrt{abN}\rceil$, $W_{ab} = W_{ab}^{\mathrm{int}}$, or any other explicitly given one; let $m \ge 1$ and $\alpha \in (\mathbb{Z}/N\mathbb{Z})^\times$ with $\operatorname{ord}_N(\alpha) > m$; and assume the list sizes below are $O(N)$. Run Harvey's Algorithm 4.2* [31] *with its Step (2) restricted to $\mathcal{L}$, the giant of the j-th block of the form (a, b) being $v_{a,b,j} := \alpha^{aN+b-n_{ab}-jm}$. If $N = pq$ and the window of some form of $\mathcal{L}$ contains $aq + bp$ – which the conclusion of* [31, Lemma 3.3] *guarantees for the standard family when $(N/r)^{1/2} \le p < N^{1/2}$ – the algorithm returns p and q; its bit complexity is $O((m + G + |\mathcal{L}|)(\log N)^{O(1)})$, where $G := \sum_{(a,b)\in\mathcal{L}} \lceil W_{ab}/m\rceil \le |\mathcal{L}| + \Sigma_W(\mathcal{L})/m$.*

Proof. In Harvey's proof of [31, Proposition 4.2], Lemma 3.3 is used only to supply a form (a_0, b_0) whose window contains $a_0q + b_0p = n_{a_0b_0} + y_0$, $0 \le y_0 < W_{a_0b_0}$, and a decomposition $y_0 = i_0 + j_0m$, $0 \le i_0 < m$, for which the giant v_{a_0,b_0,j_0} is congruent to α^{i_0} modulo p, because $\alpha^{a_0N+b_0} \equiv \alpha^{a_0q+b_0p}$ (mod p) by Fermat's theorem; by hypothesis that form belongs to $\mathcal{L}$, and nothing else in the argument depends on the position or the length of the windows. If $\operatorname{ord}_p(\alpha) < m$, Step (1) finds a factor; otherwise it certifies $\operatorname{ord}_p(\alpha) \ge m$, which makes $\alpha^0, \dots, \alpha^{m-1}$ distinct modulo p. Step (3) finds every equality $v_{a,b,j} = \alpha^i$ in $\mathbb{Z}/N\mathbb{Z}$ by sorting and uses the attached labels (a, b, j) to test the corresponding candidate $n_{ab} + jm + i$ with [31, Lemma 3.1]; after these exact matches are removed, the remaining giants satisfy the separation precondition of Algorithm 4.1, so Step (4) detects the covering collision modulo p. No other part of the correctness proof uses cells

outside $\mathcal{L}$. Step (2a) performs $O(|\mathcal{L}|)$ modular exponentiations, Step (2b) produces G giants, Step (3) sorts lists of total length $m+G$, and Step (4) has the cost of [31, Proposition 4.1] with $n \le G$. □

Theorem 13.6 (localised search, given an element of large order)**.** *Fix $C > 1$. There is a deterministic algorithm which, given $N = pq$ with $p \in J_C$ and $N \ge N_0(C)$, an interval $I \subseteq J_C$ with $O(\log N)$-bit rational endpoints containing p and of length L with $N^{3/8} \le L \le |J_C|$, and an element $\alpha \in (\mathbb{Z}/N\mathbb{Z})^\times$ with $\operatorname{ord}_N(\alpha) > m := \max\{1, \lfloor L^{3/5}N^{-1/10} \rfloor\}$, factors N in*

$$O_C\big(L^{3/5}N^{-1/10}\big)(\log N)^{O(1)}$$

bit operations: Lemma 13.5 on $\mathcal{F}(I,r)$ with $r = \lceil (N/L)^{2/5} \rceil$ and this m.

Proof. Put $A_0 := L^{3/5}N^{-1/10}$, $R_0 := (N/L)^{1/5}$ and $B_0 := L^{1/5}N^{1/20}\sqrt{\log N}$. In the stated range $N^{1/5} < r \le 2N^{1/4}$, so $16C^2 \le r \le N^{1/3}/4$ for N large, and $p \in J_C$ gives $(N/r)^{1/2} \le p < N^{1/2}$; Lemma 13.4(i) supplies a cell of $\mathcal{F}(I,r)$ satisfying Harvey's Lemma 3.3. By hypothesis $\operatorname{ord}_N(\alpha) > m$, so Lemma 13.5 applies (since $L < \sqrt{N}$ one has $A_0 < N^{1/5}$, so any element of order $> N^{1/5}$ qualifies). Lemma 13.4 gives $P \ll_C A_0 + R_0$ and $\Sigma_W \ll_C A_0^2 + B_0^2 + P$; for $L \ge N^{3/8}$, $R_0 \le A_0$ and $B_0 \le A_0\sqrt{\log N}$, and $m \asymp A_0$ gives $G \le P + \Sigma_W/m \ll_C A_0 \log N$. The enumeration of the family costs $O(P + \sqrt{r})$ operations with $\sqrt{r} \asymp R_0 \le A_0$. Lemma 13.5 therefore gives total cost $O_C(A_0)(\log N)^{O(1)}$. □

13.2. **The order element.** The hypothesis on α in Theorem 13.6 is not a formality. Harvey's Proposition 2.7 [31], quoting Hittmeir's Theorem 6.3 [39], proves that for $N^{2/5} \le D \le N$ one can deterministically return a factor, prove primality, or produce $\alpha \in (\mathbb{Z}/N\mathbb{Z})^\times$ with $\operatorname{ord}_N(\alpha) > D$, in $O(D^{1/2}\log^2 N)$ time; Harvey's Remark 2.8 reports, citing Hittmeir's Remark 6.4, that the same conclusion can be obtained for $D = \Omega(N^{1/3+o(1)})$ with suitable $o(1)$. The printed proposition therefore gives an $N^{1/5+o(1)}$ order-element construction and the reported extension an $N^{1/6+o(1)}$ one. For a given unit α, baby-step giant-step certifies $\operatorname{ord}_N(\alpha) > D$, or finds a relation of exponent at most D, in $\tilde{O}(\sqrt{D})$ operations, and with gcds inserted unequal component orders modulo p and q reveal a factor; the residues on which this neither factors nor certifies satisfy $\operatorname{ord}_p(\alpha) = \operatorname{ord}_q(\alpha) = d \le D$, and number at most $\sum_{d \le D} \varphi(d)^2 = O(D^3)$, a cardinality bound that says nothing about how many consecutive small integers the exceptional set can contain; the difficulty is to bound the number of candidates that must be tried. The hypothesis on D has since been weakened and then removed: Gao, Feng, Hu and Pan [27, Lemma 3.5] obtain $D \ge N^{1/4+o(1)}$ and Oznovich and Volk [67, Theorem 1.1] $D \ge N^{1/6}$, both by a better bound $O(\sqrt{m})$ for the number of consecutive candidates of the same order m, and Harvey and Hittmeir have proved the following.

Theorem 13.7 (Harvey–Hittmeir [33, Theorem 1.1])**.** *There is a deterministic algorithm which, given integers $N \ge 3$ and $1 \le D < N-1$, returns either $\alpha \in (\mathbb{Z}/N\mathbb{Z})^\times$ with $\operatorname{ord}_N(\alpha) > D$ or a nontrivial divisor of N – never a primality proof – in $O\big(D^{1/2}\log \bar{D}\,\log N/(\log\log \bar{D})^{1/2}\big)$ bit operations, where $\bar{D} := \max(D,3)$ (their convention for small D; for bounded D the cost is $O(\log N)$).*

Nir [65] has an independent weaker version requiring $D > \exp(\sqrt{2\log N \log\log N})$. All of these are upper bounds for one family of constructions, not lower bounds on the cost of producing such an element. Theorem 13.7 is used below as a black box. For balanced semiprimes there is a short proof of the same statement with a different recovery step, given in Section 16 as Proposition 16.1 together with a variant for explicit baby sets (Lemma 16.2); we include it because its two devices – the accumulation of the least common multiple of the orders of failing candidates, and

the recovery of p and q from $N-1=(k+l)L+klL^2$ once $L\geq N^{1/3}$ – are reused, and because the corollaries below run on either result.

13.3. **The landscape.**

Corollary 13.8 (localised search, unconditional)**.** *In Theorem 13.6 the element α may be dropped from the input: Theorem 13.7 with $D:=m$ supplies it or a nontrivial divisor of N at cost $O(m^{1/2})(\log N)^{O(1)}$ (Proposition 16.1 does the same for balanced N), and the total remains $O_C(L^{3/5}N^{-1/10})(\log N)^{O(1)}$.*

Corollary 13.9 (the partial-information landscape)**.** *Fix $C>1$, let $N=pq$ with $p\in J_C$, and suppose p is known to lie in an interval with $O(\log N)$-bit rational endpoints and length $L=N^\lambda$, $1/4\leq\lambda\leq 1/2$ (with $I\subseteq J_C$, so that $L<\sqrt{N}$ and $\lambda=1/2$ is to be read as $L\asymp_C\sqrt{N}$). Deterministic factorisation has the unconditional upper bound*

$$N^{o(1)}\min\{LN^{-1/4},\ L^{3/5}N^{-1/10}\}:$$

the first term subdivides the interval into $O_C(1+L/N^{1/4})$ subintervals of Coppersmith length [17, 42]; the second is Corollary 13.8 for $L\geq N^{3/8}$ and Corollary 13.16 for $c_6N^{1/6}\leq L<N^{3/8}$ (the Coppersmith branch is the smaller throughout $L<N^{3/8}$). In exponents the curve is $\lambda-1/4$ on $[1/4,3/8]$ and $3\lambda/5-1/10$ on $[3/8,1/2]$; it lies strictly below Harvey's $1/5$ for every $\lambda<1/2$ and below Coppersmith-plus-search for $\lambda>3/8$. By Theorem 13.2 the M1 cell-model floor has the same exponent at every $\lambda\in[1/4,1/2]$. For $\lambda\geq 3/8$ the localised search attains it inside the model, so there the materialised-giant landscape is settled up to logarithmic factors. For $\lambda<3/8$ the floor's exponent is attained by Coppersmith subdivision, which lies outside the cell model, and inside the model on the intervals of bounded depth: by Theorem 13.15 the in-model cost of an interval is $\Theta_C(\min\{\sqrt{\mathfrak{D}(I)}LN^{-1/4},L^{3/5}N^{-1/10}\})$ in terms of its depth $\mathfrak{D}(I)$, so every interval has cost $O_C(L^{3/5}N^{-1/10})$, attained by a Farey-fraction family (Corollary 13.16, an in-model algorithm down to $L\asymp N^{1/6}$), that scale is the cost for all but a vanishing fraction of the centres, and for fixed $\lambda<3/8$ the intervals of in-model cost $\Theta_C(LN^{-1/4})$ are exactly those of bounded depth. The floor is in any case not a lower bound for arbitrary algorithms.

The collision realisation behind M1 presupposes α and the cost models do not charge its selection; Theorem 13.7 shows that for consecutive babies there is nothing to charge. An explicitly materialised Harvey-type search of target cost N^δ needs $\mathrm{ord}_N(\alpha)>m$ with $m\leq N^{\delta+o(1)}$, and the theorem supplies such an element, or factors N, in $N^{\delta/2+o(1)}$ bit operations, for every $\delta>0$. For an explicitly represented baby set the relevant condition is pairwise distinctness of the babies modulo each prime, and it too can be supplied at a cost subordinate to the search: Lemma 16.2 (Section 16) returns a factor or a unit whose powers along a given set B of spread S are pairwise distinct modulo both primes in $\tilde{O}(|B|+\sqrt{S})$ operations, an overhead subordinate to a target cost T whenever $|B|,\sqrt{S}\leq TN^{o(1)}$; no claim is made for implicit baby sets, for sets of larger spread, or for collision algorithms needing more than the pairwise separation of one baby set. Within that scope the supply of an element of large order, which the constructions available before 2026 priced at $N^{1/5+o(1)}$ or (as reported) $N^{1/6+o(1)}$, is not an obstruction to Fermat-collision algorithms at any exponent – the conclusion Harvey and Hittmeir draw for general N [33]: in any exponential-time deterministic factoring algorithm, finding elements of large order is no longer a bottleneck, whatever the exponent. What remains are the covering floors and the difference-cover and product questions of Part 1.

13.4. **Below $N^{3/8}$: the landscape depends on the interval.** Theorem 13.2 is a worst case over intervals of a given length. In this subsection the interval is $I=[p_0-L/2,p_0+L/2]\subseteq J_C$ and the quantities are expressed through $\xi(p):=N/p^2$, the reciprocal of the ratio coordinate

of the proof of Theorem 13.2: a form (a,b) has $g'(p)=b-aN/p^2=a\,(b/a-\xi(p))$, its critical point $p^*=\sqrt{aN/b}$ lies in I iff $b/a\in[\xi_1,\xi_2]:=[\xi(p_0+L/2),\xi(p_0-L/2)]$, and we write $d=d(a,b):=\operatorname{dist}(b/a,[\xi_1,\xi_2])$, so that $|g'|\ge ad$ on I. Since $|\xi'(p)|=2N/p^3\le 2C^{3/2}/\sqrt{N}$ on J_C, $\eta_0:=\xi_2-\xi_1\le 2C^{3/2}L/\sqrt{N}$, and conversely $|dp/d\xi|\le (C+1)^{3/2}\sqrt{N}/2$ for $\xi\ge 1/(C+1)$. Throughout, families are as in Theorem 13.2 (they may depend on I but not on $p\in I$), T is the M1 cost, and we put

$$P_0:=L^{3/5}N^{-1/10},\qquad Q:=(N/L)^{1/5},$$

so that $LQ^2/\sqrt{N}=P_0$, $L/Q=P_0^2$ and $\sqrt{N}/Q^2=L/P_0$. Two reductions are used. *Primitive reduction:* if $(a,b)=k(a',b')$ with $k\ge 2$ then $g_{a,b}=kg_{a',b'}$, so the form (a,b) with candidate set S covers the same set as (a',b') with the candidate set S/k, of measure w/k; merging every form into its primitive representative, with the union of the rescaled candidate sets, leaves at most P distinct primitive forms of total candidate measure at most Σ_w covering the same set, and preserves every ratio b/a and hence every d; Lemma 3.1 and the monotone bound hold for measurable candidate sets, and the cost bounds below hold a fortiori for the reduced quantities, so lower bounds may assume the forms primitive and distinct. *Cost:* every nonempty family has $T\ge m+\lceil W_1/m\rceil\ge 2$; $T\ge m+\sum_i\lceil W_i/m\rceil$ gives $T\ge P$ and $T\ge 2\sqrt{\Sigma_W}\ge\sqrt{2\Sigma_w}$; conversely a family with P forms and Σ_W tested integers has $T\le P+2\sqrt{\Sigma_W}+1$ at $m=\lceil\sqrt{\Sigma_W}\rceil$, and a single form with W tested integers has $T\le 2\sqrt{W}+1$.

Lemma 13.10 (localised covering sum)**.** *Let $\rho>0$; call a form* near *if $d\le\rho$ and* far *otherwise, and put $H_\rho:=\sum_{\text{near}}1/a_i$. If the family covers I then $L\le\Sigma_w/\rho+2N^{1/4}\sqrt{\Sigma_w H_\rho}$; consequently $\Sigma_w\ge\min\{\rho L/2,\ L^2/(16\sqrt{N}H_\rho)\}$.*

Proof. A far form has $|g'|\ge a_id_i>\rho$ on I, hence covers at most w_i/ρ; a near form covers at most $2\sqrt{w_i\sqrt{N}/a_i}$ (Lemma 3.1), and Cauchy–Schwarz bounds the sum over near forms by $2N^{1/4}\sqrt{\Sigma_wH_\rho}$. If $\Sigma_w<\rho L/2$ the far forms cover less than $L/2$, so $L/2\le 2N^{1/4}\sqrt{\Sigma_wH_\rho}$. □

Lemma 13.11 (Farey spacing)**.** *An interval of length η contains at most ηD^2+1 reduced fractions with denominator at most D. Hence if the primitive forms with b/a in an interval of length η are listed by nondecreasing denominator $a_1\le a_2\le\cdots$, then $a_k\ge\sqrt{(k-1)/\eta}$, and for every P and every $D\le a_1$, $\sum_{k\le P}1/a_k\le\min\{P/D,\ 1/D+2\sqrt{\eta P}\}$.*

Proof. Distinct reduced fractions with denominators at most D differ by at least $1/D^2$. Then $\sum_{k\le P}1/a_k\le 1/a_1+\sum_{k=2}^{P}\sqrt{\eta/(k-1)}\le 1/D+2\sqrt{\eta P}$, and trivially $\le P/D$. □

Definition 13.12 (Diophantine depth)**.** The *depth* of I is

$$\mathfrak{D}(I):=\min_{a,b\ge 1}\ a\cdot\max\Big\{1,\ \frac{\sqrt{N}\,d(a,b)}{2L}\Big\}.$$

A form whose critical point lies within about L of I ($d\le 2L/\sqrt{N}$) counts with its denominator, one at p-distance δ with about $a\delta/L$; the depth is the effective denominator of the best single form. Always $1\le\mathfrak{D}(I)\le(\sqrt{N}/(2L))^{1/2}+1$: with $D:=\lceil(\sqrt{N}/(2L))^{1/2}\rceil$, Dirichlet's theorem gives b/a with $a\le D$ and $d\le 1/(aD)$, whence $a\max\{1,\sqrt{N}d/(2L)\}\le\max\{D,\sqrt{N}/(2LD)\}\le D$. A form with $b/a\notin[1/2,C+2]$ has $d\ge 1/2$ and value at least $\sqrt{N}/(4L)$, which exceeds $(\sqrt{N}/(2L))^{1/2}+1$ once $\sqrt{N}/(2L)\ge 8$ (the exact threshold is $(1+\sqrt{3})^2\approx 7.46$); so for $N\ge N_0(C)$ the minimum is attained by a form with $b/a\in[1/2,C+2]$, and by finitely many candidates.

Lemma 13.13 (one form)**.** *There is $c_7=c_7(C)$ such that for $N\ge N_0(C)$ every form (a,b) with $b/a\in[1/2,C+2]$ covers I alone at cost $T\le 3+c_7\sqrt{a\max\{1,\sqrt{N}d(a,b)/(2L)\}}\,LN^{-1/4}$; in particular $T(I)\le 3+c_7\sqrt{\mathfrak{D}(I)}\,LN^{-1/4}$.*

Proof. If $d \le 2L/\sqrt{N}$, the critical point p^* lies within $(C+2)^{3/2}L$ of I, so every $p \in I$ is within $(1+(C+2)^{3/2})L$ of p^*; on the relevant range $g'' = 2aN/p^3 \le 2a(C+2)^{3/2}/\sqrt{N}$, so $g(p) - g(p^*) \le X := a(C+2)^{3/2}(1+(C+2)^{3/2})^2L^2/\sqrt{N}$ on I, and the window of $1+\lfloor X\rfloor$ consecutive integers from $\lceil g(p^*)\rceil$ has candidate set containing $[g(p^*), g(p^*)+1+\lfloor X\rfloor) \supseteq g(I)$; the cost is at most $2\sqrt{1+\lfloor X\rfloor}+1 \le 3+2\sqrt{X}$. If $d > 2L/\sqrt{N}$, g is monotone on I with $|g'| = a|b/a - \xi(p)| \le a(d+\eta_0) \le ad(1+C^{3/2})$, so $g(I)$ has length at most $X := ad(1+C^{3/2})L$, the window of $1+\lfloor X\rfloor$ integers from $\lceil \min_I g\rceil$ covers I, and the cost is at most $3+2\sqrt{X} = 3+2\sqrt{2(1+C^{3/2})}\sqrt{a\sqrt{N}d/(2L)}\,LN^{-1/4}$. Take $c_7 := \max\{2(C+2)^{3/4}(1+(C+2)^{3/2}), 2\sqrt{2(1+C^{3/2})}\}$; the minimising form of Definition 13.12 has $b/a \in [1/2, C+2]$ for $N \ge N_0(C)$, as noted after the definition. □

Lemma 13.14 (the Farey family)**.** *There are c_2, c_3, c_4, c_6 depending on C such that for $N \ge N_0(C)$ and $c_6N^{1/6} \le L \le N^{3/8}$, every I is covered by a family of at most c_2P_0 forms with $\Sigma_W \le c_3P_0^2$, hence with $T \le c_4P_0 = c_4L^{3/5}N^{-1/10}$.*

Theorem 13.15 (the cost of an interval below $N^{3/8}$)**.** *Fix $C > 1$. There are $c_8, c_9, c_{10}, c_6 > 0$ depending only on C and $N_0(C)$ such that for $N \ge N_0(C)$, $c_6N^{1/6} \le L \le N^{3/8}$ and every interval $I \subseteq J_C$ of length L,*

$$c_8\big(1+\min\{\sqrt{\mathfrak{D}(I)}\,LN^{-1/4},\ P_0\}\big) \ \le\ T(I)\ \le\ c_9\big(1+\min\{\sqrt{\mathfrak{D}(I)}\,LN^{-1/4},\ P_0\}\big),$$

where $P_0 = L^{3/5}N^{-1/10}$ as above and $T(I)$ is the least M1 cost of a family covering I; for $L \ge N^{1/4}$ the additive 1 is absorbed, $LN^{-1/4}$ being at least 1, and $T(I) \asymp_C \min\{\sqrt{\mathfrak{D}(I)}\,LN^{-1/4}, P_0\}$. The two branches meet at $\mathfrak{D}(I) = D_1 := N^{3/10}L^{-4/5}$, and the centres p_0 with $\mathfrak{D}(I) < D_1$ form a set of measure at most $c_{10}N^{3/5}L^{-3/5}$. Consequently: every interval has $T(I) \ge c_8LN^{-1/4}$ (the first branch of Theorem 13.2 without its logarithm) and $T(I) \le c_4P_0$; an interval containing the critical point $\sqrt{aN/b}$ of a form with $a \le K$ has $\mathfrak{D}(I) \le K$, one within L of it has $\mathfrak{D}(I) \le 2C^{3/2}K$, and in either case $T(I) \le 3 + c_7\sqrt{2C^{3/2}K}\,LN^{-1/4}$; and all but a fraction $O_C(N^{1/10}L^{-3/5})$ of the centres have $T(I) \asymp_C P_0$, a fraction which tends to zero when $L/N^{1/6} \to \infty$, in particular for $L = N^\lambda$ with any fixed $\lambda > 1/6$.

Proof of Lemma 13.14. Let $R := \lfloor Q\rfloor \ge Q/2$ and let $\mathcal{F}$ consist of the reduced fractions b/a with $a \le R$ in $[\xi_1, \xi_2]$ together with the largest such fraction below ξ_1 and the smallest above ξ_2. Consecutive members of $\mathcal{F}$ are consecutive among the reduced fractions of denominator at most R on the real line – the Farey sequence of order R, extended beyond $[0,1]$ by integer translation – so by the mediant property their denominators satisfy $a+a' \ge R+1$ and their gap is $1/(aa')$; with $a_{\min} := \min(a,a')$ the gap is at most $4/(a_{\min}R) \le 8/(a_{\min}Q)$ (if $a_{\min} \le R/2$ then $aa' \ge a_{\min}R/2$, otherwise $aa' > R^2/4 \ge a_{\min}R/4$). Assign each gap to its endpoint of smaller denominator. Under $\xi \mapsto p = \sqrt{N/\xi}$ the gaps between the critical points of consecutive members partition I into pieces, each assigned to a form $f = (a,b) \in \mathcal{F}$ and of p-length at most $G_a := 4(C+1)^{3/2}\sqrt{N}/(aQ)$; every member receives at most two pieces, and the two exterior members (critical points outside I) at most one each. For an interior member the assigned pieces lie within $r_f := \min(2G_a, L)$ of p_f^*, and as in Lemma 13.13 the window of $W_f := 1+\lfloor 3a(C+1)^{3/2}r_f^2/\sqrt{N}\rfloor$ integers from $\lceil g_f(p_f^*)\rceil$ covers them; for an exterior member at distance $\delta_f \le G_a$ from I whose piece has length $\ell_f \le \min(G_a, L)$, g_f is monotone on I with $|g_f'| \le 3a(C+1)^{3/2}(\delta_f+\ell_f)/\sqrt{N}$ on the assigned piece, and the window of $W_f := 1+\lfloor 3a(C+1)^{3/2}(\delta_f+\ell_f)\ell_f/\sqrt{N}\rfloor$ integers from the least integer above the minimum of g_f on the piece covers it (an exterior member receiving no piece gets $W_f := 1$). The family covers I. Its size is $P \le \eta_0R^2+3 \le 2C^{3/2}P_0+3$ by Lemma 13.11. For Σ_W: interior members with $a \ge D_2 := 4(C+1)^{3/2}\sqrt{N}/(LQ)$ have $G_a \le L$ and contribute at

most $1+12a(C+1)^{3/2}G_a^2/\sqrt{N} = 1+192(C+1)^{9/2}\sqrt{N}/(aQ^2)$ each, and by Lemma 13.11 (with $D=D_2$ and at most η_0R^2+1 fractions) their harmonic sum is at most $1/D_2+2\eta_0R+2\sqrt{\eta_0}$; since $(\sqrt{N}/Q^2)/D_2 = L/(4(C+1)^{3/2}Q) \le P_0^2$, $(\sqrt{N}/Q^2)\eta_0R \le 2C^{3/2}L/Q = 2C^{3/2}P_0^2$ and $(\sqrt{N}/Q^2)\sqrt{\eta_0} \le \sqrt{2}C^{3/4}L^{9/10}N^{-3/20} \le P_0^2$ for $L \ge c_6N^{1/6}$, their total is $\ll_C P+P_0^2$. Interior members with $a<D_2$ contribute at most $1+12a(C+1)^{3/2}L^2/\sqrt{N}$ each and number at most $\eta_0D_2^2+1$, so their total is at most their number, which is at most P, plus $\ll_C L^2(\eta_0D_2^3+D_2)/\sqrt{N} \ll_C P_0+P_0^2$, using $L^2\eta_0D_2^3/\sqrt{N} \ll_C \sqrt{N}/Q^3 = P_0$ and $L^2D_2/\sqrt{N} \ll_C L/Q = P_0^2$. The two exterior members contribute at most $1+6a(C+1)^{3/2}G_a\min(G_a,L)/\sqrt{N}$ each, which is $\ll_C 1+P_0^2$: if $G_a \le L$ then $a \ge D_2$ and $aG_a^2/\sqrt{N} \le 16(C+1)^3\sqrt{N}/(D_2Q^2) = 4(C+1)^{3/2}P_0^2$, and otherwise $aG_aL/\sqrt{N} = 4(C+1)^{3/2}L/Q = 4(C+1)^{3/2}P_0^2$. Since $P \ll_C P_0$, choosing c_6 so that $P_0 \ge 1$ gives $\Sigma_W \le c_3P_0^2$ and $T \le P+2\sqrt{\Sigma_W}+1 \le c_4P_0$. □

Proof of Theorem 13.15. The upper bound is Lemma 13.13 or Lemma 13.14, whichever is smaller: $T(I) \le \min\{3+c_7\sqrt{\mathfrak{D}}\,LN^{-1/4}, c_4P_0\} \le \max\{3,c_7,c_4\}(1+\min\{\sqrt{\mathfrak{D}}LN^{-1/4}, P_0\})$. For the lower bound, since every family has $T \ge 2$, it suffices to show $T \ge c\min\{\sqrt{\mathfrak{D}}LN^{-1/4}, P_0\}$ with some $c=c(C)\le 2$, for then $T \ge \max\{2, c\min\} \ge (c/2)(1+\min)$. Put $\mathfrak{D} := \mathfrak{D}(I)$, $\rho := 2L/\sqrt{N}$, $\eta := \eta_0+2\rho \le c_\eta L/\sqrt{N}$ with $c_\eta := 2C^{3/2}+4$, and take $c_6 \ge (4c_\eta)^{5/3}$ in addition to the requirements of Lemma 13.14, so that $\eta D_1^2 \le c_\eta N^{1/10}L^{-3/5} \le 1/4$ for $L \ge c_6N^{1/6}$. By the definition of the depth, a form with $d>\rho$ has $ad \ge (2L/\sqrt{N})\mathfrak{D}$ and a form with $d \le \rho$ has $a \ge \mathfrak{D}$; primitive reduction preserves d and replaces a by a divisor, and the definition takes the minimum over all forms, so both statements hold for the reduced forms. Let a reduced family of P forms cover I. *Far forms* ($d>\rho$) cover at most $\sum w_i/(a_id_i) \le \Sigma_w\sqrt{N}/(2L\mathfrak{D})$. If this is at least $L/2$ then $\Sigma_w \ge L^2\mathfrak{D}/\sqrt{N}$ and $T \ge \sqrt{2\Sigma_w} \ge \sqrt{2}\sqrt{\mathfrak{D}}\,LN^{-1/4}$, which is at least $\sqrt{2}\min\{\sqrt{\mathfrak{D}}LN^{-1/4}, P_0\}$. Otherwise the *near forms* ($d \le \rho$) cover at least $L/2$, so $L/2 \le 2N^{1/4}\sqrt{\Sigma_wH}$ with $H := \sum_{\text{near}} 1/a_i$ (Lemma 3.1 and Cauchy–Schwarz), i.e. $\Sigma_w \ge L^2/(16\sqrt{N}H)$. The near forms have distinct ratios in an interval of length η and denominators at least $\mathfrak{D}$, so listing them as $a_1 \le a_2 \le \cdots$, Lemma 13.11 gives $a_k \ge \max\{\mathfrak{D}, \sqrt{(k-1)/\eta}\}$: the terms with $\sqrt{(k-1)/\eta} \le \mathfrak{D}$ number at most $\eta\mathfrak{D}^2+1$ and contribute at most $1/\mathfrak{D}$ each, and the rest contribute at most $\sum_k \sqrt{\eta/(k-1)} \le 2\sqrt{\eta P}$, so

$$H \;\le\; \eta\mathfrak{D} + \frac{1}{\mathfrak{D}} + 2\sqrt{\eta P}.$$

Case $\mathfrak{D} \le D_1$. Then $\eta\mathfrak{D} \le \eta D_1^2/\mathfrak{D} \le 1/(4\mathfrak{D})$. If $P \le 1/(16\eta\mathfrak{D}^2)$ then $2\sqrt{\eta P} \le 1/(2\mathfrak{D})$, so $H \le 7/(4\mathfrak{D})$, $\Sigma_w \ge \mathfrak{D}L^2/(28\sqrt{N})$ and $T \ge \sqrt{2\Sigma_w} \ge \sqrt{\mathfrak{D}}\,LN^{-1/4}/\sqrt{14}$. Otherwise $T \ge P > 1/(16\eta\mathfrak{D}^2) \ge 1/(16\eta D_1^2) \ge P_0/(16c_\eta) \ge \sqrt{\mathfrak{D}}\,LN^{-1/4}/(16c_\eta)$, using $\eta \le c_\eta L/\sqrt{N}$ and $\sqrt{\mathfrak{D}}LN^{-1/4} \le \sqrt{D_1}LN^{-1/4} = P_0$. In this case the minimum in the theorem is $\sqrt{\mathfrak{D}}LN^{-1/4}$, and the bound is proved. *Case* $\mathfrak{D} > D_1$. Now the minimum is P_0. One of the three terms bounding H is at least $H/3$. If it is $1/\mathfrak{D}$, then $\Sigma_w \ge \mathfrak{D}L^2/(48\sqrt{N}) \ge D_1L^2/(48\sqrt{N}) = P_0^2/48$ and $T \ge P_0/\sqrt{24}$. If it is $2\sqrt{\eta P}$, then either $P \ge P_0$, and $T \ge P_0$, or $P < P_0$ and $\Sigma_w \ge L^2/(96\sqrt{N}\sqrt{\eta P_0}) \ge L^{3/2}N^{-1/4}/(96\sqrt{c_\eta P_0}) = P_0^2/(96\sqrt{c_\eta})$, so $T \ge \sqrt{2}P_0/(\sqrt{96}\,c_\eta^{1/4})$. If it is $\eta\mathfrak{D}$, then, using $\mathfrak{D} \le 2(\sqrt{N}/(2L))^{1/2}$, $\Sigma_w \ge L^2/(48\sqrt{N}\eta\mathfrak{D}) \ge \sqrt{2}L^{3/2}N^{-1/4}/(96c_\eta)$ and $T \ge 2^{3/4}L^{3/4}N^{-1/8}/\sqrt{96c_\eta} \ge 2^{3/4}P_0/\sqrt{96c_\eta}$, since $L^{3/4}N^{-1/8} \ge P_0$ for $L \ge N^{1/6}$. This proves the lower bound with c the least of the constants displayed and $c_8 := c/2$.

Consequences. $\mathfrak{D}(I) \ge 1$ gives the universal floor, and Lemma 13.14 the universal upper bound. If the critical point of (a,b) lies in I then $d(a,b)=0$ and $\mathfrak{D}(I) \le a$; if it lies within p-distance L of I, then, as $|\xi'(p)| \le 3C^{3/2}/\sqrt{N}$ on $[\sqrt{N/C}-L, \sqrt{N}]$ for N large, $d(a,b) \le$

$3C^{3/2}L/\sqrt{N}$, so $\sqrt{N}d/(2L) \le 3C^{3/2}/2$ and $\mathfrak{D}(I) \le 2C^{3/2}a$; Lemma 13.13 applied to that form gives the cost bound.

Measure. If $\mathfrak{D}(I) < D_1$, some form with $a < D_1$ has $a\max\{1, \sqrt{N}d/(2L)\} < D_1$, hence $d < 2LD_1/(a\sqrt{N})$, and $b/a \in [1/2, C+2]$ because $2LD_1/\sqrt{N} = 2N^{-1/5}L^{1/5} < 1/2$ for N large. For a fixed such fraction the condition places its critical point within p-distance $(C+2)^{3/2}LD_1/a$ of I, i.e. p_0 in an interval of length at most $L + 2(C+2)^{3/2}LD_1/a$; there are at most $(C+2)a$ fractions per denominator, so the centres with $\mathfrak{D}(I) < D_1$ have measure at most $\sum_{a<D_1}(C+2)a(L + 2(C+2)^{3/2}LD_1/a) \le 3(C+2)^{5/2}LD_1^2 = 3(C+2)^{5/2}N^{3/5}L^{-3/5}$, which is $c_{10}N^{3/5}L^{-3/5}$ and, relative to $|J_C| \asymp \sqrt{N}$, a fraction $O_C(N^{1/10}L^{-3/5})$. □

Corollary 13.16 (an in-model algorithm for every interval)**.** *Fix $C > 1$. Given $N = pq$ with $p \in J_C$, $N \ge N_0(C)$, and an interval $I \subseteq J_C$ with $O(\log N)$-bit rational endpoints containing p, of length L with $c_6N^{1/6} \le L \le |J_C|$, N is factored deterministically in $O_C(L^{3/5}N^{-1/10})(\log N)^{O(1)}$ bit operations by a family of cells lying inside the covering model.*

Proof. For $L \ge N^{3/8}$ this is Corollary 13.8. Otherwise take the Farey family of Lemma 13.14, enumerated in $O(P + \log N)$ operations – the fractions of order R in $[\xi_1, \xi_2]$ by the mediant recursion from the two exterior neighbours, which the continued fractions of the endpoints supply – with its windows of consecutive integers, and run Harvey's Algorithm 4.2 restricted to it as in Lemma 13.5, which admits arbitrary consecutive windows. The family covers every real $p \in I$ in the sense of Definition 2.1; for the prime divisor p the integer $aq + bp$ of the covering form lies within distance less than 1 of a tested integer, hence equals it, which is the hypothesis of the lemma. With $m := \max\{1, \lfloor P_0 \rfloor\}$ the giant count $G \le P + \Sigma_W/m$ is $O_C(P_0)$, and the element of order $> m$ comes from Theorem 13.7 at cost $O(\sqrt{P_0})(\log N)^{O(1)}$. □

The construction was checked on balanced semiprimes of 40, 50 and 60 bits: for twelve random centres in J_2 at each of $\lambda = 0.26, 0.30, 0.34, 0.375$, the family of the reduced fractions of order $Q' := \mathrm{round}((N/L)^{1/5})$ in $[\xi_1, \xi_2]$ and their two neighbours was built with the gap assignment of Lemma 13.14, each member receiving, in place of the lemma's bound W_f, the least window of consecutive integers whose conservative covered set $\{p : g_f(p_f^*) \le g_f(p) \le g_f(p_f^*) + W - 1\}$ contains its assigned pieces; the coverage of I was verified from these windows (in double precision, with a slack of $10^{-9}L$ at the joints), and the cost proxy $P + 2\sqrt{\Sigma_W}$ – an upper bound for the optimised M1 cost up to an additive 1 – was found between 1.1 and 4.7 times P_0 in every one of the 144 cases, with no trend in the size; on the interval $[\sqrt{N} - 3L/2, \sqrt{N} - L/2]$, of depth 1, the single form $(1,1)$ covers at 3.0–4.2 times $LN^{-1/4}$ while the Farey family costs from 5.5 to 380 there. At these sizes the hypotheses $N \ge N_0(C)$ and $L \ge c_6N^{1/6}$ are not met with the constants of the proofs; the checks concern the two constructions and their coverage, not the theorem's constants or its lower bound.

What the theorem does and does not say. The exponent curve $\min\{\lambda - 1/4, 3\lambda/5 - 1/10\}$ of Theorem 13.2 is the least in-model cost over the intervals of a given length – a floor valid for every interval – and, because Coppersmith subdivision works for every interval, also the deterministic upper bound; it is not the typical in-model cost. For fixed $\lambda \in [1/4, 3/8)$ the classification says that the least cost over positions is $\Theta_C(N^{\lambda - 1/4})$, attained exactly by the intervals of bounded depth, while the typical, and hence the largest, in-model cost is $\Theta_C(N^{3\lambda/5 - 1/10})$; at $\lambda = 3/8$ the two coincide for every interval. In the model a known leading bit of p is therefore worth a factor $2^{3/5}$ at every $\lambda \in [1/4, 1/2]$ for a typical interval, not only above $3/8$. Nothing here is a lower bound for algorithms outside the model.

14. Arithmetic thinning and the residue landscape

Thinned families – forms carrying windows whose candidate sets satisfy the coverage condition (3), each testing the intersection of its window with a union of residue classes modulo a common modulus q – were defined in Part 1 (Definition 2.2), where Lemma 2.3 shows that thinning by the classes compatible with N is worth only a factor $N^{o(1)}$ when q is at most half the shortest window, Lemma 2.4 handles the 2-adic part of Hittmeir's modulus, and Theorem 2.8 prices the enumerated-hypothesis thinning of Harvey and Hittmeir [32] at $(\varphi(\mathsf{q})/\mathsf{q})^{3/5}$. Here the floors for thinned families are stated in the localised form, and thinning then acts at full strength with genuine side information – the residue of p modulo M – which gives the partial-information landscape a second axis.

Proposition 14.1 (floors for thinned families)**.** *A thinned family of density θ, with q at most half the shortest window, whose windows satisfy (3) on J_C, has $\Sigma_W \ge (\theta c_C/4)^{2/3}N^{1/3}$, $T \ge (\theta c_C)^{2/5}N^{1/5}$ in M1 and $T \ge 2(\theta c_C/4)^{1/3}N^{1/6}$ in M2. If the windows satisfy (3) on an interval I as in Theorem 13.2, then in M1*

$$T \;\ge\; \min\Big\{\Big(\frac{\theta}{906C}\Big)^{2/5}L^{3/5}N^{-1/10},\; \Big(\frac{\theta}{128C\log_2 N}\Big)^{1/2}LN^{-1/4}\Big\}.$$

Proof. Lemma 3.2 applied to the windows gives $\Sigma_{\widetilde{w}} \ge c_C\sqrt{N}/\sqrt{P}$, so $\Sigma_{\widetilde{W}} \ge c_C\sqrt{N}/(2\sqrt{P})$ and $\Sigma_W \ge \theta c_C\sqrt{N}/(4\sqrt{P})$; with $\Sigma_W \ge P$ (Definition 2.2) this gives $\Sigma_W^{3/2} \ge \theta c_C\sqrt{N}/4$. In M1, every ceiling is at least one, so $T \ge P$, and $T^2 \ge 4\Sigma_W \ge \theta c_C\sqrt{N}/\sqrt{P} \ge \theta c_C\sqrt{N}/\sqrt{T}$, whence $T^{5/2} \ge \theta c_C\sqrt{N}$; in M2, $T \ge 2\sqrt{\Sigma_W}$. Under localisation, the proof of Theorem 13.2 applied to the windows gives in the first near sub-case $\Sigma_{\widetilde{w}}\sqrt{P} \ge L^{3/2}N^{-1/4}/(906C)$ and in the logarithmic sub-case $\Sigma_{\widetilde{w}} \ge L^2N^{-1/2}/(128C\log_2 N)$, while the far case gives $\Sigma_{\widetilde{w}} \ge L^2/\sqrt{CN}$; with $\Sigma_{\widetilde{w}} \le 2\Sigma_{\widetilde{W}} \le 4\Sigma_W/\theta \le T^2/\theta$ and $P \le T$ these are $T^{5/2} \ge \theta L^{3/2}N^{-1/4}/(906C)$, $T^2 \ge \theta L^2N^{-1/2}/(128C\log_2 N)$ and $T^2 \ge \theta L^2/\sqrt{CN}$, the last stronger than the second. □

Proposition 14.2 (a residue-class variant of Harvey's search)**.** *Let $N = pq$ with $p \in J_C$ and $N \ge N_0(C)$, and let $p_0 \equiv p \pmod M$ be given with $\gcd(p_0, M) = 1$ and $1 \le M \le N^{1/7}$. Then N is factored deterministically in $N^{1/5+o(1)}M^{-2/5}$ bit operations. If moreover p is known to lie in an interval I of length L with $N^{3/8}M^{1/4} \le L \le |J_C|$ and $M \le (L^{3/5}N^{-1/10})^{5/7}$, the cost is $N^{o(1)}L^{3/5}N^{-1/10}M^{-2/5}$. The construction adapts Harvey's Algorithms 4.1 and 4.2 [31].*

Proof. Since $q \equiv Np_0^{-1} \pmod M$, the true value $aq + bp$ of every cell is congruent to $\omega_{ab} := aNp_0^{-1} + bp_0 \pmod M$, which the algorithm can compute. With $c_{ab} = \lceil 2\sqrt{abN}\rceil$ and $\widetilde{W}_{ab} = \lceil\sqrt{N}/(4r\sqrt{ab})\rceil$ the standard window, let $j_{ab} \in \{0, \dots, M-1\}$ be $\equiv \omega_{ab} - c_{ab} \pmod M$ and $K_{ab} := 1 + \lfloor(\widetilde{W}_{ab} - 1 - j_{ab})/M\rfloor$ if $j_{ab} < \widetilde{W}_{ab}$, else 0; the retained integers are $c_{ab} + j_{ab} + Mt$, $0 \le t < K_{ab} \le \widetilde{W}_{ab}/M + 1$, and the true $u = c_{ab} + y$ of the covering cell, having $y \equiv j_{ab}$, is among them. Writing $t = i + m'k$ with $0 \le i < m'$, the babies are β^i, $\beta := \alpha^M$, and the giants $\alpha^{aN+b-c_{ab}-j_{ab}}\beta^{-m'k}$; since $\alpha^{aN+b} \equiv \alpha^{aq+bp} \pmod p$ (from $\alpha^p \equiv \alpha$), the collision $\beta^i \equiv$ giant $\pmod p$ occurs at the true (i, k), as in Harvey's Step (2). Steps (1), (3) and (4) run with β in place of α and the labels (a, b, k) attached to the giants: Step (1) certifies $\mathrm{ord}_p(\beta) \ge m'$ or factors N, Step (3) removes the exact matches modulo N and decodes them by Lemma 3.1, and Step (4) is Algorithm 4.1 on the remaining giants; this is the adaptation of Lemma 13.5 to windows reindexed as arithmetic progressions. The order hypothesis $\mathrm{ord}_N(\beta) > m'$ follows from $\mathrm{ord}_N(\alpha) > Mm'$, as $\mathrm{ord}_N(\beta) = \mathrm{ord}_N(\alpha)/\gcd(\mathrm{ord}_N(\alpha), M)$, and Theorem 13.7 supplies it with $D = Mm'$ at cost $O(\sqrt{Mm'})(\log N)^{O(1)}$. The search costs $m' + O(P) + O(\Sigma_{\widetilde{W}}/(Mm'))$ with $P = O(r\log r)$ and $\Sigma_{\widetilde{W}} = O(N^{1/2}r^{-1/2}\log r)$ for the standard family ($\sum_{n\le r} d(n)/\sqrt{n} \sim 2\sqrt{r}\log r$);

with $r = m' = \lceil N^{1/5}M^{-2/5} \rceil$ all three terms are $N^{1/5+o(1)}M^{-2/5}$, and $\sqrt{Mm'} = N^{1/10}M^{3/10} \le N^{1/5}M^{-2/5}$ at the power-law level exactly when $M \le N^{1/7}$. In the localised case use $\mathcal{F}(I, r)$: with $P \ll_C Lr/\sqrt{N} + \sqrt{r}$ and $\Sigma_{\widetilde{W}} \ll_C L/\sqrt{r} + \sqrt{N}\log(6Cr)/r$, the choice $r = \lceil (N/(LM))^{2/5} \rceil$ and $m' \asymp A_0 M^{-2/5}$, $A_0 := L^{3/5}N^{-1/10}$, balances P against $\Sigma_{\widetilde{W}}/(Mm')$ at $A_0 M^{-2/5}$; the terms $\sqrt{r}$ and $\sqrt{N}\log(6Cr)/(rMm')$ are both $(N/(LM))^{1/5}$ up to the logarithm and are dominated when $L \ge N^{3/8}M^{1/4}$, the selection cost $\sqrt{Mm'}$ is dominated when $M \le A_0^{5/7}$, and r lies in the range $[16C^2, N^{1/3}/4]$ of Lemma 13.4 for N large. □

An implementation of the search, with Step (4) realised as a product tree and multipoint evaluation modulo N, was run on six 40-bit balanced semiprimes from a seeded generator for the five values $M = 1, 4, 16, 64, 256$, the residue $p \bmod M$ supplied, with $r = m'$ the nearest integer to $N^{1/5}M^{-2/5}$ (between 227 and 247 at $M = 1$, between 25 and 27 at $M = 256$) and the element of large order supplied by the procedure of Section 16, which factored none of the moduli itself. All thirty runs factored their modulus, 17 through an exact match modulo N decoded by Harvey's Lemma 3.1 and 13 through a collision modulo one prime. The work reported is an object count – babies plus giants plus cells, excluding the order selection and the arithmetic of the product tree, the multipoint evaluation, the sorting and the gcds – and its mean over the six moduli was $2966, 1550, 802, 409, 209$, i.e. ratios $0.52, 0.27, 0.14, 0.07$ to the unthinned run against $M^{-2/5} = 0.57, 0.33, 0.19, 0.11$. The case $M = 1$ is Harvey's algorithm itself; the measurement checks the implementation, not the asymptotics, and the moduli and per-run counts are archived with the paper.

Corollary 14.3 (the two-axis landscape)**.** *Let $p \in J_C$ be known to lie in an interval $I = [A, A + L]$, $L = N^{\lambda}$, $1/4 \le \lambda \le 1/2$, and modulo $M = N^{\mu}$, $\mu \ge 0$; put $u := 1/2 - \lambda$. Compute $\gcd(M, N)$; if it is proper, N is factored. Otherwise let s be the first integer of I congruent to p_0 modulo M, so that $p = s + Mx$ with $0 \le x \le L/M$; multiplying $s + Mx$ by $M^{-1} \bmod N$ (which exists since $\gcd(M, N) = 1$) gives the monic polynomial $x + M^{-1}s$, which has the root x modulo the unknown divisor $p \gg_C N^{1/2}$, and the Howgrave-Graham–Coppersmith theorem [17, 42] recovers it in polynomial time when $L/M = O_C(N^{1/4})$, i.e. when $\lambda - \mu \le 1/4$. Otherwise subdividing the range of x into pieces of length $N^{1/4}$ costs $N^{\lambda-\mu-1/4+o(1)}$, and in the range of Proposition 14.2 the deterministic cost is at most*

$$N^{o(1)} \min\bigl\{N^{\lambda-\mu-1/4},\ L^{3/5}N^{-1/10}M^{-2/5}\bigr\}.$$

The second term is smaller exactly when $8u + 12\mu < 1$. That region lies inside the range of Proposition 14.2: it implies $8u + 2\mu < 1$, i.e. $L \ge N^{3/8}M^{1/4}$; $\mu < 1/12 < 1/7$; and $\mu < (1-3u)/7$, i.e. $M \le (L^{3/5}N^{-1/10})^{5/7}$. In that range the second term is, up to logarithmic factors, the M1 floor of Proposition 14.1 for thinned families of density $1/M$ (the residue restriction keeps one class per window, and M is below half the shortest window in the algorithm's range). Along the residue axis ($u = 0$) the curve is $\min\{1/5 - 2\mu/5,\ 1/4 - \mu\}$, with crossover at $\mu = 1/12$. In the cell model a known trailing bit of p is thus worth a factor $2^{2/5}$ and a known leading bit $2^{3/5}$, whereas in the Coppersmith-subdivision branch either bit is worth a factor 2; polynomial-time recovery begins once $\frac{1}{4}\log_2 N + O_C(1)$ suitable contiguous bits are known, about half of the bits of a balanced factor. Which method is better depends on the shape of the information, not only on its amount.

Remark 14.4 (lattices inside cells). If a cell (a, b) at radius r covers p, Lemma 13.3 gives $|p - \sqrt{aN/b}| \ll_C \sqrt{N}/(b\sqrt{r})$, so a univariate Coppersmith call becomes available when $b\sqrt{r} \gg_C N^{1/4}$, and for smaller b the cell's interval must first be divided into $O(N^{1/4}/(b\sqrt{r}))$ Coppersmith-sized pieces. With at most $3\ell b + 1$ localised cells of denominator b, $\ell \asymp_C L/\sqrt{N}$, the exhaustive

implementation costs $\tilde{O}_C(Lr/\sqrt{N}+\sqrt{r}+L/r+N^{1/4}r^{-1/2}\log(2+N^{1/4}/\sqrt{r}))$, which for $L \geq N^{3/8}$ is minimised at the power-law level by $r \asymp N^{1/4}$, where it is $\tilde{O}_C(LN^{-1/4})$ – Coppersmith subdivision again. For the bivariate equation $(c_{ab}+x)^2-y^2 = 4abN$ of a shell cell $ab \asymp r \leq N^{1/3}$ one has $XY \asymp N/r^2$ and height $\mathsf{W} \asymp N/r$, and the degree-two case of Coppersmith's bivariate integer theorem [17], $XY \leq \mathsf{W}^{1/3}$, would require $r \gg N^{2/5}$, outside the range in which the height estimate holds and where the cell count already exceeds Lehman's. These two per-cell lattice implementations reproduce the $LN^{-1/4}$ branch rather than improving it; this is not a lower bound against every use of lattices inside or across cells, and whether any algorithm beats $\lambda - 1/4$ below $\lambda = 3/8$ is open – the known-bits methods [17, 42, 36] do not – (inside the model, Theorem 13.15 shows that typical intervals cost the larger $3\lambda/5 - 1/10$).

15. Related work, and what is not formal substitution

Coppersmith's method [17], in Howgrave-Graham's formulation [42], factors N in polynomial time once one factor is known to within an interval of length $O_C(N^{1/4})$, and with a residue $p \equiv p_0 \pmod{M}$ once $L/M = O_C(N^{1/4})$. Factoring with partial information about the factors begins with Rivest and Shamir [72], who factor N in polynomial time given two-thirds of the bits of p, and Maurer [57], who factors in polynomial time given an oracle answering $\varepsilon \log N$ arbitrary yes/no questions; Coppersmith's method reduced the contiguous prefix to half the bits; Herrmann and May [36] factor heuristically in polynomial time from a fraction $\ln 2 \approx 0.70$ of the bits of p in arbitrary positions; and the partial-key-exposure line [7, 35] recovers the private key from a fraction of its bits – a random 27% of the bits of p, q, d, d_p and d_q for a small public exponent in [35]. All of these are polynomial-time results past a threshold of known bits. Below that threshold we have found only subdivision into Coppersmith-sized pieces, of cost $N^{\lambda-\mu-1/4+o(1)}$, Hittmeir's Strassen-type search on a residue class [39, Theorem 3.1], of cost $N^{1/4}/\sqrt{M}$, and the elementary baseline, Pollard–Strassen's product-tree divisibility test over the interval [69, 78], of cost $\tilde{O}(L^{1/2}) = N^{\lambda/2+o(1)}$, which subdivision dominates throughout $\lambda \leq 1/2$ (as $\lambda/2 \geq \lambda - 1/4$) and the localised search dominates at every λ (as $\lambda/2 \geq 3\lambda/5 - 1/10$); none of the three depends on Harvey's cell structure. We have not found a published deterministic bound that depends on the interval length L between $LN^{-1/4}$ and $N^{1/5}$, nor a residue-class bound of the shape $N^{1/5}M^{-c}$; the known-bits and partial-key-exposure literature cited above gives polynomial-time results past a threshold and nothing below it. An unrefereed note of Carella [13] states bounds for moduli whose two factors both lie in prescribed congruence classes with small offsets; we have not verified its arguments, and it does not cover the statements here. Harvey's own covering lemma [31, Lemma 3.3] is local in the sense that it approximates p^2/N by a ratio a/b, but his algorithm sweeps every cell with $ab \leq r$ at a single radius $r \asymp N^{1/5}$, and its complexity is a function of N alone. The element of large order is not ours: for every bound it is supplied by Theorem 13.7 [33], after [27] and [67]; Section 16 gives the balanced-semiprime proof because its recovery step is reused.

What in the present results is more than running Harvey's search on a shorter range. Four steps. (i) The proximity lemma (Lemma 13.3): a cell of the standard family that covers p has $|a/b - p/q| < 2/(b\sqrt{r})$, so the cells that matter for an interval I are those whose ratio lies within $2/(b\sqrt{r})$ of the ratio interval $R(I)$, a set of $\ll_C Lr/\sqrt{N} + \sqrt{r}$ cells rather than $r \log r$ (Lemma 13.4); this is what makes the radius r a free parameter again, and its re-optimisation to $r \asymp (N/L)^{2/5}$ is where the exponent $3\lambda/5 - 1/10$ comes from. (ii) The list-restricted form of Harvey's correctness proof (Lemma 13.5): the cell family enters that proof only through the covering lemma, so any explicitly enumerated list containing a covering cell can replace the full family; the residue-class variant (Proposition 14.2) is the same observation for windows reindexed as arithmetic progressions, its content being that the order hypothesis passes to $\beta = \alpha^M$ and that

the balance of the three costs moves r and m' to $N^{1/5}M^{-2/5}$. (iii) On the floor side, the near/far split in the proof of Theorem 13.2 – a form is stationary inside I only if its ratio is within ℓ of $R(I)$, and the near forms number $\ll \ell 4^k + 2^k$ at denominator scale 2^k – is what raises the floor's exponent from $4\lambda/5 - 1/5$, the bound that substituting the shorter set into Lemma 3.2 gives (Proposition 13.1), to the $\min\{\lambda - 1/4, 3\lambda/5 - 1/10\}$ that the algorithms achieve. The matching of upper and lower exponents is the content of Corollary 13.9: inside the covering model for $\lambda \geq 3/8$, and for $\lambda < 3/8$ only through Coppersmith's method, which lies outside the model and is what makes the curve's first branch attainable. (iv) Below $3/8$ the Diophantine analysis of Section 13.4 – the depth of an interval, the Farey-fraction family, the localised covering sum with its near/far split at the scale of the depth, and the measure estimate for the centres of small depth – shows that the curve of least costs is not the typical in-model cost: the cost of covering an interval is determined by the position of N/p_0^2 relative to the fractions of small denominator, which no function of λ alone can express, and the typical exponent is $3\lambda/5 - 1/10$ throughout.

16. Appendix to Part 2: order-element selection for balanced semiprimes

Hittmeir's Algorithm 6.2 [39], with a bound D on the order (his δ), tries $\alpha = 2, 3, \ldots$, computes $\mathrm{ord}_N(\alpha)$ by baby-step giant-step when it is at most the bound, tests the prime divisors of the order by gcds, accumulates the least common multiple L of the orders of the candidates on which every gcd is trivial (for such α the orders modulo p and q coincide, so $L \mid p-1$ and $L \mid q-1$), skips candidates with $\alpha^L \equiv 1$, and recovers the factors from $p \equiv 1 \pmod{L}$ once $L \geq D^{1/3}$ by a Strassen-type search of cost $N^{1/4}/\sqrt{L}$. In his proof, the root bound modulo the smallest prime factor shows that at most L consecutive candidates can be skipped while the current lcm is $L < D^{1/3}$, which controls the preliminary work by $D^{1/3+o(1)}$; the hypothesis $D \geq N^{2/5}$ is used in the recovery estimate $N^{1/4}/\sqrt{L} \leq N^{1/4}D^{-1/6} \leq D^{11/24}$. For $N = pq$ with $p \in J_C$ both estimates can be replaced.

Proposition 16.1 (order-element selection for every bound; balanced semiprimes)**.** *Fix $C > 1$. There are $N_0(C)$ and a deterministic algorithm which, given $N = pq$ with $p \in J_C$, $N \geq N_0$, and $1 \leq D \leq N$, returns a factor of N or an element $\alpha \in (\mathbb{Z}/N\mathbb{Z})^\times$ with $\mathrm{ord}_N(\alpha) > D$, in $O(D^{1/2}(\log N)^{O(1)})$ bit operations. The algorithm adapts the lcm-accumulation mechanism of Hittmeir's Algorithm 6.2: it works under the balanced-semiprime hypothesis, searches directly for the exact order up to D, uses the threshold $N^{1/3}$ for the accumulated lcm, and replaces the Strassen-type recovery by the identity $N - 1 = (k+l)L + klL^2$.*

Proof. Set $L := 1$ and, for $\alpha = 2, 3, \ldots$: (a) put $g := \gcd(\alpha, N)$; if $1 < g < N$ return g, and if $g = N$ pass to the next candidate; (a′) if $\alpha^L \equiv 1 \pmod{N}$, pass to the next candidate; (b) by baby-step giant-step with $h = \lceil \sqrt{D} \rceil$ babies α^i, $0 \leq i < h$, and giants α^{jh}, $1 \leq j \leq \lceil D/h \rceil$, either certify $\mathrm{ord}_N(\alpha) > D$ and return α, or obtain a relation exponent $M = jh - i \leq h^2$ with $\alpha^M \equiv 1$; factor M by trial division and strip its prime factors by modular powering to obtain the exact order E, returning α if $E > D$; (c) for every prime $r \mid E$ compute $g_r := \gcd(\alpha^{E/r} - 1, N)$ and return g_r if it is proper; (d) set $L := \mathrm{lcm}(L, E)$, and if $L \geq N^{1/3}$ put $X := (N-1)/L$, $s := X \bmod L$, $P := (X - s)/L$, $\Delta := s^2 - 4P$, and return $1 + L(s \mp \sqrt{\Delta})/2$.

Correctness. In (c) no g_r equals N, since E is the exact order; if all are 1, then by [39, Lemma 2.3] $\mathrm{ord}_p(\alpha) = \mathrm{ord}_q(\alpha) = E$, so $E \mid p-1$ and $E \mid q-1$, and inductively $L \mid p-1$, $L \mid q-1$. Write $p = 1 + kL$, $q = 1 + lL$; then $X = (k+l) + klL$ is an integer, and $k + l \leq (p+q)/L \leq (1 + \sqrt{C})\sqrt{N}/L \leq (1 + \sqrt{C})N^{1/6} < L$ once $L \geq N^{1/3}$ and $N^{1/6} > 1 + \sqrt{C}$; so $k + l = s$, $kl = P$, and (d) returns p and q.

Termination and cost. A candidate reaching (d) has $E \nmid L$ (otherwise $\alpha^L \equiv 1$ and it was skipped), so (d) at least doubles L; as $L < N^{1/3}$ before every non-terminating execution of (d),

at most $\lceil \log_2 N^{1/3} \rceil + 1$ candidates execute (d), and at most one more reaches (b) and terminates there or in (c). Each costs $O(\sqrt{D})(\log N)^{O(1)}$: the scan, the trial division of $M \le D + 2\sqrt{D} + 1$, and $\omega(E)$ gcds. Suppose, for contradiction, that every integer $2 \le \alpha \le y$, $y := (\log N)^4$, is processed without termination, and let L_y be the value of L afterwards. Then $L_y < N^{1/3}$, because an update producing $L \ge N^{1/3}$ invokes the recovery; $L_y \mid p-1$; and every processed candidate lies modulo p in $H := \{x \in \mathbb{F}_p^\times : x^{L_y} = 1\}$ (a skipped one satisfies $\alpha^{L_e} \equiv 1$ with $L_e \mid L_y$; one executing (d) has $E \mid L_y$), a subgroup with exactly L_y elements. Every y-smooth integer in $[1, p-1]$ is a product of the processed primes $\le y$, and distinct such integers give distinct residues modulo p; hence $\Psi(p-1, y) \le L_y < N^{1/3}$. Put $u := \lfloor \log(p-1)/\log y \rfloor$; for N large $u \le \pi(y)$, and the products of u distinct primes $\le y$ are distinct y-smooth integers below p, so $\Psi(p-1, y) \ge \binom{\pi(y)}{u} \ge (\pi(y)/u)^u$. Since $\pi(y) > y/\log y$ for $y \ge 17$ and $u \log y \le \log(p-1)$, $\pi(y)/u > (\log N)^4/\log(p-1) \ge (\log N)^3$, while $u \ge \log(p-1)/(4 \log\log N) - 1$; therefore

$$\Psi(p-1, y) \ge (\log N)^{3u} \ge \frac{(p-1)^{3/4}}{(\log N)^3} \gg_C \frac{N^{3/8}}{(\log N)^3} > N^{1/3}$$

for $N \ge N_0(C)$, a contradiction (alternatively, $\Psi(x, (\log x)^4) = x^{3/4 + o(1)}$ [80, Ch. III.5]). So the procedure terminates within $(\log N)^4$ candidates, the skipped ones costing a gcd and one modular exponentiation each. □

The constant C enters only through $N_0(C)$. The statement is a special case of Theorem 13.7, which holds for every $N \ge 3$ with the sharper cost $O(D^{1/2} \log \bar{D} \log N/(\log\log \bar{D})^{1/2})$; for $D \ge N^{2/5}$ it was already Hittmeir's Theorem 6.3. The observation that drives the termination argument – every y-smooth integer below p lies in the subgroup $\{x^{L_y} = 1\}$ of $\mathbb{F}_p^\times$, so $\Psi(p-1, y) \le L_y$ – is also the key observation of [33], where the lower bound for Ψ is Konyagin–Pomerance's; what differs is the second stage: theirs is a trial division along the progression $kM + 1$ over a range of length $(\log D)^{O(1)} M$, valid for every N, whereas the proof above uses the balanced hypothesis to recover p and q from $N - 1 = (k+l)L + klL^2$ as soon as $L \ge N^{1/3}$. A Burgess-type argument – the failing candidates lie in a proper subgroup of $\mathbb{F}_q^\times$ – gives the weaker bound $N^{1/(8\sqrt{\mathrm{e}}) + \varepsilon} D^{1/2}$ for the same loop without the recovery step; the counting argument is both simpler and stronger.

Computational verification. On 60 random 36- and 44-bit moduli with $D = \lfloor N^{1/6} \rfloor$ the first candidate $\alpha = 2$ is accepted; on 22 moduli $N = pq$ built from two prime factors of $2^d - 1$ with $\mathrm{ord}_p(2) = \mathrm{ord}_q(2) = d$ the candidate 2 fails, as it must, and 3 is accepted or factors N; and on 40 constructed moduli $p = 1 + 4L_0$, $q = 1 + 12L_0$ with $L_0 \equiv 1 \pmod 3$ prime and 2 a cubic residue modulo q – so that $p, q \equiv 5 \pmod 8$ make 2 a non-residue modulo both and $\mathrm{ord}_p(2) = \mathrm{ord}_q(2) = 4L_0$ exactly, checked directly – run with $D = N \ge 4L_0$, the candidate 2 fails and, as $4L_0 \ge N^{1/3}$, the recovery step returns p and q. The moduli and traces of these runs are archived with the paper.

For an explicitly represented baby set the relevant condition is pairwise distinctness of the babies modulo each prime, and it too can be supplied at a cost subordinate to the search.

Lemma 16.2 (order selection for an explicit baby set)**.** *Let $N = pq$ with $p \in J_C$, $N \ge N_0(C)$, and let $B \subset \mathbb{Z}$ be an explicitly given set of integers of $O(\log N)$ bits, nonempty, with spread $S := \max B - \min B \le N$. A deterministic algorithm returns either a proper factor of N or a unit α whose residues α^b, $b \in B$, are pairwise distinct modulo p and modulo q, in $O((|B| + \sqrt{S})(\log N)^{O(1)})$ bit operations.*

Proof. If $|B| = 1$ the condition is vacuous and the first unit candidate is returned; so let $|B| \ge 2$, whence $S \ge 1$, and replace B by $B - \min B \subseteq [0, S]$, which multiplies every baby by the same

unit and preserves distinctness. Set $L := 1$ and for $\alpha = 2, 3, \ldots$: (a) as in Proposition 16.1, return a proper $\gcd(\alpha, N)$ and skip α with $\alpha^L \equiv 1$; (b) with $F_B(Y) := \prod_{b\in B}(Y - \alpha^b)$ compute $g_B := \gcd(\prod_b F_B'(\alpha^b), N)$. Since $\prod_b F_B'(\alpha^b) = \pm\prod_{\{b,b'\}}(\alpha^b - \alpha^{b'})^2$, $g_B = 1$ exactly when the babies are distinct modulo both primes (return α), g_B is proper when they coincide modulo exactly one (return it), and $g_B = N$ when they coincide modulo each, possibly at different pairs; the product is computed division-free in $\tilde{O}(|B|)$ ring operations by a product tree for F_B, formal differentiation and monic multipoint evaluation of F_B' at the α^b. (c) If $g_B = N$, both $\mathrm{ord}_p(\alpha)$ and $\mathrm{ord}_q(\alpha)$ divide differences of elements of B, hence are at most S. With $H := \lceil\sqrt{S}\rceil$ form the babies α^i, $0 \le i < H$, and the giants α^{jH}, $1 \le j \le \lceil S/H\rceil$; the differences $jH - i$ cover $[1, S]$. Sort both lists. An exact equality $\alpha^i = \alpha^{jH}$ in $\mathbb{Z}/N\mathbb{Z}$ gives a relation exponent $jH - i \le S + H$, from which the exact order $E = \mathrm{ord}_N(\alpha)$ follows by trial division and modular powering. For every prime $\ell \mid E$ compute $g_\ell := \gcd(\alpha^{E/\ell} - 1, N)$; since E is the exact order no g_ℓ equals N, so any $g_\ell > 1$ is a proper factor, and if all $g_\ell = 1$ then [39, Lemma 2.3] gives $\mathrm{ord}_p(\alpha) = \mathrm{ord}_q(\alpha) = E$; in the latter case set $L := \mathrm{lcm}(L, E)$ and, if $L \ge N^{1/3}$, recover p and q from $N - 1 = (k + l)L + klL^2$ as in Proposition 16.1. If there is no exact equality, then $\mathrm{ord}_N(\alpha) > S \ge H$, since an order $d \le S$ would be some $jH - i$; as both component orders are at most S they differ, so $e := \min(\mathrm{ord}_p \alpha, \mathrm{ord}_q \alpha) \le S$ has $\gcd(\alpha^e - 1, N)$ proper, i.e. the pair with $jH - i = e$ collides modulo exactly one prime, and Harvey's Algorithm 4.1, whose preconditions ($\mathrm{ord}_N(\alpha) > H$, no exact equalities) are met, finds a factor in $\tilde{O}(\sqrt{S})$ operations. Termination and cost are as in Proposition 16.1: a candidate reaching the lcm update has $E \nmid L$, so $O(\log N)$ candidates do; every processed candidate lies modulo p in $\{x : x^{L_y} = 1\}$ with $L_y < N^{1/3}$, so the smooth-number count bounds the candidates by $(\log N)^4$; skipped candidates cost a gcd and an exponentiation, the others $\tilde{O}(|B| + \sqrt{S})$. □

Part 3. The additive structure of the Lehman–Harvey start set

Summary. Harvey's deterministic factoring algorithm tests, for each cell (a, b) with $ab \le r$, a window of integers beginning at the *start* $\lceil 2\sqrt{abN}\rceil - aN - b$, and Part 1 shows that, within a covering model containing that algorithm, a lower-bound certificate for representing the tested set as a difference $B - G$ of two explicit sets is governed by one statistic of the starts: the largest number $D_{\max}$ of ordered pairs of windows whose start difference lies within a window length of a common value. This part studies that statistic. For the $a = 1$ shell the starts form a linear point system in $u = 2\sqrt{N}$; its exact repeated speeds are classified and its transient coincidence mechanisms are analysed. In the transfer range $r \le N^{1/5-\varepsilon}$ the statistic is modulus-free and lies between $c\,r^{1/3}$ and $2.1\,r^{2/3}$; in the planar regime $N^{1/5} \le r \le N^{1/3}$ offset resonance produces clusters of size $\asymp N^{1/12}$ at $r = N^{1/3}$, where $6N^{2/9}$ bounds the whole shell. For the balanced restriction the half-offset families force the statistic up to order $rN^{-1/4}$ in every family retaining the constructed cells – unconditionally for the primitive family – which rules out a bounded statistic there and leaves an envelope conjecture with conditional exponents 5/28 and 2/11. Exact-start censuses of the resonant families to 96 bits, and exhaustive sweeps at the sizes marked exact, are reported, and the open envelope problem – the exponent of $D_{\max}$ between the proved 2/3 and the conjectured 1/4 – is stated precisely. Conditional and heuristic statements are marked, and none is a lower bound on the cost of factoring.

17. Introduction

Let $N = pq$ be a balanced semiprime and let $r \ge 1$. Lehman's method [49] and Harvey's exponent-one-fifth algorithm [31, 32] test, for cells (a, b) of positive integers with $ab \le r$, whether an integer close to $2\sqrt{abN}$ equals $aq + bp$; Harvey organises the test as a babystep–giantstep collision in the exponent, so that each cell contributes a window of consecutive integers beginning

at the start $\lceil 2\sqrt{abN}\rceil - aN - b$. Part 1 formalises the class of such algorithms as a covering model, proves the floors $N^{1/5}$ (materialised giants) and $N^{1/6}$ (free giants) inside it, and reduces the question whether the second floor is attainable, within the explicit-difference representation model and subject to the order and decoding conditions of Proposition 4.1, to a combinatorial one: can the set of tested integers be written as a difference $B-G$ of two explicit sets of total size $N^{1/6+o(1)}$? Its Lemma 4.3 answers that a small representation forces approximate coincidences among the starts, measured by the statistic $D_{\max}$, and its Theorem 4.4 converts an envelope $D_{\max} \le r^{\eta}$ into the floor $N^{(1-\eta)/(5-4\eta)}$. The rest of this part is about that statistic – what the starts' coincidences are, how large the clusters provably are, how large they can provably be, and what remains open – and, in one section, about the covers that a direct search finds.

The starts of the $a = 1$ shell are the rounded chirp $\lceil 2\sqrt{kN}\rceil$, $k \in (r/2, r]$, and their pairwise differences move linearly in $u = 2\sqrt{N}$ with speeds $\sqrt{k} - \sqrt{k'}$. The additive structure is therefore that of the square roots of the integers of a dyadic range, seen at a resolution far below their mean spacing, together with the integer offsets $k - k'$ that the modulus introduces. Section 18 fixes the notation of this part. Section 19 examines the balanced restriction $a \le b \le Ca$: the rays, the equal-product pairs, the excision to one cell per product, the half-offset families that survive it, and the conditional exponents $5/28$ and $2/11$. Section 20 treats the $a = 1$ shell in the transfer range $r \le N^{1/5-\varepsilon}$, where the statistic is modulus-free: the classification of exact repeated speeds, Beatty chains, near-chains, two-progression families, the symmetric family with its sharp constant, and the envelope $7\,r^{1/3}$ for every drift-free family. Section 21 treats the planar regime $N^{1/5} \le r \le N^{1/3}$: the cap, offset resonance and the theorems that make it rigorous, the planar two-progression envelope, the unconditional upper bounds $6N^{2/9}$ and $2.1\,r^{2/3}$, and the additive energy. Section 22 states the conjectures and the refined-prime programme, and Section 23 reports the computations. Section 24 lists what is proved, what is conjectured and what is open.

This part depends on Part 1 through the covering-sum lemma (Lemma 3.2), the two floors (Theorems 3.4 and 3.6), the explicit-difference representation model with Lemma 4.3 and Theorem 4.4; the standard results of analytic number theory that are used are cited where they occur. Throughout, $\omega(n)$ is the number of distinct prime factors of n, $d(n)$ the number of divisors, $\pi(x)$ the number of primes up to x, $\|x\|$ the distance from x to the nearest integer, v_p the p-adic valuation, log the natural logarithm, and "bits" the bit length $\lfloor \log_2 N \rfloor + 1$ of a modulus.

18. Setting and notation

The model, the standard family with its windows $W_{ab} := \lceil \sqrt{N}/(4r\sqrt{ab})\rceil$, the shifted tested exponents $n - aN - b$ and the starts $\lceil 2\sqrt{abN}\rceil - aN - b$ are those of Definition 2.1 and Section 4.1. Harvey's algorithm materialises babies α^i, $0 \le i < m$, and giants $\alpha^{aN+b-\lceil 2\sqrt{abN}\rceil - jm}$, and a collision between a baby and a giant modulo p detects the shifted tested exponent $z = \text{start} + jm + i$ as the difference $i - \gamma$ of the element i of $B := [0, m) \cap \mathbb{Z}$ and the element $\gamma := -(\text{start} + jm)$ of $G := \{-(\text{start} + jm)\}$; we call this pair (B, G) *Harvey's block representation*. In the explicit-difference representation model an algorithm materialises two explicit sets $B, G \subset \mathbb{Z}$ with the set Z of shifted tested exponents contained in $B - G$ at cost $|B| + |G|$; since $Z \subseteq B - G$ gives $|B|\,|G| \ge |Z|$, a representation of a set of pairwise distinct shifted tested exponents costs at least $2\sqrt{\Sigma_W} \ge 2(c_C/2)^{1/3} N^{1/6}$ (Theorem 3.6), and when the exponents are not distinct Σ_W is to be replaced by the number of distinct ones, which is how the hypothesis appears in Proposition 19.5. The standard family has pairwise disjoint windows for $r \le N^{1/3}$ and N large, so its tested exponents are distinct. Harvey's algorithm attains the first floor; whether any representation attains the second is the question that the statistic $D_{\max}$ of Lemma 4.3 addresses. The statements about $D_{\max}$ below are combinatorial and need no covering hypothesis.

For a family $\mathcal{F}$ of windows $I = [s_I, s_I+W_I)\cap\mathbb{Z}$ we write, as in Lemma 4.3, $D_{\max}(\mathcal{F}, W)$ for the approximate-Sidon statistic; for a family of cells, $D_{\max}(\cdot, W)$ denotes the statistic of their starts, each start regarded as the first element of a window of the stated width; with $W = 1$ it is the statistic of the singleton starts, the largest number of ordered pairs of distinct cells with the same nonzero start difference. We also write $D(h) := \#\{(I, I') \in \mathcal{F}^2 : I \neq I',\ |s_I - s_{I'} - h| < W\}$ for the count at a given difference h, so that $D_{\max} = \max_{h\neq 0} D(h)$. Harvey's block representation (B an interval of length m, G the negated window starts minus multiples of m) satisfies Lemma 4.3 with room; the bound bites when a representation uses fewer giants than windows, which forces approximate coincidences among the starts.

For the cells $(1, k)$ write $s_k = \lceil 2\sqrt{kN}\rceil - N - k$ for the first tested exponent, $\mathcal{F}_1(r) = \{(1, k) : r/2 < k \le r\}$ for the *full* $a = 1$ *shell*, $\mathcal{F}_1^{\square}(r)$ for its squarefree members, $\mathcal{F}_1^{P}(r)$ for its prime members, and, for the standard family at parameter R, $W_R(N; \rho) := \max_{\rho/2<k\le\rho}\lceil\sqrt{N}/(4R\sqrt{k})\rceil$ for Lemma 4.3's window on the shell at ρ; we abbreviate $W(N, \rho) := W_\rho(N; \rho) = \lceil W_{\mathrm{real}}(\rho)(1 + O(1/\rho))\rceil$, with $W_{\mathrm{real}}(\rho) := \sqrt{N}/(2\sqrt{2}\,\rho^{3/2})$ as in Theorem 4.4, and write W for $W(N, r)$ and W_{real} for $W_{\mathrm{real}}(r)$ when r is fixed.

19. The difference-cover statistic and the balanced restriction

The balanced-cell restriction for J_C. Theorem 4.4 concerns the standard family, whose $a = 1$ shell covers the unbalanced range $p \asymp \sqrt{N/k}$. A search required to cover only J_C need not retain that shell. We therefore examine the natural balanced-cell restriction

$$\mathcal{F}_{\mathrm{bal}}(r) := \{(a, b) : a \le b \le Ca,\ ab \le r\}.$$

Lemma 4.3 applies to every disjoint sub-family of its candidate set, but its whole-family statistic contains a ray contribution absent from the $a = 1$ shell.

Lemma 19.1 (rays)**.** *Assume $N \ge 4$. Let $\mathcal{F}_{\mathrm{bal}}(r)$ be the balanced family at parameter r. Then, for $W \ge 2$,*

$$D_{\max}(\mathcal{F}_{\mathrm{bal}}(r), W) \ge \lfloor\sqrt{r}\rfloor - 1, \qquad \textit{and} \qquad D_{\max}(\mathcal{F}_{\mathrm{bal}}(r), 1) \ge \frac{\lfloor\sqrt{r}\rfloor - 1}{2}.$$

Proof. Put $M := \lfloor\sqrt{r}\rfloor$. The diagonal cells (m, m), $1 \le m \le M$, are balanced and belong to the family because $m^2 \le r$. With the sign convention of Theorem 4.4, their starts are

$$s_m = \lceil 2m\sqrt{N}\rceil - m(N + 1) = \lceil m\theta\rceil - m(N + 1), \qquad \theta := 2\sqrt{N};$$

equivalently, after negating all starts, $\widetilde{s}_m := m(N+1)-\lceil m\theta\rceil = \lfloor m(\sqrt{N}-1)^2\rfloor$, and simultaneous negation does not change $D_{\max}$, since it sends the parameter t to $-t$. For every real x, write $x = n - f$ with $n = \lceil x\rceil$ and $0 \le f < 1$; then $\lceil x + \theta\rceil - \lceil x\rceil = \lceil\theta - f\rceil \in \{\lfloor\theta\rfloor, \lceil\theta\rceil\}$. Thus $s_{m+1} - s_m \in \{\lfloor\theta\rfloor - (N + 1), \lceil\theta\rceil - (N + 1)\}$, two integers differing by at most 1. Put $t := \lceil\theta\rceil - (N + 1)$; since $N \ge 4$, both possible increments are nonzero. For each $1 \le m < M$ the ordered pair $(I, I') = ((m + 1, m + 1), (m, m))$ satisfies $|s_I - s_{I'} - t| \le 1 < W$; these are $M - 1$ distinct ordered pairs of distinct windows, proving the first assertion. For $W = 1$ the condition $|s_I - s_{I'} - t| < 1$ is exact equality; one of the at most two increment values occurs at least $(M - 1)/2$ times, and that nonzero value may be used as t. □

More generally every ray $m(a_0, b_0)$ is a Beatty sequence of starts $\lceil m\cdot 2\sqrt{a_0b_0N}\rceil - m(a_0N+b_0)$ with two-valued increments, and the diagonal is the longest. The balanced family therefore has $D_{\max} \gg \sqrt{r}$, the value of a generic difference set $B - G$ with $|B| = |G| \asymp \sqrt{r}$, so the small-$D_{\max}$ condition that makes Lemma 4.3 stronger than the counting floor is unavailable for the whole balanced family: the Sidon branch of the certificate furnished by Lemma 4.3 for it is at most $2(n'^2/(4W\sqrt{r}))^{1/3} \asymp N^{1/6}/W^{1/3}$ at $r = N^{1/3}$, no more than the free-giant floor M2 and, with

the family's own maximal window $W \asymp N^{1/6}$ (the cells with $ab = O(1)$), less than the trivial bound $2\sqrt{n'}$.

Removing the rays is not enough: the twist $\lceil 2\sqrt{abN} \rceil$ depends on the product ab alone, so two divisor pairs of the same product have starts differing by an integer with no rounding at all.

Lemma 19.2 (equal products)**.** *Let $h < k$ be coprime integers with $k < Ch$, let $m \geq 1$, and put $\lambda_0 := \max\{\lceil hm/(k-h) \rceil, \lceil km/(Ch-k) \rceil\}$. For every integer λ with $\lambda_0 \leq \lambda \leq \sqrt{r/(hk)} - m$ the cells $(h\lambda, k(\lambda+m))$ and $(h(\lambda+m), k\lambda)$ belong to $\mathcal{F}_{\mathrm{bal}}(r)$, have the same product $hk\lambda(\lambda+m)$, and their starts differ by exactly $m(hN-k)$. For $(k,h,m) = (3,2,6)$ and $C = 2$ (the displayed bound remains valid for every $C \geq 2$) both cells are primitive exactly when λ is odd and $3 \nmid \lambda$; hence*

$$D_{\max}(\mathcal{F}_{\mathrm{bal}}^{\mathrm{prim}}(r), 1) \;\geq\; \tfrac{1}{3}\big(\sqrt{r/6} - 24\big) - 1,$$

where $\mathcal{F}_{\mathrm{bal}}^{\mathrm{prim}}(r)$ denotes the primitive ($\gcd(a,b) = 1$) balanced cells.

Proof. For the first cell, $a \leq b$ holds since $h < k$, and $b \leq Ca$ is $k(\lambda+m) \leq Ch\lambda$, i.e. $\lambda \geq km/(Ch-k)$; for the second, $a \leq b$ is $h(\lambda+m) \leq k\lambda$, i.e. $\lambda \geq hm/(k-h)$, and $b \leq Ca$ is $k\lambda \leq Ch(\lambda+m)$, true since $k < Ch$. Both products equal $hk\lambda(\lambda+m) \leq hk(\lambda+m)^2 \leq r$. The twists agree, and with the cells ordered as above,

$$s_{h\lambda,k(\lambda+m)} - s_{h(\lambda+m),k\lambda} = \big(h(\lambda+m) - h\lambda\big)N + \big(k\lambda - k(\lambda+m)\big) = m(hN-k),$$

so all these ordered pairs contribute to the same signed difference $t = m(hN-k)$ in Lemma 4.3. For $(3,2,6)$: a prime dividing 2λ and $3(\lambda+6)$ is 2 (forcing λ even), 3 (forcing $3 \mid \lambda$) or a prime $\ell \geq 5$ dividing λ and $\lambda+6$, hence 6, impossible; the same for $2(\lambda+6)$ and 3λ. So both cells are primitive iff λ is odd and $3 \nmid \lambda$, a set of density $1/3$, and $\lambda_0 = 18$ for $C = 2$. The λ in $[18, \sqrt{r/6} - 6]$ with λ odd and $3 \nmid \lambda$ number at least $(\sqrt{r/6} - 24)/3 - 1$, and each gives an ordered pair of primitive cells with the same start difference. □

The census of Section 23.6 confirms that this construction is the maximiser of the primitive family from 44 bits on: the maximal exact cluster has $t = 12N - 18$ and sizes $13, 15$; $19, 19$; $25, 25$ at $44, 46, 48$ bits. For the corresponding radii $22524, 25327$; $38227, 36007$; $56099, 54374$ the displayed lower bound gives the integers $12, 13$; $18, 17$; $24, 23$; counting the admissible λ exactly in the lemma's sufficient interval gives $13, 13$; $19, 18$; $24, 24$, and using the exact product condition $6\lambda(\lambda+6) \leq r$ gives $13, 15$; $19, 19$; $25, 25$, exactly the measured cluster sizes (at ≤ 42 bits the family is still short and a different structure, described below, is maximal). Thus the primitive balanced family also sits at $D_{\max} \asymp \sqrt{r}$, and Lemma 4.3 applied to it gives nothing above $N^{1/6}$.

Both structures are excised by keeping *one cell per product*: let $\mathcal{F}_{\mathrm{uniq}}(r)$ consist of, for each $n \leq r$ with a primitive balanced divisor pair, the pair with the smallest a. Then $\#\mathcal{F}_{\mathrm{uniq}}(r) \geq r^{1-o(1)}$: indeed $\#\mathcal{F}_{\mathrm{bal}}^{\mathrm{prim}}(r) = (3\log C/\pi^2 + o(1))\, r$, while a product $n \leq r$ has at most $2^{\omega(n)} = r^{o(1)}$ ordered coprime divisor pairs (the census finds between $0.196\, r$ and $0.210\, r$, the former at the largest sample). A cover of Z covers the singleton starts of $\mathcal{F}_{\mathrm{uniq}}(r)$, so Lemma 4.3 applies to them with $W = 1$. The census gives, for two moduli at each even size from 30 through 46 bits and one modulus at 48 bits (r from 958 to 56 099),

$$D_{\max}(\mathcal{F}_{\mathrm{uniq}}(r), 1) = 3, 4;\; 4, 5;\; 5, 6;\; 6, 6;\; 5, 5;\; 7, 6;\; 7, 6;\; 7, 7;\; 9, 6;\; 9,$$

while $\sqrt{r}$ grows from 31 to 237 and $r^{1/3}$ from 9.9 to 38. At every sampled size a maximising cluster is realised by pairs of consecutive cells $(a,b), (a,b-1)$ on the line $2b = 3a + 1$ (the unique maximiser at thirteen of the nineteen moduli, one of several at the other six); with $c_{a,b} := \lceil 2\sqrt{abN} \rceil$ its signed difference is exactly $t = s_{a,b} - s_{a,b-1} = c_{a,b} - c_{a,b-1} - 1$, of order

$\sqrt{N}$. The mechanism is explicit: on that line $b+(b-1)=3a$ and $b(b-1)=(9a^2-1)/4$, so the increment of the twist is

$$2\sqrt{N}\big(\sqrt{ab}-\sqrt{a(b-1)}\big)=\frac{2\sqrt{N}\sqrt{a}}{\sqrt{b}+\sqrt{b-1}}=\sqrt{\tfrac{2N}{3}}\Big(1+\frac{1}{72a^2}+O(a^{-4})\Big),$$

so that $t=\sqrt{2N/3}\,(1+1/(72a^2)+O(a^{-4}))-1+O(1)$, stationary to second order along the family. The a^{-2} correction has derivative of order $\sqrt{N}/a^3$, so the natural scale on which one rounded increment can remain coherent is $\Delta a \asymp a^3/\sqrt{N}$; at $a\asymp\sqrt{r}$ and $r=N^{1/3}$ this scale is bounded, and in the largest samples of the census the maximising clusters contain about nine cells. Writing $h=2h'$, the general half-offset line is $hb=ka+h/2$, with increment $\sqrt{Nh/k}\,(1+h^2/(32k^2a^2)+O(a^{-4}))$; the derivative of the correction gives a unit-drift interval $\Delta a\asymp k^{5/2}a^3/(h^{5/2}\sqrt{N})$, which at $a\asymp\sqrt{rh/k}$, $r=N^{1/3}$, and after the congruence spacing h of the admissible a, is of order k/h^2 cells, largest at $(k,h)=(3,2)$ for $C=2$. This identifies the local coherence scale of one mechanism; it does not classify all exact coincidences or exclude pile-ups from several families. The same calculation shows that the statistic of the primitive balanced family is not bounded, and that of the unique-product family is not bounded whenever the constructed cells are its representatives. The coherence scale $\Delta a\asymp a^3/\sqrt{N}$ was computed for the shortest such families; a family whose line direction (k,d) grows with N trades length for coherence, and the trade-off is won at a direction of size $N^{1/12}$.

Lemma 19.3 (half-offset families)**.** *Let $m\equiv 1 \pmod 4$, $m\ge 5$, and put $k:=(m+1)/2$, $d:=m-1$, $\alpha:=(m-1)/4$, $\beta:=(m-3)/2$, so that $k\beta-d\alpha=-1$ and $k(\beta+2)-d(\alpha+1)=1$. For $\lambda\ge 1$ the cells $P_\lambda:=(k\lambda+\alpha,d\lambda+\beta)$ and $P'_\lambda:=(k\lambda+\alpha+1,d\lambda+\beta+2)$ are primitive, satisfy $a\le b\le 2a$, and have products*

$$n_\lambda=kd\,(\lambda+s)^2-\frac{1}{4kd},\qquad n'_\lambda=kd\,(\lambda+s')^2-\frac{1}{4kd},\qquad s:=\frac{k\beta+d\alpha}{2kd}\in\big(0,\tfrac12\big),\quad s':=s+\frac{m}{kd},$$

pairwise distinct over λ and interlaced ($n_\lambda<n'_\lambda<n_{\lambda+1}$). Their twist increment is

$$2\sqrt{N}\big(\sqrt{n'_\lambda}-\sqrt{n_\lambda}\big)=\frac{2\sqrt{N}\,m}{\sqrt{kd}}\Big(1+\frac{1}{8(kd)^2(\lambda+s)(\lambda+s')}+\theta_\lambda\Big),\qquad |\theta_\lambda|\le\frac{1}{40\,(kd)^4\lambda^4}.$$

Consequently, for N sufficiently large and $47N^{1/4}\le r\le N^{1/3}$, every family of cells containing the P_λ,P'_λ with $\lambda_{\max}/2\le\lambda\le\lambda_{\max}$, where m is the least positive integer congruent to $1 \pmod 4$ with $m^2\ge 11\sqrt{N}/r$ and $\lambda_{\max}$ the largest λ with $n'_\lambda\le r$, has $D_{\max}(\cdot,1)\ge rN^{-1/4}/15-1$. In particular, when the balance constant satisfies $C\ge 2$, this holds for $\mathcal{F}^{\mathrm{prim}}_{\mathrm{bal}}(r)$, and it holds for $\mathcal{F}_{\mathrm{uniq}}(r)$ provided that both cells in each counted pair are selected as the unique-product representatives (a census of the half-offset directions at 42 and 48 bits finds this for 87–100% of the pairs of the directions it tabulates, and for 97–100% when $m\ge 9$; it is not proved). At $r=\lfloor N^{1/3}\rfloor$ it gives $D_{\max}\ge N^{1/12}/15-1-O(N^{-1/4})$, hence $D_{\max}\ge N^{1/12}/15-2$ for all sufficiently large N.

Proof. Put $K:=kd=(m^2-1)/2$. Since $m\equiv 1 \pmod 4$, $k=(m+1)/2$ is odd, and every common divisor of k and $d=m-1$ divides $2k-d=2$; thus $\gcd(k,d)=1$. Direct calculation gives $k\beta-d\alpha=[(m+1)(m-3)-(m-1)^2]/4=-1$ and $k(\beta+2)-d(\alpha+1)=-1+2k-d=1$. A common divisor of the coordinates of P_λ therefore divides $k(d\lambda+\beta)-d(k\lambda+\alpha)=-1$, and the corresponding determinant for P'_λ equals 1; hence both cells are primitive. For $P_\lambda=(a,b)$, $b-a=(d-k)\lambda+(\beta-\alpha)=\frac{m-3}{2}\lambda+\frac{m-5}{4}\ge 0$, and for P'_λ this difference is larger by 1; moreover $2a-b=(2k-d)\lambda+(2\alpha-\beta)=2\lambda+1>0$, and adding $(1,2)$ leaves $2a-b$ unchanged. Thus both cells satisfy $a\le b\le 2a$.

Now $k\beta + d\alpha = (m^2 - 2m - 1)/2$, so $s = (m^2 - 2m - 1)/(2(m^2 - 1)) \in (0, \frac{1}{2})$, and completing the square, $n_\lambda = K(\lambda + s)^2 - (k\beta - d\alpha)^2/(4K) = K(\lambda + s)^2 - 1/(4K)$. For the shifted cell $s' - s = (2k + d)/(2K) = m/K = 2m/(m^2 - 1)$, so $s' = (m^2 + 2m - 1)/(2(m^2 - 1))$, $0 < s < \frac{1}{2} < s' < 1$, and the determinant $k(\beta + 2) - d(\alpha + 1) = 1$ gives $n'_\lambda = K(\lambda + s')^2 - 1/(4K)$. Since $0 < s' - s < 1$, $n_\lambda < n'_\lambda < n_{\lambda+1}$, so all the displayed products are distinct.

For the twist increment put $x := \lambda + s$, $y := \lambda + s'$, $\delta := y - x = m/K$ and $A := 1/(4K^2)$, so that $\sqrt{n_\lambda} = \sqrt{K}\sqrt{x^2 - A}$. With $Q(z) := \sqrt{z^2 - A} - z + A/(2z)$ one has the identity

$$\sqrt{y^2 - A} - \sqrt{x^2 - A} = \delta + \frac{A\delta}{2xy} + Q(y) - Q(x),$$

hence $\sqrt{n'_\lambda} - \sqrt{n_\lambda} = (m/\sqrt{K})\big(1 + 1/(8K^2xy) + \theta_\lambda\big)$ with $\theta_\lambda := (Q(y) - Q(x))/\delta$. For $u = A/z^2$, $Q'(z) = (1 - u)^{-1/2} - 1 - u/2$; here $K \geq 12$ and $z \geq \lambda \geq 1$, so $u \leq 1/576$, and the nonnegative binomial series gives $0 \leq Q'(z) \leq (\frac{3}{8} + u/(1 - u))u^2 < \frac{2}{5}u^2$. The mean-value theorem yields $0 \leq \theta_\lambda \leq \frac{2}{5}A^2/\lambda^4 = 1/(40K^4\lambda^4)$, and multiplication by $2\sqrt{N}$ proves the displayed formula.

For the count, note that the least positive $m \equiv 1 \pmod 4$ with $m^2 \geq 11\sqrt{N}/r$ satisfies $m \leq (11\sqrt{N}/r)^{1/2} + 4$. Since $r \geq 47N^{1/4}$,

$$\frac{20m}{\sqrt{2r}} \leq \frac{20\sqrt{11/2}}{rN^{-1/4}} + \frac{80}{\sqrt{2r}} \leq \frac{20\sqrt{11/2}}{47} + o(1) < 1,$$

so $m \leq \sqrt{2r}/20$ for all sufficiently large N. Because $K \leq m^2/2$ and $s' < 1$, the inequality $\lambda + s' \leq \sqrt{2r}/m$ implies $n'_\lambda \leq r$; hence, with $A_0 := \sqrt{2r}/m \geq 20$, $\lambda_{\max} \geq \lfloor A_0 - s' \rfloor > A_0 - 2 \geq 18$, so $\lambda_{\max} \geq 19$ and $\lambda_{\max}^2 > (A_0 - 2)^2 \geq 0.8A_0^2 = 1.6\, r/m^2$.

Over the integer range $\lambda_{\max}/2 \leq \lambda \leq \lambda_{\max}$ the correction $1/((\lambda + s)(\lambda + s'))$ varies by at most $3/\lambda_{\max}^2$: after scaling by $\lambda_{\max}$, the difference $1/((\frac{1}{2} + a)(\frac{1}{2} + b)) - 1/((1 + a)(1 + b))$ is decreasing in each of $a, b \geq 0$ and is at most 3. The θ_λ terms contribute $o(K^{-2}\lambda_{\max}^{-2})$, uniformly in the stated range. Thus the twist increment varies by at most

$$\begin{aligned}\frac{2\sqrt{N}\, m}{\sqrt{K}} \cdot \frac{3}{8K^2\lambda_{\max}^2}\,(1 + o(1)) \leq \frac{15}{32}\,\frac{\sqrt{N}\, m^3}{rK^{5/2}}\,(1 + o(1)) &= \Big(\frac{15}{4\sqrt{2}} + o(1)\Big)\frac{\sqrt{N}}{rm^2} \\ &< 2.93\,\frac{\sqrt{N}}{rm^2} \leq \frac{2.93}{11} < 0.27,\end{aligned}$$

using $K = (m^2 - 1)/2$, the fact that the selected m tends to infinity with N, and $m^2 \geq 11\sqrt{N}/r$.

For real x, y the integer $\lceil y \rceil - \lceil x \rceil$ is the floor or the ceiling of $y - x$. Since the real increments above lie in an interval of length less than 1, the rounded increments $\lceil 2\sqrt{n'_\lambda N} \rceil - \lceil 2\sqrt{n_\lambda N} \rceil$ take at most three values over the range, and so do the start differences $s_{P'_\lambda} - s_{P_\lambda} = (\lceil 2\sqrt{n'_\lambda N} \rceil - \lceil 2\sqrt{n_\lambda N} \rceil) - N - 2$; these are nonzero for N large, the rounded increment being $O(\sqrt{N})$. The range contains at least $\lambda_{\max}/2$ integers, so one nonzero start difference occurs at least $\frac{1}{3}(\sqrt{2r}/(2m) - 1)$ times. Finally $m \leq (11\sqrt{N}/r)^{1/2} + 4$ and $r \leq N^{1/3}$ give $\sqrt{2r}/(2m) \geq rN^{-1/4}(1 - o(1))/\sqrt{22}$, and $1/(3\sqrt{22}) > 1/15$, whence the bound. □

At $r = \lfloor N^{1/3} \rfloor$ the direction selected in Lemma 19.3 has $m \asymp N^{1/12}$ and $\lambda_{\max} \asymp N^{1/12}$, and the resulting family alone has a cluster of size $\asymp N^{1/12}$. By contrast the fixed direction $(2, 3)$ with shift $(0, -1)$ has unit-scale drift at $r = N^{1/3}$ and contributes only $O(1)$ cells to one rounded increment. The numerical scale $N^{1/12} = 30$ corresponds to moduli of about 59 bits, beyond the 48-bit census; this is a scale comparison, not a proved crossover for the full statistic. At 48 bits the directions $m = 13, 17, 21, 25$ have $25, 19, 15, 12$ cells and maximal clusters $4, 6, 6, 6$; these are illustrative directions rather than the m selected by the lemma,

which is about 53 at that size, and their top halves span at most $V + 3$ rounded values with $V = 2.88, 1.64, 1.09, 1.01$ the increment's variation (they span $4, 3, 2, 2$), as the formula predicts. The $C = 2$ balanced restriction therefore exhibits the $N^{1/12}$ stationary scale of Theorem 21.2 by a different mechanism, and the appropriate conjecture is an envelope, not a bounded-statistic conjecture.

Conjecture 19.4 (envelope of the unique-product statistic)**.** *For every $C > 1$ and $\varepsilon > 0$ there is $r_0 = r_0(C, \varepsilon)$ such that, for all sufficiently large N and all integers $r_0 \le r \le N^{1/3}$, $D_{\max}(\mathcal{F}_{\mathrm{uniq}}(r), 1) \le N^{\varepsilon} \max\{1, rN^{-1/4}\}$, where $\mathcal{F}_{\mathrm{uniq}}(r)$ is formed from the balanced family with the displayed balance constant C.*

By Lemma 19.3 the envelope is attained up to constants for $r \ge 47N^{1/4}$ whenever the half-offset cells are representatives. The corresponding Lemma-4.3 values of the singleton starts at $r = \lfloor N^{1/3} \rfloor$ are $30, 28$; $37, 34$; $48, 44$; $57, 58$; $82, 81$; $102, 102$; $142, 142$; $179, 194$; $233, 257$; 300 at 30–48 bits, against $N^{1/6} = 31$, 39, 50, 60, 76, 98, 125, 150, 196, 237 and $N^{1/5} = 62$, 82, 111, 136, 182, 245, 330, 409, 562, 707, and the energy form gives $30, 32$; $40, 39$; $55, 52$; $66, 68$; $93, 91$; $128, 120$; $166, 161$; $209, 226$; $296, 286$; 381.

Proposition 19.5 (conditional floor for controlled-size families testing the balanced starts)**.** *Assume Conjecture 19.4. Let $\mathcal{A}$ be an oblivious Lehman-cell family covering J_C, and let Z be the multiset of its shifted tested exponents $n - aN - b$, n ranging over the tested integers of the form (a, b). Suppose that, for some integer $r \le N^{1/3}$: (1) for every $(a, b) \in \mathcal{F}_{\mathrm{bal}}(r)$ the standard shifted start $z_{a,b} := \lceil 2\sqrt{abN} \rceil - aN - b$ belongs to $\mathrm{supp}(Z)$; (2) the number $P(\mathcal{A})$ of forms of $\mathcal{A}$ satisfies $P(\mathcal{A}) \ll_C r \log(2r)$; (3) $|\mathrm{supp}(Z)| \ge N^{-o(1)}\Sigma_W$. Then every explicit difference representation $\mathrm{supp}(Z) \subseteq B - G$ satisfies $|B| + |G| \gg_C N^{5/28-o(1)}$. Condition (3) holds in particular when the shifted tested exponents are pairwise distinct. The proposition applies to the standard Lehman–Harvey family at every parameter for which it covers J_C, and to the standard balanced restriction whenever that restriction is used as an oblivious covering family of J_C.*

Proof. If $r \le N^{2/7}$, Lemma 3.2, $\Sigma_W \ge \Sigma_w/2$ and (2) give $\Sigma_W \gg_C \sqrt{N}/\sqrt{r \log(2r)} \ge N^{5/14-o(1)}$; by (3), $|B|\,|G| \ge |\mathrm{supp}(Z)| \ge N^{-o(1)}\Sigma_W$, so $|B| + |G| \ge 2\sqrt{|B|\,|G|} \gg_C N^{5/28-o(1)}$. Assume now $r \ge N^{2/7}$. By (1) the standard starts of $\mathcal{F}_{\mathrm{uniq}}(r)$ lie in $\mathrm{supp}(Z)$, and they are distinct for N large: cells with different a have starts differing by $N - O_C(\sqrt{rN}) > 0$, and for the same a and $b' = b + h > b$, using $b, b' \le Ca$,

$$z_{a,b'} - z_{a,b} = \lceil 2\sqrt{ab'N} \rceil - \lceil 2\sqrt{abN} \rceil - h > \frac{2\sqrt{aN}\,h}{\sqrt{b'} + \sqrt{b}} - 1 - h \ge h\big(\sqrt{N/C} - 1\big) - 1 > 0.$$

Since $r \ge N^{2/7} > N^{1/4}$, the conjecture gives $D := D_{\max}(\mathcal{F}_{\mathrm{uniq}}(r), 1) \le N^{o(1)} rN^{-1/4}$, and $n' := \#\mathcal{F}_{\mathrm{uniq}}(r) \ge r^{1-o(1)}$. Lemma 4.3 applied to these disjoint singleton windows yields $|B| + |G| \ge \min\{n'/2, 2(n'^2/(4D))^{1/3}\}$; the second term is $\gg N^{-o(1)} r^{1/3} N^{1/12}$, the first is larger in this range, and $r^{1/3}N^{1/12} \ge N^{2/21+1/12} = N^{5/28}$. For the standard family $P = O(r \log(2r))$ and for the balanced restriction $P = O_C(r)$; their shifted standard windows are pairwise disjoint for N large, so (3) holds for both, and (1) holds by construction. □

Three remarks. First, the exponent $5/28 = 0.178\ldots$ lies between the free-giant floor $1/6$ and the $2/11$ that a bounded exact statistic would have given: with a bounded statistic the two scales are $N^{1/4}r^{-1/4}$ and $r^{2/3}$, balancing at $r = N^{3/11}$ to give $N^{2/11}$; under the second branch of Conjecture 19.4, motivated for $C \ge 2$ by Lemma 19.3, the second scale becomes $N^{1/12}r^{1/3}$, balancing at $r = N^{2/7}$ to give $N^{5/28}$. Second, Lemma 4.3 must here be applied to the singleton starts with $W = 1$: at the family's own window length the half-offset families cohere – the

census measures $D_{\max}(\mathcal{F}_{\text{uniq}}(r), W) = 14, 24, 34, 52$ at $W = 8, 13, 20, 32$, namely 0.42–0.48 times $\sqrt{r}$ – so $D_{\max}$ is of the same scale as W and the $W^{1/3}$ gain that produces the exponent $1/5$ in Theorem 4.4's $\eta = 0$ case is unavailable. Third, the proposition offers a route to a floor for the standard family resting on a different hypothesis from the envelope conjectures (S″) and (P′): an exact-statistic envelope for the unique-product sub-family, with an explicit half-offset mechanism that attains the proposed scale in every family containing the constructed cells – for $\mathcal{F}_{\text{uniq}}$ itself this attainment is conditional on those cells being selected as representatives – rather than a tolerance-statistic envelope for the shell. The two conjectural routes give $5/28$ and $2/11$ respectively, neither upper envelope is proved, the remark on the provable envelopes applies to both, and Theorem 4.4 remains unconditional only in its stated dependence on the actual statistic.

Excision and the thinned statistic. Lemma 4.3 may be applied to the singleton starts of any sub-family of the tested set. Thus a useful dense-subfamily statistic is, for $\#\mathcal{F}_{\text{uniq}}(r) \geq 2$,

$$D^{\flat}(r;\epsilon) := \min\Big\{ D_{\max}(\mathcal{F}', 1) : \mathcal{F}' \subseteq \mathcal{F}_{\text{uniq}}(r),\ \#\mathcal{F}' \geq \max\big\{2, \lceil (1-\epsilon)\,\#\mathcal{F}_{\text{uniq}}(r)\rceil\big\}\Big\},$$

the balance constant C being suppressed from the notation. This statistic records the best exact-difference statistic obtainable while retaining a prescribed positive proportion of the singleton starts; the fully optimised Lemma-4.3 certificate also depends on the retained cardinality.

A possible sparse-excision heuristic is as follows. Put $K := kd$. A family in direction (k, d) has $O_C(\sqrt{r/K})$ admissible cells, so a family capable of contributing M pairs must have $K = O_C(r/M^2)$. Summing $O_C(\sqrt{r/K})$ over the primitive balanced directions $k \leq d \leq Ck$, $K \leq r/M^2$, gives $O_C(r/M)$ incidences if one charges only $O_C(1)$ relevant shifts per direction. For an integer shift (k', d') of the direction (k, d) write $X := kd' + dk'$ and $Y := kd' - dk'$. For a fixed $\Delta = Y/2$ the window condition is $\Delta^2|X| \leq 4K/M$, while the two congruences place X in a prescribed lattice class; a generic lattice-density calculation predicts about $16/M$ admissible shifts per direction after summing over Δ, with exceptional small representatives and boundary classes requiring separate treatment. Thus, under the additional heuristic assumptions that $\#\mathcal{F}_{\text{uniq}}(r) \asymp_C r$ and that overlaps do not change the scale, the coarse count suggests an $O_C(1/M)$ deleted fraction. This is a heuristic upper-scale calculation, not a proved asymptotic or a proved lower bound.

The excision census starts from the primitive balanced family and keeps, for each represented product, its smallest-a cell. For each threshold M it scans coprime directions $k \leq d \leq Ck$, $kd \leq 4r/M^2$. For an integer shift (k', d'), with X and Y as above, it requires X and Y to be even and nonzero, $X + Y \equiv 0 \pmod{2k}$, $X - Y \equiv 0 \pmod{2d}$, and $4kd/((Y/2)^2|X|) \geq M$. It then scans the affine lines $(a, b) = (k\lambda + \alpha, d\lambda + \beta)$ over the integer solutions (α, β) of $k\beta - d\alpha = -Y/2$ (one representative modulo translation by (k, d)), and marks both (a, b) and $(a + k', b + d')$ whenever both cells are positive, satisfy $a \leq b \leq Ca$, and have product at most r. The partner need not itself be primitive or belong to $\mathcal{F}_{\text{uniq}}(r)$. The residual set is obtained by deleting from $\mathcal{F}_{\text{uniq}}(r)$ every endpoint marked in this way. (The census works with the negated starts $aN + b - \lceil 2\sqrt{abN}\rceil$; simultaneous negation leaves the exact statistic unchanged.)

At 42 and 48 bits the removed counts for $M = 3, 4, 6, 8$ are respectively 3064, 2933, 1943, 1164 of 3149 and 10839, 10447, 7111, 4496 of 11009, namely 97.3%, 93.1%, 61.7%, 37.0% and 98.5%, 94.9%, 64.6%, 40.8%. The residual exact statistics are $3, 3, 5, 5$ at both sizes. These are the statistics of the four particular residual sets produced by the scan, not computed values of the minimum $D^{\flat}$. For $M = 8$ the residual sets retain 1985 of 3149 and 6513 of 11009 cells, and therefore give the finite-sample bounds $D^{\flat}(15728; \frac{1}{2}) \leq 5$ and $D^{\flat}(56099; \frac{1}{2}) \leq 5$. The deletion fractions decrease on a rough $1/M$ scale, but "$\approx 3/M$" is not a quantitative fit: $3/M$ predicts 100%, 75%, 50%, 37.5%, substantially underestimating the observed fractions at $M = 4$ and

$M = 6$. At the 48-bit point the residual maximisers for $M = 3, 4$ are the three same-line pairs $(146 + j, 291 + 2j) \mapsto (147 + j, 293 + 2j)$, $0 \le j \le 2$, whose shift has $Y = 0$ and is therefore excluded by the scan; for $M = 6, 8$ the five pairs are $(187 + 3j, 249 + 4j) \mapsto (188 + 3j, 251 + 4j)$, $0 \le j \le 4$, the direction $(3, 4)$ with shift $(1, 2)$ and predicted window $4 \cdot 3 \cdot 4/(1^2 \cdot 10) = 4.8$. The 42-bit residual maximisers at $M = 6, 8$ are carried by other pairs. The census records only one maximising difference and its pairs at each threshold; it does not enumerate or classify every cluster of $\mathcal{F}_{\mathrm{uniq}}(r)$.

Conjecture 19.6 (thinned statistic)**.** *For every $C > 1$ and every fixed $\epsilon \in (0, 1)$ there is $r_0 = r_0(C, \epsilon)$ such that, for every $\delta > 0$, there is $N_0 = N_0(C, \epsilon, \delta)$ for which $D^\flat(r; \epsilon) \le N^\delta$ for every $N \ge N_0$ and every integer $r_0 \le r \le N^{1/3}$.*

Corollary 19.7. *Assume Conjecture 19.6. Let $\mathcal{A}$, Z, r, B and G satisfy hypotheses* (1)–(3) *and the explicit-difference-representation hypothesis of Proposition 19.5. Then $|B| + |G| \gg_C N^{2/11-o(1)}$. Thus the conclusion of Proposition 19.5 holds with* $2/11$ *in place of* $5/28$.

Proof. Suppose first that $r \le N^{3/11}$. By Lemma 3.2, hypothesis (2), and $\Sigma_W \ge \Sigma_w/2$, $\Sigma_W \gg_C \sqrt{N}/\sqrt{r \log(2r)} \ge N^{1/2-3/22-o(1)} = N^{4/11-o(1)}$. Hypothesis (3) and $\mathrm{supp}(Z) \subseteq B - G$ give $|B|\,|G| \ge |\mathrm{supp}(Z)| \ge N^{-o(1)}\Sigma_W \ge N^{4/11-o(1)}$, hence $|B| + |G| \ge 2\sqrt{|B|\,|G|} \gg_C N^{2/11-o(1)}$. Now suppose that $r \ge N^{3/11}$. Apply Conjecture 19.6 with $\epsilon = \frac{1}{2}$, and choose $\mathcal{F}' \subseteq \mathcal{F}_{\mathrm{uniq}}(r)$ realising $D^\flat(r; \frac{1}{2})$. By hypothesis (1) the standard start of every cell of $\mathcal{F}'$ belongs to $\mathrm{supp}(Z)$; these starts are pairwise distinct by the argument in the proof of Proposition 19.5; thus they form $n' := \#\mathcal{F}' \ge \frac{1}{2}\#\mathcal{F}_{\mathrm{uniq}}(r) \ge r^{1-o(1)}$ disjoint singleton windows contained in $B - G$. The conjecture gives $D^\flat(r; \frac{1}{2}) \le N^{o(1)} = r^{o(1)}$ in this range, and Lemma 4.3 with $W = 1$ yields $|B| + |G| \ge \min\{n'/2, 2(n'^2/(4D^\flat(r; \frac{1}{2})))^{1/3}\}$. The first term is at least $r^{1-o(1)}$ and the second at least $r^{2/3-o(1)}$; moreover $D^\flat \ge 1$, so for large N the second is the smaller branch. Consequently $|B| + |G| \gg r^{2/3-o(1)} \ge N^{2/11-o(1)}$. □

The two conjectures control different aspects of the statistic, and neither formally implies the other as stated. Conjecture 19.4 bounds the unthinned maximum, allowing it to grow as large as $N^{o(1)}rN^{-1/4}$; Lemma 19.3 shows that this scale is attained up to constants whenever the relevant half-offset cells are selected as the unique-product representatives. It yields the exponent $5/28$. Conjecture 19.6, by contrast, permits the deleted exceptional set to have a much larger statistic, but asserts that for every fixed deletion allowance $\epsilon > 0$ one can retain at least $(1 - \epsilon)\#\mathcal{F}_{\mathrm{uniq}}(r)$ cells with residual statistic $N^{o(1)}$. It yields the exponent $2/11$.

A proof of Conjecture 19.6 would require at least two ingredients. First, a uniform incidence bound for the cells of the actual representative set $\mathcal{F}_{\mathrm{uniq}}(r)$ that participate in high-window half-offset configurations, including the exceptional congruence classes in which the prescribed class of X modulo $2kd$ has an unusually small representative and including overlaps between different families; the heuristic target is an $O_C(r/M)$ bound for the cells deleted at threshold M. What is provable cheaply falls short of it: for a direction (k, d) the admissible X is $\equiv Y\hat{c} \pmod{2kd}$ with $\hat{c} \equiv -1 \pmod{2k}$, $\hat{c} \equiv 1 \pmod{2d}$, so a resonant shift is an even $Y \ne 0$ with $0 < \|Y\theta\| \le 8/(MY^2)$, $\theta := \hat{c}/(2kd)$; by Legendre's theorem such Y are multiples of convergent denominators q_i of θ with $q_{i+1} > Mq_i^2/16$, at most $(16q_{i+1}/(Mq_i^2))^{1/3}$ of them each, whence at most $5(32kd/M)^{1/3} + 8/M + 1$ resonant shifts per direction and $O(r^{4/3}/M^2)$ deleted cells in all – $o(r)$ only for $M \gg r^{1/6}$, which through Lemma 4.3 gives an exponent below $5/28$. The heuristic average $16/M$ per direction would follow from equidistribution of the inverses $\bar{d} \bmod k$ in windows of length $16k/(MY^2)$, which the Weil bound for Kloosterman sums cannot see once $Y \gtrsim (k/M)^{1/4}$. Second, an inverse or residual estimate proving that, after this deletion, every exact cluster has size $N^{o(1)}$; such an estimate must cover low-predicted-window half-offset

families, the same-line $Y = 0$ configurations omitted by the scan, pile-ups of several families, and clusters not carried by a single half-offset family. The 42- and 48-bit census records only one residual maximising difference for each of four thresholds; it illustrates the proposed excision but does not enumerate all clusters and gives no evidence that non-half-offset residual clusters are absent.

Returning to Theorem 4.4: the exponent of its uniform-shell consequence is $1/5$ at $\eta = 0$, $2/11$ at $\eta = 1/3$ and the floor $1/6$ at $\eta = 1/2$. The sub-family to which Lemma 4.3 is applied is a free choice: thinning by a factor f costs $f^{2/3}$ in the Sidon branch, so a polylogarithmic thinning costs $N^{o(1)}$, and the exponent obtained is a supremum over admissible sub-families. The rest of this part determines the envelopes η of the full, squarefree and prime shells.

20. The additive structure of the window starts

20.1. **A linear point system.** For the cells $(1, k)$, with $u := 2\sqrt{N}$ and $\lceil 2\sqrt{kN}\rceil = u\sqrt{k} + \varepsilon_k$, $\varepsilon_k \in [0, 1)$,

$$s_k - s_{k'} = u\,L - \delta + (\varepsilon_k - \varepsilon_{k'}), \qquad L := \sqrt{k} - \sqrt{k'}, \quad \delta := k - k'.$$

The ordered pairs of the shell are points $p_i(u) = uL_i - \delta_i$ moving linearly in u, and $D_{\max}(\mathcal{F}_1(r), W)$ is the largest number of them inside an integer window $|p - t| < W$, $t \neq 0$, at time $u(N)$, up to the ceiling perturbation.

Proposition 20.1 (speeds)**.** *Two distinct ordered pairs (k, k'), (k'', k''') of the shell with $k > k'$, $k'' > k'''$ have the same speed if and only if there exist a squarefree integer j and integers $m, m', s \geq 1$ with $k = jm^2$, $k' = j(m - s)^2$, $k'' = jm'^2$, $k''' = j(m' - s)^2$ (the negative-speed case follows by reversing both pairs). In particular the speeds of the squarefree shell are pairwise distinct.*

Proof. Write each k as jm^2 with j squarefree. The relation $\sqrt{k} - \sqrt{k'} - \sqrt{k''} + \sqrt{k'''} = 0$ is an integer combination of square roots of squarefree integers, which are linearly independent over $\mathbb{Q}$, so the coefficients of equal radicals cancel; checking the possible pairings of the four radicals leaves only the stated case. □

Equal speeds give *persistent* clusters (for every u or for none); distinct speeds give *transient* ones, within $2W$ only for u in an interval of length $< 4W/|L_i - L_j|$. Write $D(t; u) := \#\{i : |uL_i - \delta_i - t| < W\}$ for the rounding-free count at the window centred at t. For u uniform on an interval of length U,

$$\mathbb{E}_u D(t; u) \leq \sum_i \min\Big(1, \frac{2W}{|L_i|U}\Big) \leq \frac{4W\sqrt{r}}{U} \sum_{k \neq k'} \frac{1}{|k - k'|} \leq \frac{4W\sqrt{r}}{U} R'(2 + 2\ln R'),$$

using $|L| \geq |k - k'|/(2\sqrt{r})$, with R' the number of cells of the shell under consideration; for $N \in [X, 2X]$, u uniform on $[2\sqrt{X}, 2\sqrt{2X}]$, $W \leq \sqrt{X}/(2r^{3/2}) + 1$ and $R' \leq r/2$ this is $(2.42 + o(1))(1 + \ln r)$ uniformly in t. The statistic is therefore a maximum over $\asymp r^2$ windows of a quantity with logarithmic mean.

20.2. **Beatty chains.** For $k = jm^2$ the chirp $\lceil m \cdot 2\sqrt{jN}\rceil$ is a Beatty sequence in m, so the cells $(1, jm^2)$ and $(1, j(m-1)^2)$ have start differences $2\sqrt{jN} - j(2m - 1) + O(1)$, decreasing by $2j + O(1)$ per step. For $j = 1$ the shell contains $Q = (1 - 2^{-1/2})\sqrt{r} + O(1)$ squares, whose $Q - 1$ consecutive differences have gaps in $\{1, 2, 3\}$; hence

$$D_{\max}(\mathcal{F}_1(r), W) \;\geq\; \min\big(Q - 1, \lfloor (2W - 1)/3 \rfloor\big).$$

Since $W \gtrsim \sqrt{r}$ exactly when $r \lesssim N^{1/4}$, this layer alone forces $D_{\max} \gg \sqrt{r}$ there and neutralises Lemma 4.3 on the full $a = 1$ shell down to the M2 floor; it is absent from the squarefree shell (Proposition 20.1), which is why Theorem 4.4 is stated for $\mathcal{F}_1^{\square}$.

20.3. **Near-chains.** The speeds of the squarefree shell are distinct, but approximately equal speeds occur along arithmetic families. For $r \le N^{1/5}$ the offsets are a bounded fraction of a window, and the statistic reduces, up to a window enlargement (Lemma 20.9), to the *modulus-free* statistic

$$D_\theta^*(r) := \max_{\tau\in\mathbb{R}} \#\{(k,k') : k \ne k' \text{ squarefree in } (r/2, r],\ |\sqrt{k} - \sqrt{k'} - \tau| < \theta\rho_r\},$$
$$\rho_r := \frac{W_{\text{real}}}{u} = \frac{1}{4\sqrt{2}\, r^{3/2}},$$

and we write $D^*(r) := D_1^*(r)$.

Proposition 20.2 ($e = 1$ near-chains)**.** *Let j be odd and squarefree and $\tau_j := \sqrt{j/2}$. The odd M with $M^2 \equiv 1 \pmod{8j}$ form $2^{\omega(j)}$ residue classes modulo $2j$; for M in one class put $k_M := (M^2-1)/(8j)$. Then (i) every k_M is an integer; (ii) consecutive members satisfy*

$$\sqrt{k_{M+2j}} - \sqrt{k_M} = \tau_j + \frac{\tau_j}{2M(M+2j)} + O\Big(\frac{\tau_j}{M^4}\Big);$$

(iii) a window of width $2\rho_r$ on the chain contains $M^3/(4r^{3/2}j^{3/2}) + O(1)$ consecutive links per class, which is 2 at the bottom and 5.66 at the top of the shell; (iv) each class has $(\sqrt{2}-1)\sqrt{r/j} + O(1)$ members in the shell.

Proof. (i): $(M_0+2jn)^2 - M_0^2 = 4jn(M_0+jn)$ with $n(M_0+jn)$ even. (ii): $\sqrt{k_M} = (8j)^{-1/2}(M - 1/(2M) - \ldots)$ and $\frac{1}{2}(1/M - 1/(M+2j)) = j/(M(M+2j))$. (iii) and (iv) follow from the decrease $\sim 2j\tau_j/M^3$ of the link speed per link and $M^2 = ajr$, $a \in (4, 8]$, on the shell. □

The chains have $\tau_j^2 = j/2 \notin \mathbb{Z}$, so τ_j is never an exact repeated speed; they are the nearest squarefree-compatible relatives of the Beatty chains, with $8jk+1$ a square instead of $8jk$. Conditional on a uniform positive squarefree-link density along the chains, which we do not prove and do not need, their clusters grow like $2^{\omega(j)}$ with $j \lesssim r/190$, i.e. like $r^{(\ln 2 - o(1))/\ln\ln r}$; this would exclude a polylogarithmic envelope.

20.4. **Two-progression families.**

Proposition 20.3. *For an ordered pair $k < k'$ put $d := k' - k$ and $s := k + k'$. (i) If $(k, k') = (P_1(n), P_2(n))$ with $P_i(n) = an^2 + b_i n + c_i$, $a > 0$, $b_1 \ne b_2$, then $s = Ad^2 + Bd + C$ with $A = 2a/(b_2 - b_1)^2 > 0$ and $B = [4a(c_1 - c_2) + b_2^2 - b_1^2]/(b_2 - b_1)^2$; conversely any set of pairs on a parabola $s = Ad^2 + C$ is the family $((Ad^2 - d + C)/2, (Ad^2 + d + C)/2)$. Call the case $B = 0$* drift-free. *(ii) For a drift-free family, with $\Delta_d := 1/4 - AC$,*

$$\sqrt{k'} - \sqrt{k} = \frac{d}{\sqrt{k'} + \sqrt{k}} = \frac{1}{\sqrt{2A}}\Big(1 + \frac{\Delta_d}{2A^2d^2} + O(d^{-4})\Big).$$

Proof. (i) is algebra. (ii): $\sqrt{k'} + \sqrt{k} = \sqrt{2s}\,(1 - d^2/(8s^2) + O(d^4/s^4))$ with $s = Ad^2 + C$. □

The near-chains are the case $A = 1/j$, $C = (j^2-1)/(4j)$; the *symmetric family* below is $A = j$, $C = 0$.

Lemma 20.4 (integrality)**.** *Let $\mathcal{F} = \{((Ad^2 - d + C)/2, (Ad^2 + d + C)/2) : d \equiv D_0 \pmod q\}$ consist of pairs of integers, with q the least positive period of the set of admissible d. Then $\alpha := Aq^2$ and $\gamma := Cq^2$ are integers with $q \mid 2\alpha D_0$, $q^2 \mid \alpha D_0^2 + \gamma$, and*

$$\Delta_d = \frac{M}{4q^2}, \qquad M := \frac{q^4 - 4\alpha\gamma}{q^2} \in \mathbb{Z}.$$

Moreover (i) $C = 0 \iff M = q^2 \iff \Delta_d = 1/4$; (ii) if $C \neq 0$ then $\alpha \le q|q^2 - M|/2$; (iii) the polynomials $2k_\mp(m)$, $d = D_0 + qm$, have discriminant M, so a square M forces all but finitely many members to be composite; infinitely many prime members require a non-square M, hence $C \neq 0$, and since $M \equiv \beta^2 \pmod 8$ for the integral linear coefficient β of $2k_\mp$ (its constant term being even), $M \equiv 0, 1, 4 \pmod 8$ and $|M| \ge 4$; (iv) $M = 0$ is the Beatty case of Proposition 20.1.

Proof. With $d = D_0 + qm$, $k_\mp(m) = \frac{\alpha}{2}m^2 + \frac{1}{2}(2\alpha D_0/q \mp q)m + \frac{1}{2}(AD_0^2 \mp D_0 + C)$. A quadratic is integer-valued iff twice its leading coefficient, the sum of its leading and linear coefficients and its constant term are integers, which gives $\alpha \in \mathbb{Z}$, $2\alpha D_0/q \in \mathbb{Z}$, $AD_0^2 + C \in \mathbb{Z}$, hence $\gamma \in \mathbb{Z}$ and $q^2 \mid \alpha D_0^2 + \gamma$; then $4\alpha\gamma \equiv -(2\alpha D_0)^2 \equiv 0 \pmod{q^2}$. (i) is immediate. (ii): with $g = \gcd(q, D_0)$, $q = gq'$, one finds $2\alpha = q'\alpha_1$, $\gamma = g^2\gamma_1$, $2\gamma_1 = q'\gamma_2$ and $\alpha_1\gamma_2 = q^2 - M$, so $\alpha \le q'|q^2 - M|/2$. (iii): the discriminant of $2k_\mp(m)$ is $q^2 - 4\alpha C = M$ for both signs; a square discriminant makes $2k_\mp$ factor over $\mathbb{Q}$. (iv): $\Delta_d = 0$ gives $4kk' = (Ad^2 - 1/(4A))^2$. □

To first order, a drift-free family on one class has at most $1 + \sqrt{2}\theta A/(q|\Delta_d|)$ members in a window of half-width $\theta\rho_r$ and $(\sqrt{2} - 1)\sqrt{r/A}/q$ members in the shell; maximising over A gives a potential $0.62\,\theta^{1/3}|\Delta_d|^{-1/3}r^{1/3}/q$ per class. Hence, per class, $C = 0$ families are maximal at $q = 1$ (the symmetric family), and families with $C \neq 0$ and $M \neq 0$ – in particular every family with infinitely many prime members – have first-order occupancy $O(n|M|^{-3/7}r^{2/7})$, where n denotes the number of residue classes modulo q on which the family is integral, the $r^{2/7}$ *law* (first order); the factor n makes multi-class families competitive at finite r (Section 20.5).

20.5. **The symmetric family and Theorem Q′.** For odd j and $t \ge 1$ let $f_\pm(t) = (jt^2 \pm t)/2 = t(jt \pm 1)/2$, the two integers at distance $t/2$ from $(j/2)t^2$. Exactly

$$\sqrt{f_+} - \sqrt{f_-} = \frac{t}{\sqrt{f_+} + \sqrt{f_-}} = \frac{1}{\sqrt{2j}}\Big(1 + \frac{1}{8j^2t^2} + O(j^{-4}t^{-4})\Big),$$

so the family has speed $1/\sqrt{2j}$ with a t^{-2} drift; among one-class integral two-progression pairs in the normal form $at^2 \pm bt + c$ it has the largest per-class capacity (the smallest $b|\Delta|$); multi-class families with $C \neq 0$ can have larger total clusters (Conjecture 20.12).

Lemma 20.5 (local conditions)**.** *For an odd prime p, $p^2 \mid f_\mp(t)$ iff $p^2 \mid t$ or $p^2 \mid jt \mp 1$; the set B_p of bad residues modulo p^2 is $\{0, \pm j^{-1}\}$ for $p \nmid j$ and $\{0\}$ for $p \mid j$. Modulo 8, $B_2 = \{0, \pm j\}$. Hence $f_-(t), f_+(t)$ are both squarefree iff $t \notin B_p$ for every prime p.*

Proof. $\gcd(t, jt \pm 1) = 1$; for $p = 2$ check the parities of t and use $j^2 \equiv 1 \pmod 8$. □

Lemma 20.6 (sieve)**.** *Let I be a set of L consecutive positive integers with $\max I = T_2$, and $z \ge 3$. Then*

$$\#\{t \in I : f_-(t), f_+(t) \text{ squarefree}\} \;\ge\; \kappa_j L - \frac{3L}{z} - 4^{\pi(z)} - \pi(\sqrt{T_2}) - 2\pi\big(\sqrt{jT_2 + 1}\big),$$

$$\kappa_j := \frac{5}{8}\prod_{p \mid j}\Big(1 - \frac{1}{p^2}\Big)\prod_{p \nmid 2j}\Big(1 - \frac{3}{p^2}\Big) \;\ge\; \kappa_0 := \frac{5}{8}\prod_{p \ge 3}\Big(1 - \frac{3}{p^2}\Big) = 0.313717\ldots$$

(equality at $j = 1$; $\kappa_{15} = 0.456316\ldots$).

Proof. Let G be the set of $t \in I$ avoiding B_p for all $p \le z$, and let E_G be the set of those $t \in G$ lying in B_p for some $p > z$. Inclusion–exclusion over squarefree d composed of primes $\le z$ gives $|G| = L\prod_{p\le z}(1-|B_p|/m_p) + \theta'$ with $|\theta'| \le \prod_{p\le z}(1+|B_p|) \le 4^{\pi(z)}$, and the product is $\ge \kappa_j$. For $p > z \ge 3$ odd, $t \in B_p$ means $p^2 \mid t$ (so $p \le \sqrt{T_2}$) or $p^2 \mid jt \mp 1$; for $t \in G$ the integer $jt \mp 1$ is nonzero (the only zero, $(j,t) = (1,1)$, is excluded by B_2), so $p \le \sqrt{jT_2+1}$ and the solutions form one class modulo p^2; summing $L/p^2 + 1$ over these primes with $\sum_{p>z} p^{-2} < 1/z$ gives the bound on $|E_G|$. □

Lemma 20.7 (window fit)**.** *Let $I_j(r) := \{t \ge 1 : r/2 < f_-(t),\ f_+(t) \le r\}$ with $jr \ge 4$, and $\varepsilon = 1/(jt)$. Every $t \in I_j(r)$ has $\sqrt{r/j} < t \le \sqrt{2r/j}$, hence $1/(2jr) \le \varepsilon^2 < 1/(jr)$, and $s(t) := \sqrt{f_+} - \sqrt{f_-}$ satisfies $1 + \varepsilon^2/8 \le s(t)\sqrt{2j} \le 1 + \varepsilon^2/8 + \varepsilon^4/8$. The speeds of all pairs of $I_j(r)$ lie in an interval of length at most $(1+2/(jr))/(16\sqrt{2}\,j^{3/2}r)$, whose ratio to a window of half-width $\theta\rho_r$ is $r^{1/2}(1+2/(jr))/(8\theta j^{3/2})$; so all of them lie in one window whenever $j \ge (1+2/r)(8\theta)^{-2/3}r^{1/3}$. Moreover $|I_j(r)| = \lfloor B\rfloor - \lfloor A\rfloor \ge (\sqrt{2}-1)\sqrt{r/j} - 3$ with $A = (1+\sqrt{1+4jr})/(2j)$, $B = (-1+\sqrt{1+8jr})/(2j)$.*

Proof. For $\varepsilon \le 1/2$, $\sqrt{1+\varepsilon} + \sqrt{1-\varepsilon} = 2 - \varepsilon^2/4 - \eta$ with $0 \le \eta \le \varepsilon^4/9$ (the even binomial coefficients $\binom{1/2}{2m}$, $m \ge 1$, are negative and decreasing in absolute value); the rest is arithmetic. □

Theorem 20.8 (Q′)**.** *For every fixed $\theta \in (0,1]$, as $r \to \infty$,*

$$D^*_\theta(r) \ \ge\ \big(\kappa_0(\sqrt{2}-1)(8\theta)^{1/3} - o(1)\big)\, r^{1/3};$$

numerically $D^*_1(r) \ge (0.2598 - o(1))r^{1/3}$ and $D^*_{1/2}(r) \ge (0.2062 - o(1))r^{1/3}$; with j restricted to odd multiples of 15 the constant at $\theta = 1$ is 0.378.

Proof. Let j be the smallest odd integer $\ge (1+2/r)(8\theta)^{-2/3}r^{1/3}$ (or the smallest such odd multiple of 15, in which case κ_0 is replaced by κ_{15} below). By Lemma 20.7 every pair of $I_j(r)$ lies in one window of half-width $\theta\rho_r$, $L = |I_j(r)| = (\sqrt{2}-1)(8\theta)^{1/3}r^{1/3}(1 + O(r^{-1/3}))$ and $T_2 = \sqrt{2}(8\theta)^{1/3}r^{1/3}(1+O(r^{-1/3}))$, so $j = O_\theta(T_2)$ and $L \asymp_\theta T_2$. Lemma 20.6 with $z = \ln T_2$ has error $O_\theta(L/\ln L)$, so the squarefree pairs in the window number $\ge \kappa_0 L(1 - O(1/\ln L))$. □

Lemma 20.9 (transfer)**.** *For fixed $\varepsilon > 0$ and $\theta < 1 < \theta'$, there is $N_0 = N_0(\theta,\theta',\varepsilon)$ such that, whenever $r \le N^{1/5-\varepsilon}$ and $N \ge N_0$,*

$$D^*_\theta(r) \le D_{\max}(\mathcal{F}_1^{\square}(r), W) \le D^*_{\theta'}(r),$$

and the same holds for every subset $K \subseteq (r/2, r]$ of cells $(1,k)$ with its own statistics.

Proof. The offsets $(k-k')/u$, the ceiling terms and the rounding of the real centre $u\tau$ to an integer move each point by at most $r/u + 2/u = o(\rho_r)$, since $r/W \approx 2\sqrt{2}\,r^{5/2}/\sqrt{N} \to 0$. Thus every set of speeds inside a real window of half-width $\theta\rho_r$ gives a set of the same cardinality of starts inside an integer window of half-width W, and conversely. This is a window-enlargement inequality, not an additive $O(1)$ statement. □

Corollary 20.10. *For fixed $0 < \alpha \le 1/5 - \varepsilon$ and $r = N^\alpha$, $D_{\max}(\mathcal{F}_1^{\square}(r), W) \ge (0.20 - o(1))N^{\alpha/3}$. In particular the envelopes $D_{\max} \le (\log N)^{O(1)}$ and $D_{\max} \le N^{o(1)}$ are false for the squarefree $a = 1$ shell, and Theorem 4.4 cannot give more than $N^{2/11-o(1)}$ through it.*

The measured clusters of the symmetric family are $D_{\mathrm{sym}}/r^{1/3} \in [0.347, 0.391]$ for $2^{18} \le r \le 2^{32}$, at $j \approx 0.25\,r^{1/3}$ and squarefree density ≈ 0.46 (Section 23). The constant is not a free parameter: the mechanism has a sharp asymptotic constant, attained from below as the small-prime content of the maximising j grows.

Theorem 20.11 (Q″, the sharp constant of the symmetric mechanism)**.** *Let* $\kappa_\infty := \frac{5}{8}\prod_{p\ge3}(1-p^{-2}) = 5/\pi^2 = 0.50661\ldots$ *and* $c_\theta := (\sqrt{2}-1)(8\theta)^{1/3}\kappa_\infty$*, so that* $c_1 = 0.4196\ldots$ *and* $c_{1/2} = 0.3331\ldots$*. Let* $D^*_{\mathrm{sym},\theta}(r)$ *be the largest number of symmetric pairs* $((jt^2-t)/2,(jt^2+t)/2)$*,* $t\in I_j(r)$*, of squarefree cells of the shell whose speeds lie in one window of half-width* $\theta\rho_r$*, the maximum being taken over the odd* j *and the window. Then (i)* $D^*_\theta(r) \ge D^*_{\mathrm{sym},\theta}(r) \ge (c_\theta - o(1))\,r^{1/3}$*; (ii)* $D^*_{\mathrm{sym},\theta}(r) \le (c_\theta+o(1))\,r^{1/3}$*. Thus* $D^*_{\mathrm{sym},\theta}(r) = (c_\theta+o(1))\,r^{1/3}$*; the theorem does not bound windows that collect pairs from several* j*-families at once.*

Proof. (i) Run the proof of Theorem 20.8 with j the smallest odd multiple of $P_y := \prod_{3\le p\le y} p$, $y := \lfloor(\log r)/10\rfloor$, that is $\ge (1+2/r)(8\theta)^{-2/3}r^{1/3}$. Since $\vartheta(y) = \sum_{p\le y}\log p < 1.02y$ for y large, $P_y \le r^{0.102}$, so $j = (8\theta)^{-2/3}r^{1/3}(1+O(r^{-0.23}))$ and $L = |I_j(r)| = (\sqrt{2}-1)(8\theta)^{1/3}r^{1/3}(1+o(1))$ as before. Every prime $3\le p\le y$ divides j, so $\kappa_j \ge \frac{5}{8}\prod_{3\le p\le y}(1-p^{-2})\prod_{p>y}(1-3p^{-2}) \ge \kappa_\infty\prod_{p>y}\frac{1-3p^{-2}}{1-p^{-2}} \ge \kappa_\infty\bigl(1-2.1\sum_{p>y}p^{-2}\bigr) \ge \kappa_\infty(1-2.1/y)$, and Lemma 20.6 with $z = \ln T_2$ gives at least $\kappa_j L(1-O(1/\ln L))$ squarefree pairs in one window. (ii) Fix odd j with $jr\ge4$ and a window. Exactly, $s(t) = \sqrt{2}/(\sqrt{j+1/t}+\sqrt{j-1/t})$, which is strictly decreasing in t, so the members with speeds in the window are the consecutive $t\in[t_1,t_2]\cap I_j(r)$ for some $t_1\le t_2$ with $s(t_1)-s(t_2) < 2\theta\rho_r$. By Lemma 20.7, $s(t)\sqrt{2j} = 1+\varepsilon^2/8+O(\varepsilon^4)$ with $\varepsilon = 1/(jt) \le (jr)^{-1/2}$, and the $O(\varepsilon^4)$ terms contribute $O(j^{-5/2}r^{-2})$ to $s(t_1)-s(t_2)$, which is $o(\theta\rho_r)$; hence $t_1^{-2}-t_2^{-2} < \Delta := 4\theta j^{5/2}r^{-3/2}(1+o(1))$, i.e. $t_1 > (t_2^{-2}+\Delta)^{-1/2}$. The function $\Phi_\Delta(t) := t-(t^{-2}+\Delta)^{-1/2}$ has derivative $1-(1+\Delta t^2)^{-3/2} > 0$, so with $B_0 := \sqrt{2r/j} \ge t_2$ we get $t_2-t_1 < \Phi_\Delta(t_2) \le \Phi_\Delta(B_0)$, and since $\Delta B_0^2 = x(1+o(1))$ with $x := 8\theta j^{3/2}/\sqrt{r}$, the window contains at most

$$M \;\le\; \min\Bigl\{\sqrt{2r/j}\,\bigl(1-(1+x)^{-1/2}\bigr),\ (\sqrt{2}-1)\sqrt{r/j}\Bigr\}(1+o(1))+1$$

consecutive members. Let $\kappa_{j,z}$ be the finite local product over $p\le z$. The inclusion–exclusion in the proof of Lemma 20.6 gives, for any M consecutive t, at most $\kappa_{j,z}M + 4^{\pi(z)}$ values with both members squarefree, and $\kappa_j \le \kappa_{j,z} \le \kappa_j\prod_{p>z}(1-3p^{-2})^{-1} = \kappa_j(1+O(1/z))$ uniformly in j. Put $j_1 := 4(\sqrt{2}-1)^2c_\theta^{-2}r^{1/3}$. For $j\ge j_1$ the trivial bound $M \le (\sqrt{2}-1)\sqrt{r/j}+2 \le c_\theta r^{1/3}/2+2$ suffices. For $j\le j_1$ we have $M = O_\theta(r^{1/3})$, and taking $z := \log r$ gives at most $\kappa_j M + o(r^{1/3})$ squarefree pairs uniformly in j, since $4^{\pi(\log r)} = r^{o(1)}$. Writing $j = (x\sqrt{r}/(8\theta))^{2/3}$, $\sqrt{2r/j} = \sqrt{2}(8\theta)^{1/3}x^{-1/3}r^{1/3}$. Let $F(x) := x^{-1/3}(1-(1+x)^{-1/2})$ and $G(x) := x^{-1/3}(1-2^{-1/2})$: F is increasing on $0<x\le1$, G is decreasing on $x>0$, $F(1)=G(1)$, $F\le G$ for $x\le1$ and $F\ge G$ for $x\ge1$; hence $\min(F,G)$ is maximal at $x=1$, where it equals $1-2^{-1/2}$. Therefore the count is at most $(\kappa_j+o(1))\sqrt{2}(8\theta)^{1/3}(1-2^{-1/2})r^{1/3} = (\kappa_j+o(1))(\sqrt{2}-1)(8\theta)^{1/3}r^{1/3} \le (c_\theta+o(1))r^{1/3}$, since $1-3p^{-2}\le1-p^{-2}$ gives $\kappa_j\le\kappa_\infty$ for every odd j. □

Conjecture 20.12 (S‴, two-progression extremality)**.** *Let* $D^*_{2P,\theta}(r)$ *be the largest number of pairs of squarefree cells of one drift-free two-progression family (Proposition 20.3, Lemma 20.4) – a fixed* (A,C) *together with all the residue classes* $d\equiv D_0 \pmod q$ *on which its members are integers, of any least positive period* q *(families with* $q>r$ *have at most* $2^{\omega(q)} = r^{o(1)}$ *shell pairs altogether by Lemma 20.15 and do not affect the statement) – whose speeds lie in one window of half-width* $\theta\rho_r$*. For every fixed* $\theta\in(0,1]$*, as* $r\to\infty$*,* $D^*_\theta(r) = D^*_{2P,\theta}(r) + o_\theta(r^{1/3})$*.*

(S‴) implies (S″) of Section 22 in the transfer range, because Theorem 20.16 below proves $D^*_{2P,\theta}(r) \le 7\,r^{1/3}$ uniformly over the class; the remaining content of (S‴) is therefore not a capacity estimate but the inverse assertion that one family captures the maximum up to $o(r^{1/3})$. By Theorem 20.11, $D^*_{2P,\theta}(r) \ge (c_\theta-o(1))r^{1/3}$; at the sampled sizes below, the symmetric family is not the largest member of the class. To first order, a drift-free family with n integral classes

has capacity approximately $0.98\,\theta^{1/3} n\, q^{-1/3} |M|^{-1/3} r^{1/3}$ (from the rounded per-class coefficient 0.62) at its balance point $A^* = ((\sqrt{2}-1)|M|/(4\sqrt{2}\theta))^{2/3} r^{1/3} q^{-4/3}$, so families with $|M| = 1$ and $n = 2^{\omega(q)}$ classes have potential $nq^{-1/3} = 1.62$ ($q = 15$) and 1.70 ($q = 105$) times the symmetric family's. For fixed (q, M) with $C \neq 0$ the admissible α divide $q^2(q^2 - M)/4$, so such a family can track the $r^{1/3}$ law only while a divisor lies near $A^* q^2$: the largest divisor gives the necessary horizons $r \approx 2^{40.6}$ for $q = 15$ and $2^{68.7}$ for $q = 105$, a suitable divisor may fail earlier, and beyond the horizon the family is $O(1)$; the symmetric family, with $C = 0$, has no such constraint. The census of Section 23 (Table 6), restricted to $q \leq 120$, $0 < |M| \leq 64$ and α within a factor 6 of $A^* q^2$, is consistent with this first-order picture: at the sampled values $r = 2^{16}, 2^{18}, \ldots, 2^{30}$ its maximum exceeds the symmetric cluster by ratios from $172/157 = 1.095$ to $145/89 = 1.629$ (the counts at 2^{26} and 2^{30} are 200-bit values), attained by exactly these $q = 15$ and $q = 105$ families, and equals the exact D^* at $r = 2^{18}$. At each of the three values $r = 2^{16}, 2^{17}, 2^{18}$ the certified exact sweep returns a drift-free maximising family: near-chains at 2^{16} and 2^{17}, the $q = 15$ family at 2^{18}. Whether $\limsup D^*_\theta(r)/r^{1/3}$ exists, and whether it equals c_θ or is larger, is open: the divisor obstruction only shows that each fixed non-symmetric (q, M) family eventually leaves the $r^{1/3}$ balance law (for $\theta = |M| = 1$ the last smooth q with $2^{\omega(q)} q^{-1/3} \geq 1$ is forced out around $r \approx 2^{181}$, to first order), not that families with q or M varying with r, or windows collecting several families, cannot do better; the census maximum exceeds the finite symmetric cluster at every sampled size, and its normalised value exceeds c_θ at seven of the eight sizes (not at $r = 2^{16}$, where $16/r^{1/3} = 0.397$), and the per-j agreement of Table 7 corroborates the single-family density calculation without bearing on extremality.

20.6. **The two-progression envelope.** The supremum over drift-free families of the modulus-free contribution of a single family has exponent exactly $1/3$: Theorem 20.11 supplies a family of order $r^{1/3}$, Theorem 20.16 below bounds every family of least period $q \leq r$ by $7\, r^{1/3}$, and Lemma 20.15 disposes of $q > r$. Consequently Conjecture 20.12 implies the modulus-free $r^{1/3+o(1)}$ envelope, and hence the transfer-range part of (S″); the planar part of (S″) requires separate control of the offsets.

Lemma 20.13 (exact speed of a drift-free family)**.** *For a drift-free family* $(k_-(d), k_+(d)) = ((Ad^2 - d + C)/2, (Ad^2 + d + C)/2)$ *put* $S := Ad^2 + C = k_+ + k_-$, $w := 2\sqrt{k_+ k_-} = \sqrt{S^2 - d^2}$ *and* $P := 2AS - 1$. *The speed* $v(d) := \sqrt{k_+} - \sqrt{k_-}$ *satisfies* $v^2 = S - w$ *and, whenever* $2Aw + P \neq 0$,

$$1 - 2Av^2 = \frac{-4\Delta_d}{2Aw + P}, \qquad \frac{d}{dd} v^2 = \frac{-4\Delta_d\, d}{w\,(2Aw + P)}, \qquad \Delta_d = \tfrac{1}{4} - AC.$$

Proof. $v^2 = k_+ + k_- - 2\sqrt{k_+ k_-} = S - w$, so $1 - 2Av^2 = 2Aw - P$, and $(2Aw)^2 - P^2 = 4A^2(S^2 - d^2) - (2AS - 1)^2 = 4AS - 4A^2 d^2 - 1 = 4AC - 1 = -4\Delta_d$, which gives the first identity. Differentiating, $(v^2)' = S' - (SS' - d)/w = d(2Aw - 2AS + 1)/w = d(1 - 2Av^2)/w$. □

In particular, for $d > 0$ and $2Aw + P > 0$, v is strictly decreasing when $\Delta_d > 0$ and strictly increasing when $\Delta_d < 0$; and if $\Delta_d \neq 0$ then $2Aw + P$ never vanishes.

Lemma 20.14 (class count)**.** *Let a drift-free family have least positive period* q*, in the sense that the set* $\mathcal{A}$ *of integers* d *for which both members are integers has least positive period* q*, and let* n *be the number of residue classes modulo* q *contained in* $\mathcal{A}$*. Then* $n \leq 2^{\omega(q)}$.

Proof. Put $S(d) := Ad^2 + C$ and $H(d) := S(d) - d$, so that $d \in \mathcal{A}$ iff $S(d) \in \mathbb{Z}$ and $H(d)$ is even. Write $\alpha = Aq^2$, $\gamma = Cq^2$ and $f(x) := \alpha x^2 + \gamma = q^2 S(x)$. Each class $D \bmod q$ in $\mathcal{A}$ contains the q distinct residues $D + qm \bmod q^2$, $0 \leq m < q$, all roots of f modulo q^2, so $nq \leq \rho_f(q^2)$, where $\rho_f(m)$ denotes the number of roots of f modulo m; this uses only $S(d) \in \mathbb{Z}$, not the parity condition. For $p^e \parallel q$ let R_p be the set of residues $D \bmod p^e$ all of whose lifts modulo p^{2e} are

roots of f (the *robust root classes* at p). By the Chinese remainder theorem $\mathcal{A}$ is the intersection of the odd-prime local integrality sets with the joint 2-adic integrality-and-parity set; its least period is the product of the local least periods, and n is the product of the local class counts, so it suffices to prove a local bound of 2 at every prime dividing q. Put $a := v_p(\alpha)$ and $c := v_p(\gamma)$ (with $c = \infty$ if $\gamma = 0$).

Odd p. The parity condition cannot introduce an odd factor into the least period: on every arithmetic progression on which S is integral, H is an integer-valued quadratic whose reduction modulo 2 has period dividing 4. Hence the p-part of q is the p-part of the least period of the local integrality set $\{d : p^{2e} \mid f(d)\}$. If $a \geq 2e$, then $c < 2e$ gives no roots, while $c \geq 2e$ makes the local condition vacuous, contradicting the minimality of the p-part of q. If $a < 2e \leq c$, the roots are the x with $p^u \mid x$, $u := \lceil (2e-a)/2 \rceil \leq e$, so the local integrality set has least p-power period p^u; minimality gives $u = e$ and $|R_p| = 1$. If $a, c < 2e$, there is no root when $c < a$ or when $c - a$ is odd; otherwise write $c - a = 2t$, $\alpha = p^a\alpha'$, $\gamma = p^c\gamma'$ with α', γ' units, and $L := 2e - c$. Every root is $x = p^t u$ with u a unit satisfying $\alpha' u^2 \equiv -\gamma' \pmod{p^L}$, a congruence with $s \leq 2$ solutions modulo p^L whose solution set has no non-zero additive period modulo p^L; so the local integrality set has least p-power period p^{t+L}, minimality gives $e = t + L$, i.e. $a + t = e$, and the robust classes modulo p^e are exactly the s classes obtained from the solutions $u \bmod p^L$: $|R_p| = s \leq 2$.

$p = 2$. Here the parity condition is analysed jointly with integrality. Let $\mathbb{Z}_2$ denote the ring of 2-adic integers and $\mathcal{A}_2 := \{d : S(d) \in \mathbb{Z}_2,\ S(d) \equiv d \pmod 2\}$; we show that $\mathcal{A}_2$ consists of at most two classes modulo its least period 2^e. Recall $v_2(A) = a - 2e$ and $v_2(C) = c - 2e$. If $a, c \geq 2e$ then $A, C \in \mathbb{Z}_2$ and, since $d^2 \equiv d \pmod 2$, $S(d) - d \equiv (A-1)d + C \pmod 2$: $\mathcal{A}_2$ is $\mathbb{Z}$, one class modulo 2, or empty. If $a < 2e \leq c$ then $S(d) \in \mathbb{Z}_2$ iff $2^u \mid d$ with $u := \lceil (2e-a)/2 \rceil \geq 1$; along $d = 2^u j$ one has $v_2(A2^{2u}) \in \{0, 1\}$ and $S(d) - d \equiv A2^{2u}j + C \pmod 2$, so the parity condition rejects every j, accepts every j (one class modulo the least period 2^u), or accepts one parity class of j (one class modulo the least period 2^{u+1}): at most one class. If $c < 2e \leq a$ there is no integral value. If $a, c < 2e$, integrality forces $c \geq a$ and $c - a = 2t$ even, and the integral d are $d = 2^t u$ with u odd and $\alpha' u^2 + \gamma' \equiv 0 \pmod{2^L}$, where $\alpha = 2^a\alpha'$, $\gamma = 2^c\gamma'$ with α', γ' odd and $L := 2e - c \geq 1$. Then $S(d) = 2^{-L}(\alpha' u^2 + \gamma') \cdot (\text{odd unit})$ and d is odd exactly when $t = 0$, so the parity condition depends on u only through $(\alpha' u^2 + \gamma')/2^L \bmod 2$, i.e. through $u^2 \bmod 2^{L+1}$. For $L \leq 2$, $u^2 \equiv 1 \pmod{2^{L+1}}$ for every odd u: the condition is constant on the single class $\{v_2(d) = t\}$ modulo 2^{t+1}, at most one class. For $L = 3$, $u^2 \bmod 16$ is 1 or 9 according as $u \equiv \pm 1$ or $\pm 3 \pmod 8$; these values differ by 8 and α' is odd, so the parity of $(\alpha' u^2 + \gamma')/8$ is opposite on the two groups and the parity condition keeps exactly one of the symmetric pairs $\{u \equiv \pm 1\}$, $\{u \equiv \pm 3\} \pmod 8$ (or integrality fails for every u): at most two classes modulo the least period 2^{t+3}. For $L \geq 4$ the unit congruence has exactly the four solutions $\pm u_0$, $\pm u_0 + 2^{L-1} \pmod{2^L}$, so the integral d form the two classes $d \equiv \pm 2^t u_0 \pmod{2^m}$, $m := t + L - 1$. Along either class, $d = D + 2^m j$ gives $S(d) = S(D) + 2^{m+1}ADj + 2^{2m}Aj^2$ with

$$v_2(2^{m+1}AD) = (m+1) + (a-2e) + t = 2t + L + a - 2e = 0, \qquad v_2(2^{2m}A) = 2m + a - 2e = L - 2 \geq 2,$$

so $S(d) - d \equiv S(D) - D + j \pmod 2$ alternates with j on *both* classes, and the parity condition keeps exactly one sub-class modulo 2^{m+1} of each: exactly two classes modulo 2^{m+1}. Translation by 2^m reverses membership, so 2^m is not a period and the local least period is 2^{m+1}; minimality of q gives $e = m + 1 = t + L$. In every case $\mathcal{A}_2$ has at most two classes modulo its least period, and multiplying the local bounds gives $n \leq 2^{\omega(q)}$. □

A cruder treatment of $p = 2$ – bounding the parity refinement of the period by 4 without tracking which sub-classes survive – would allow a factor 4 there; the factor 2 is the truth, because the parity of $S(d) - d$ along the integrality classes has the same slope on both classes,

which is the observation that closes the even case. Two computations agree: over the 16 299 distinct pairs (A, C) of the census box $q \le 300$, $0 < |M| \le 32$, and over the 3 093 220 distinct families with $\alpha \le 200$, $|\gamma| \le 200$, $q \le 40$ (443 944 of them with even least period), the class count of the least period never exceeded $2^{\omega(q)}$. The lemma enters the constants below through $\sup_q 2^{\omega(q)} q^{-1/3} = 16 \cdot 210^{-1/3} = 2.6918\ldots$.

Lemma 20.15 (long periods)**.** *A drift-free family with at least two pairs of cells of the shell $(r/2, r]$ has least positive period $q < r^2/2$, and if $q > r$ it has at most $2^{\omega(q)}$ shell pairs altogether – hence at most $2^{\omega(q)} = r^{o(1)}$ in any window, modulus-free or planar, with no hypothesis on Δ_d.*

Proof. Let $(k_i, k_i + d_i)$, $i = 1, 2$, be two shell pairs of the family, $d_1 \ne d_2$, and $s_i := 2k_i + d_i = Ad_i^2 + C$. Then $A = (s_1 - s_2)/(d_1^2 - d_2^2)$ and $C = s_1 - Ad_1^2$ are rationals whose denominators divide $h := |d_1^2 - d_2^2| < r^2/4$, since $0 < d_i < r/2$. Write $A = a/h$, $C = c/h$ with $a, c \in \mathbb{Z}$. For every integer d, $s(d + 2h) - s(d) = 4a(d + h) \in \mathbb{Z}$ and $k_-(d + 2h) - k_-(d) = 2a(d + h) - h \in \mathbb{Z}$, so the admissible set $\mathcal{A}$ is invariant under translation by $2h$, and its least positive period q divides $2h < r^2/2$. If $q > r$, each of the $n \le 2^{\omega(q)}$ admissible classes (Lemma 20.14) meets the interval $(0, r/2)$ of shell parameters in at most one integer, so the family has at most n shell pairs. Finally $2^{\omega(q)} \le \exp((1 + o(1)) \log 2 \, \log q / \log\log q) = r^{o(1)}$ for $q < r^2/2$. □

Theorem 20.16 (V, the two-progression envelope)**.** *Let $0 < \theta \le 1.2$. For all sufficiently large r, every drift-free two-progression family with $\Delta_d \ne 0$, of any least positive period q, has at most $7\, r^{1/3}$ pairs of cells of the shell $(r/2, r]$ whose speeds lie in one window of half-width $\theta\rho_r$. In particular $D^*_{2P,\theta}(r) \le 7\, r^{1/3}$; hence, for $0 < \theta \le 1$, Conjecture 20.12 implies $D^*_\theta(r) \le (7 + o_\theta(1))\, r^{1/3}$, and by Lemma 20.9 (a window of half-width $1.1\rho_r$ being covered by two of half-width ρ_r, so $D^*_{1.1} \le 2D^*_1$) the $r^{1/3+o(1)}$ part of (S′) in the transfer range $r \le N^{1/5-\varepsilon}$. It does not by itself imply the planar part of (S′).*

Proof. If $q > r$, Lemma 20.15 bounds the family's shell pairs by $2^{\omega(q)} = r^{o(1)} < 7\, r^{1/3}$ for large r (or the family has at most one shell pair); so assume $q \le r$. A member in the shell has $r/2 < k_- < k_+ \le r$, so $S \in (r, 2r)$, $w \in (r, 2r)$, $\sigma := \sqrt{k_+} + \sqrt{k_-} \in (\sqrt{2r}, 2\sqrt{r})$, $d = k_+ - k_- < r/2$ and $v = d/\sigma$. (a) *Population.* If $C \ge 2r$ there are no members in the shell. Otherwise $Ad^2 = S - C \in (r - C, 2r - C)$ confines d to an interval of length at most $\sqrt{(2r - C)/A} - \sqrt{(r - C)^+/A} \le \sqrt{r/A}$ (for $C < r$ the difference equals $r/(\sqrt{2r - C} + \sqrt{r - C}) \le \sqrt{r}$), which contains at most $n(\sqrt{r/A}/q + 1)$ admissible d. (b) *Small A.* If $A < 1/(4r)$ then $Ad^2 < r/16$, and the shell parameters form an interval, since on $0 < d < r/2$ both $S - r - d$ and $2r - S - d$ are decreasing ($(S - r - d)' = 2Ad - 1 < -3/4$). From $(\sigma^2)' = d(2Aw + P)/w$ one has $v' = \sigma^{-1}(1 - d\sigma'/\sigma)$ with $d\sigma'/\sigma = d^2(2Aw + P)/(2w\sigma^2)$, which is ≤ 0 when $2Aw + P \le 0$ and otherwise less than $2A(w + S)d^2/(2w\sigma^2) = Ad^2/w < 1/16$; hence $v' > 15/(32\sqrt{r}) > 2\theta\rho_r$, distinct admissible d have speeds more than a window apart, and the window holds at most one pair. (c) *Capacity.* If $A \ge 1/(4r)$ then $2Aw + P = 2A(w + S) - 1 > 0$, so by Lemma 20.13 v is strictly monotone in d (decreasing for $\Delta_d > 0$, increasing for $\Delta_d < 0$). Let J be the convex hull of the shell parameters whose speeds lie in the window; by monotonicity their extreme speeds differ by less than $2\theta\rho_r$. The bounds $S < 2r$, $w < 2r$, $d < r/2$ and $\sigma > \sqrt{2r}$ hold throughout J: if the selected parameters lie in one component of the shell set this is immediate, and if they lie in two components (which requires $C > r$, k_- being convex in d) then $S - r = C - r + Ad^2 > d^2/(4r)$ on J, which is equivalent to $S + w > 2r$, i.e. $\sigma > \sqrt{2r}$. On J,

$$|v'| = \frac{|(v^2)'|}{2v} = \frac{2|\Delta_d|\, d}{v\, w\, (2Aw + P)} \;\ge\; \frac{2|\Delta_d|\, d}{(d/\sqrt{2r}) \cdot 2r \cdot 8Ar} = \frac{|\Delta_d|}{4\sqrt{2}\, Ar^{3/2}},$$

so $|J| \le 2\theta\rho_r \cdot 4\sqrt{2}\,Ar^{3/2}/|\Delta_d| = 2\theta A/|\Delta_d| = 8\theta Aq^2/|M|$, and J contains at most $n(8\theta Aq/|M|+1)$ admissible d. (d) By (a) and (c) the count is at most $n\big(1+\min\{8\theta Aq/|M|,\ \sqrt{r/A}/q\}\big) \le n\big(1+2\theta^{1/3}q^{-1/3}|M|^{-1/3}r^{1/3}\big)$, the minimum being largest where its two terms agree. Now $|M| \ge 1$, and by Lemma 20.14 $n\,q^{-1/3} \le 2^{\omega(q)}q^{-1/3}$; since q is at least the product P_k of the first $k = \omega(q)$ primes, this is at most $2^k P_k^{-1/3}$, which is maximal at $k = 4$ with the value $16 \cdot 210^{-1/3} = 2.6918\ldots$. Moreover $2^{\omega(q)} \le \max\{256, q^{1/3}\}$ for every q: for $\omega(q) \le 8$ this is clear, and for $k = \omega(q) \ge 9$ one has $q \ge P_k \ge 2^{3k}$, because $P_9 > 2^{27}$ and every further prime exceeds 8. Hence $n \le \max\{256, q^{1/3}\} \le r^{1/3}$ for $q \le r$ and $r \ge 2^{24}$, and the count is at most $r^{1/3} + 2 \cdot 2.6918 \cdot 1.2^{1/3}\,r^{1/3} = 6.73\ldots\,r^{1/3} < 7\,r^{1/3}$. □

The constant is not optimised: Theorem 20.11 gives $D^*_{2P,1}(r) \ge 0.4196\,r^{1/3}$ and the census of Section 23 reaches $0.57\,r^{1/3}$, so the two-progression class has exponent exactly $1/3$; at $\theta = 1$ its constant lies between 0.42 (Theorem 20.11) and the uniform 7 of Theorem 20.16, which the absorption of the additive class term lowers to 6 for large r. The hypothesis $\Delta_d \ne 0$ excludes only the Beatty case, in which $4k_-k_+$ is a square; two distinct squarefree integers cannot have a square product, so such a family contributes no pair of squarefree cells and the bound on $D^*_{2P,\theta}$ is unaffected. Families of least period $q > r$ are covered by Lemma 20.15: they have at most $2^{\omega(q)} = r^{o(1)}$ shell pairs altogether, so $D^*_{2P,\theta}$ may be taken over all drift-free families and the two-progression class has exponent exactly $1/3$ without the restriction $q \le r$.

21. The planar regime

For $N^{1/5} \ll r \le N^{1/3}$ the offsets $k' - k$ exceed the window and the modulus can no longer be removed. We prove two additional effects: the offset caps each symmetric family, and offset resonance – the drift of the speed cancelling the drift of the offset – produces clusters that the modulus-free statistic cannot see; these are not claimed to exhaust planar clustering.

21.1. **The cap.**

Proposition 21.1 (planar cap)**.** *Fix N, r and an odd j, and let $(k_-, k_+) = ((jt^2-t)/2, (jt^2+t)/2)$ run over the symmetric pairs in the shell. The exact start difference $d(t) = \lceil 2\sqrt{k_+N}\rceil - \lceil 2\sqrt{k_-N}\rceil - t$ is non-increasing and satisfies $d(t) - d(t+m) > m-2$, so an integer window of $2W-1$ consecutive integers contains the start differences of at most $2W$ pairs of the family.*

Proof. $|d(t) - f(t)| < 1$, where $f(t) := u(\sqrt{k_+} - \sqrt{k_-}) - t$; the speed decreases with t and the offset decreases by exactly one, so $f(t) - f(t+1) > 1$, whence $d(t+1) - d(t) < 1$, i.e. d is non-increasing (its increments are integers), and $d(t) - d(t+m) > f(t) - f(t+m) - 2 > m-2$. □

The cap falls below the modulus-free law $\asymp r^{1/3}$ exactly when $r \gtrsim 1.4\,N^{3/11}$ – the same point at which the two branches of Theorem 4.4 cross at $\eta = 1/3$, since both compare $r^{1/3}$ with $\sqrt{N}r^{-3/2}$. The cap is per family; the full statistic may collect several families in one window.

21.2. **Offset resonance.**

For a two-progression family $((s-d)/2, (s+d)/2)$, $s = Ad^2 + Bd + C$, write $v(d) := d/(\sqrt{k'} + \sqrt{k})$. The exact start difference is $D(d) = u\,v(d) - d + \vartheta(d)$ with $|\vartheta| < 1$. The unrounded phase $u\,v(d) - d$ is *stationary* where $u\,v'(d_*) = 1$ and, to leading order, varies near d_* by $\frac{1}{2}u|v''(d_*)|(d-d_*)^2$, so, to first order, the members with $|d - d_*| < \sqrt{2W/(u|v''(d_*)|)}$ share one window. For drift-free families $v'(d) = -(2A)^{-1/2}\Delta_d/(A^2d^3) + O(d^{-5})$ by Proposition 20.3, so resonance needs $\Delta_d < 0$, and then $d_*^3 = u|\Delta_d|/(\sqrt{2A}A^2)$ and $v''(d_*) = -3(1 + O(d_*^{-2}))/(ud_*)$. The stationary point lies in the shell iff $|M|/(4\alpha) \approx 1/(\sqrt{2}\,W_{\mathrm{real}})$ (first order), and a primitive family with n classes then has

$$\text{members} \approx 1.18\,n|M|^{-1/2}r^{5/4}N^{-1/4}, \qquad \text{range} \approx 1.94\,n\,q^{-1/2}|M|^{-1/4}N^{1/8}r^{-1/8}$$

(first order), equal at $r \approx 1.44\,N^{3/11}(q^{-2}|M|)^{2/11}$, where both are $\asymp N^{1/11}$; at $r = N^{1/3}$ the range is $\asymp N^{1/12}$. The next three theorems make the mechanism rigorous for specific families.

21.3. **Theorem W.**

Theorem 21.2 (W). *For every $\varepsilon > 0$ there is $N_0(\varepsilon)$ such that for every integer $N \ge N_0(\varepsilon)$, with $r = \lfloor N^{1/3} \rfloor$ (where every $a = 1$ window has $W = 1$),*

$$D_{\max}(\mathcal{F}_1(r), 1) \;\ge\; (0.633 - \varepsilon)\, N^{1/12}.$$

Proof. (a) $W_{1k} = \lceil \sqrt{N}/(4r\sqrt{k}) \rceil = 1$ for every k in the shell when $N \ge 27$, so $D_{\max}(\mathcal{F}_1(r), 1)$ is the largest number of ordered pairs of distinct cells with the same start difference $t \ne 0$. (b) Let $k_{\mp}(d) = (d^2 \mp d + 2)/2$, integers for every d, with k_+ strictly increasing on $d \ge 1$. Put $d_* := (7\sqrt{N}/(2\sqrt{2}))^{1/3} = 1.3526\,N^{1/6}$ and $L := \lambda\sqrt{d_*}$ with $0 < \lambda < \sqrt{2/3}$; every $d \in [d_* - L, d_* + L]$ has $k_{\pm}(d) = 0.9148\,N^{1/3}(1 + O(N^{-1/12}))$, so the pairs lie in the shell for $N \ge N_0$. (c) With $u = 2\sqrt{N}$ the start difference is $D(d) = f(d) + \vartheta(d)$, $f(d) := u\,v(d) - d$, $v(d) := \sqrt{k_+} - \sqrt{k_-}$, $|\vartheta| < 1$. From $k_+ + k_- = d^2 + 2$ and $4k_+k_- = d^4 + 3d^2 + 4$, $v(d)^2 = d^2 + 2 - \sqrt{d^4 + 3d^2 + 4} = \frac{1}{2} - \frac{7}{8}d^{-2} + O(d^{-4})$, so $v(d) = 2^{-1/2}(1 - \frac{7}{8}d^{-2} + \frac{119}{128}d^{-4} + O(d^{-6}))$, a convergent expansion for $d \ge 2$ that may be differentiated termwise. With $A := 7u/(4\sqrt{2})$, $f'(d) = Ad^{-3}(1 - \frac{17}{8}d^{-2} + O(d^{-4})) - 1$ and $f''(d) = -3Ad^{-4}(1 - \frac{85}{24}d^{-2} + O(d^{-4}))$. The nominal point $d_* = A^{1/3}$ is not the critical point: $f'(d_*) = -\frac{17}{8}d_*^{-2} + O(d_*^{-4})$, and the unique zero $\hat{d}$ of f' near it is $\hat{d} = d_* - 17/(24d_*) + O(d_*^{-3})$, at distance $O(d_*^{-1}) \ll L$, with $f''(d) = -(3/d_*)(1 + O(L/d_*))$ uniformly on the window. Taylor at $\hat{d}$ gives, for $|d - d_*| \le L$, $0 \le f(\hat{d}) - f(d) \le \frac{3}{2d_*}(1 + o(1))L^2$, so the values $f(d)$ on the window lie in an interval of length $V \le \frac{3}{2}\lambda^2(1 + O(N^{-1/12})) < 1$ for $N \ge N_0(\lambda)$. (d) The integers $D(d)$, $|d - d_*| \le L$, lie in an open interval of length $V + 2 < 3$ and take at most three values; the parameter interval contains $M \ge 2L - 1$ integers d, so some t is the start difference of at least $\lceil M/3 \rceil \ge (2\lambda\sqrt{d_*} - 1)/3$ distinct ordered pairs. Finally $t \ne 0$: $\sqrt{k_+} + \sqrt{k_-} \le \sqrt{2(k_+ + k_-)} \le 2d$ for $d \ge 2$, so $v \ge 1/2$ and $f(d) \ge \sqrt{N} - 2N^{1/6} > 1$ on the window. Thus $D_{\max} \ge \frac{2\lambda}{3}\sqrt{d_*} - \frac{1}{3} = (\frac{2\lambda}{3} \cdot 1.1630 - o(1))N^{1/12}$, and $\lambda \uparrow \sqrt{2/3}$ gives 0.633. □

No primality of N is used. Lemma 4.3 applied to the full $a = 1$ shell at $r \ge N^{1/3}$ has Sidon branch $(2^{-1/3} + o(1))r^{2/3}D^{-1/3}$, which with $D \ge 0.633N^{1/12}$ is at most $0.925\,N^{7/36} = N^{0.194}$: the full $a = 1$ shell cannot certify more than this through Lemma 4.3 for $r \ge N^{1/3}$. The squarefree version of Theorem 21.2 would need a positive proportion of d in an interval of length $\asymp d_*^{1/2}$ about $d_* \asymp N^{1/6}$ with both $(d^2 \mp d + 2)/2$ squarefree; no positive-proportion result for squarefree values of a fixed irreducible quadratic in intervals $[x, x + x^{\vartheta}]$ is known, to our knowledge, for any $\vartheta < 1$ (the known results [22, 41, 63, 34] are global), so that statement remains open for every individual N. Averaged over N it is provable, because the family is N-independent and only the window moves with N.

Theorem 21.3 (W_{sf}, almost all N). *Let $g(d) := (d^2 - d + 2)/2$, so that the cells of the family of Theorem 21.2 are $(g(d), g(d+1))$, and let*

$$\kappa_g := \frac{1}{2} \prod_{\substack{p \ge 11 \\ p \equiv 1,2,4 \pmod 7}} \left(1 - \frac{4}{p^2}\right) = 0.4715\ldots,$$

the density of the integers d for which $g(d)$ and $g(d+1)$ are both squarefree. For every $\lambda \in (0, \sqrt{2/3})$ and $\varepsilon \in (0, \kappa_g)$ there are $C = C(\lambda, \varepsilon)$ and X_0 such that for every $X \ge X_0$, all but at

most $CX/\log X$ of the integers $N \in [X, 2X]$ satisfy

$$D_{\max}\big(\mathcal{F}_1^{\square}(\lfloor N^{1/3}\rfloor), 1\big) \;\ge\; (\kappa_g - \varepsilon)\,\frac{2\lambda}{3}\Big(\frac{7}{2\sqrt{2}}\Big)^{1/6} N^{1/12} - 1.$$

In particular, for every $\varepsilon > 0$ the bound $(0.298-\varepsilon)N^{1/12}$ holds for all N outside a set of relative density $O_\varepsilon(1/\log X)$ in $[X, 2X]$.

Proof. (a) *Local densities.* Put $G(d) := 2g(d) = d^2 - d + 2$, of discriminant -7. For an odd prime $p \neq 7$, $p^2 \mid g(d)$ iff $G(d) \equiv 0 \pmod{p^2}$, which has exactly two solutions modulo p^2 if $(-7/p) = 1$ (simple roots lift by Hensel's lemma) and none otherwise; for $p = 7$ the double root $d \equiv 4$ does not lift, since $G(4 + 7t) \equiv 14 \pmod{49}$. For odd p the roots of $G(d)$ and of $G(d+1)$ modulo p are disjoint, because $G(d+1) - G(d) = 2d$ and $p \mid d$ forces $G(d) \equiv 2 \pmod{p}$. Finally $4 \mid g(d)$ iff $d \equiv 3, 6 \pmod 8$. Hence the density of d with $g(d)$ and $g(d+1)$ both squarefree is κ_g as displayed, with 4 bad residues modulo p^2 for $p \equiv 1, 2, 4 \pmod 7$, $p \ge 11$, none for the other odd primes, and 4 bad residues modulo 8. That κ_g is the natural density of this set follows from (b) and (c) below applied to the whole range $[0, 3D]$: the periodic set gives upper density $\le \kappa_z$ and the large-prime count gives lower density $\ge \kappa_z - O(1/z + 1/\log D)$; let $D \to \infty$ and then $z \to \infty$. Numerically the product to 10^6 is 0.47154 (tail below 10^{-6}), and the empirical density over $d \le 4 \cdot 10^5$ is 0.47146. (b) *Small primes are periodic.* For $z \ge 3$ let $P_z := 8\prod_{3\le p\le z} p^2$ and let $\mathcal{G}_z$ be the set of d with $4 \nmid g(d)$, $4 \nmid g(d+1)$ and $p^2 \nmid g(d)$, $p^2 \nmid g(d+1)$ for $3 \le p \le z$. It is periodic modulo P_z with density $\kappa_z := \frac{1}{2}\prod_{3\le p\le z}(1 - \rho(p)/p^2) \ge \kappa_g$, $\rho(p) \in \{0, 4\}$ as in (a), so every interval of M consecutive integers contains at least $\kappa_z M - P_z$ elements of $\mathcal{G}_z$. (c) *Large primes on average.* Let $D \ge L \ge 1$, $M := 2L + 1$, and for $c \in [D, 2D]$ let $I_c := [c - L, c + L]$ and $T(c) := \#\{d \in I_c : p^2 \mid g(d + s) \text{ for some prime } p > z \text{ and } s \in \{0, 1\}\}$. Every such d lies in $[0, 3D]$ and has $g(d + s) < (3D + 2)^2$, so only $p \le 3D + 1$ occur. For fixed s and p the $d \in [0, 3D]$ with $p^2 \mid g(d + s)$ lie in at most two residue classes modulo p^2, so they number at most $2(3D/p^2 + 1)$, and at most 2 when $p^2 > 3D$. Hence, for $D \ge D_0$,

$$\begin{aligned}&\#\{d \in [0, 3D] : p^2 \mid g(d + s) \text{ for some } p > z,\ s \in \{0, 1\}\}\\ &\quad\le\; 2\Big[\sum_{z<p\le\sqrt{3D}} 2\Big(\frac{3D}{p^2} + 1\Big) + 2\pi(3D + 1)\Big] \;\le\; \frac{12D}{z} + \frac{16D}{\log D},\end{aligned}$$

using $\sum_{p>z} p^{-2} < 1/z$ and $\pi(y) \le 1.26\, y/\log y$. Each such d is counted in $T(c)$ for at most M values of c, so $\sum_{D\le c\le 2D} T(c) \le M(12D/z + 16D/\log D)$, and by Markov's inequality

$$\#\{c \in [D, 2D] : T(c) \ge \varepsilon M\} \;\le\; \frac{D + 1}{\varepsilon}\Big(\frac{12}{z} + \frac{16}{\log D}\Big).$$

(d) *Combination.* Let $S(c) := \#\{d \in I_c : g(d), g(d + 1) \text{ squarefree}\}$. Then $S(c) \ge \#(\mathcal{G}_z \cap I_c) - T(c)$, so $S(c) \ge (\kappa_g - \varepsilon)M - P_z$ for every $c \in [D, 2D]$ outside the exceptional set of (c). Take $z := \lfloor(\log L)/5\rfloor$: since $\sum_{p\le z} \log p < 1.02z$, $P_z \le \sqrt{L}$ for $L \ge L_0$, and $12/z \le 61/\log L$. (e) *Transfer to N.* Put $d_*(N) := (7\sqrt{N}/(2\sqrt{2}))^{1/3}$, fix $\delta > 0$ so small that $\lambda' := \lambda(1+\delta)^{1/12} < \sqrt{2/3}$, and cover $[X, 2X]$ by $\lceil 1/\delta\rceil$ intervals $[Y, (1+\delta)Y]$. On such an interval put $L := \lfloor \lambda'\sqrt{d_*(Y)}\rfloor - 1$ and, for $N \in [Y, (1 + \delta)Y]$, let $c(N)$ be the integer nearest to $d_*(N)$. Then $I_{c(N)} \subset [d_*(N) - \lambda'\sqrt{d_*(N)}, d_*(N) + \lambda'\sqrt{d_*(N)}]$, so by steps (b)–(d) of the proof of Theorem 21.2 (with λ' in place of λ) the start differences $D(d)$, $d \in I_{c(N)}$, take at most three values, none zero, for $N \ge N_0(\lambda')$, and every $d \in I_{c(N)}$ with $g(d), g(d+1)$ squarefree gives an ordered pair $(g(d), g(d+1))$ of distinct squarefree cells of the shell. Hence $D_{\max}(\mathcal{F}_1^{\square}(\lfloor N^{1/3}\rfloor), 1) \ge S(c(N))/3$. As N runs through $[Y, (1+\delta)Y]$, $c(N)$ runs monotonically through an interval contained in $[D, 2D]$ with $D := c(Y) \ge L$, and each value c is taken by at most $8Y^{5/6}$ integers N, since $dN/dd_* = 6N^{5/6}/1.3526$. By

(c)–(d) the exceptional N in the interval number at most $8Y^{5/6}(D+1)(61/\log L+16/\log D)/\varepsilon = O(Y/(\varepsilon\log Y))$, and every other N has $D_{\max} \ge ((\kappa_g-\varepsilon)(2L+1)-\sqrt{L})/3 \ge (\kappa_g-2\varepsilon)\frac{2\lambda}{3}\sqrt{d_*(N)}-1$ for N large, because $2L+1 = 2\lfloor\lambda'\sqrt{d_*(Y)}\rfloor - 1 \ge 2\lambda'\sqrt{d_*(Y)} - 3 \ge 2\lambda\sqrt{d_*(N)} - 3$, and this $O(1)$ loss and the term $\sqrt{L}/3 = o(\sqrt{d_*(N)})$ are absorbed in the ε-margin. Summing over the $\lceil 1/\delta\rceil$ intervals and renaming 2ε as ε gives the theorem, and $\sqrt{d_*(N)} = (7/(2\sqrt{2}))^{1/6}N^{1/12}$. Finally $\kappa_g\cdot\frac{2}{3}\sqrt{2/3}\cdot(7/(2\sqrt{2}))^{1/6} = 0.4715\cdot 0.6330 = 0.2985 > 0.298$. □

Two consequences. For all N outside the exceptional set, Lemma 4.3 applied to the squarefree $a=1$ shell at $r \ge N^{1/3}$ has Sidon branch $(2^{1/3}(3/\pi^2)^{2/3}+o(1))u^{2/3}D^{-1/3} = (0.5696 + o(1))u^{2/3}D^{-1/3} \le (0.86+o(1))N^{7/36}$; and the squarefree resonance lower bound at the top of the range, which for every individual N rests on the open short-interval problem above, is unconditional for almost all N – a statement incompatible with any $r^{o(1)}$ planar envelope for the squarefree shell, and consistent with (P′) of Section 22. The same averaging applies verbatim to the families $nt^2 \mp mt + 1$ of Theorem 21.4, since they too are N-independent; the constants are omitted.

21.4. **Theorem W′.**

Theorem 21.4 (W′). *For every $C \ge 1$ there are $c(C) > 0$ and $N_0(C)$ such that for every integer $N \ge N_0(C)$ and every integer r with $\lceil N^{1/3}/C\rceil \le r \le \lfloor N^{1/3}\rfloor$,*

$$D_{\max}(\mathcal{F}_1(r), W) \;\ge\; c(C)\, N^{1/12}.$$

Proof. Put $w := W_{\mathrm{real}}(r) \in [2^{-3/2}, 2^{-3/2}C^{3/2}]$, so W is bounded. For integers $n, m \ge 1$ with $\delta := 4n - m^2 > 0$ consider $k_\mp(t) = nt^2 \mp mt + 1$ (integers; k_+ strictly increasing on $t \ge 0$; $k_+ - k_- = 2mt$). From $k_+ + k_- = 2nt^2 + 2$ and $4k_+k_- = 4(nt^2+1)^2 - 4m^2t^2$,

$$v(t) := \sqrt{k_+} - \sqrt{k_-} = \frac{m}{\sqrt{n}}\Big(1 - \frac{\delta}{8n^2t^2} + O(t^{-4})\Big),$$

uniformly for bounded n, m; the phase $f(t) := u\,v(t) - 2mt$ has $f'(t) = \frac{um\delta}{4n^{5/2}}t^{-3}(1+O(t^{-2})) - 2m$ and $f''(t) = -\frac{3um\delta}{4n^{5/2}}t^{-4}(1+O(t^{-2}))$, so its critical point is $\hat{t} = t_*(1+O(t_*^{-2}))$ with $t_*^3 := u\delta/(8n^{5/2}) \asymp N^{1/2}$, i.e. $t_* \asymp N^{1/6}$, and $f''(\hat{t}) = -(6m/t_*)(1+o(1))$. Writing $z := nt_*^2$, $(z/r)^{3/2} = w\delta/(\sqrt{2}\,n)$ exactly, so the nominal point z lies in the shell iff $\rho := \delta/n = 4 - m^2/n$ lies in $J_w := (1/(2w), \sqrt{2}/w)$, strictly inside with a margin when ρ lies in the middle third of J_w ($z/r \in (0.686, 0.850)$); the actual members have $k_\pm(t)/r = z/r + O(t_*^{-1/2})$ on the window below and so lie in the shell for N large. The set $\{4 - m^2/n\}$ is dense in $(0,4)$, and $J_w \subset (0,4]$ for $w \ge 2^{-3/2}$, so every w in the compact range admits a pair with ρ in the open middle third of J_w. Each pair serves an open set of values of w, so finitely many pairs $S(C)$ cover the range with a uniform margin; all O-terms are uniform over $S(C)$, which fixes $N_0(C)$. For r served by (n,m), the argument of Theorem 21.2 applies: over $|t - t_*| \le \lambda\sqrt{t_*/(3m)}$, $\lambda < 1$, the phase varies by $\lambda^2(1+o(1)) < 1$; the integers $D(t)$ take at most three values; one value is shared by $\ge \lceil M/3\rceil \ge (2\lambda\sqrt{t_*/(3m)} - 1)/3$ distinct ordered pairs; and that value exceeds $W-1$, since $D(t) = (2m/\sqrt{n})\sqrt{N} + O(N^{1/6})$ while $W = O_C(1)$, so the window containing it is not centred at 0. Since $t_* \ge c_1(C)N^{1/6}$ and $m \le \max S(C)$, this is $\ge c(C)N^{1/12}$. □

At the endpoints $r \asymp N^{1/3}/C$ the pairs $(n,m) = (1,1), (3,3), (7,5), (13,7)$ serve for $C = 1,2,3,4$ respectively (Section 23); $(1,1)$ is the family $(m^2 - m + 1, m^2 + m + 1)$ that the exact census returns as the maximiser at $r = N^{1/3}$ on every modulus tested.

21.5. **Theorem X.**

Theorem 21.5 (X). *Let $r = r(N)$ satisfy $N^{1/5} \le r \le N^{1/3}$ and $W_{\mathrm{real}} \to \infty$, let $W = W_{\mathrm{real}}(1+o(1))$ be the window width, fix $x > 0$, let $j = j(N)$ be odd with $j = x\,r^{1/3}(1+o(1))$, and put $L := |I_j(r)| = (\sqrt{2}-1)x^{-1/2}r^{1/3}(1+o(1))$. Then*

$$D_{\max}(\mathcal{F}_1(r), W) \;\ge\; \frac{L(2W-1)}{V+2W-1},$$
$$V \;:=\; \frac{W_{\mathrm{real}}(1+2/(jr))r^{1/2}}{4j^{3/2}} + L + 1 = \frac{W_{\mathrm{real}}}{4x^{3/2}}(1+o(1)) + L + 1,$$

and $D_{\max}(\mathcal{F}_1^{\square}(r), W) \ge \kappa_j L(1-o(1))\,(2W-1)/(V+2W-1)$.

Proof. The pairs $(k_-, k_+) = ((jt^2-t)/2, (jt^2+t)/2)$, $t \in I_j(r)$, are L distinct ordered pairs of cells with start differences $d(t) = u\,s(t) - t + \vartheta(t)$, $|\vartheta| < 1$. By Lemma 20.7 the speeds $s(t)$ lie in an interval of length at most $(1+2/(jr))/(16\sqrt{2}\,j^{3/2}r)$; multiplying this length by $u = 2\sqrt{N} = 4\sqrt{2}\,r^{3/2}W_{\mathrm{real}}$ gives $S := W_{\mathrm{real}}(1+2/(jr))r^{1/2}/(4j^{3/2})$; the offsets $-t$ span $L-1$ and the ϑ lie in $(-1,1)$, so the L integers $d(t)$ lie in an open interval of length $V = S+L+1$, hence in a set J of at most $V+1$ consecutive integers. There are at most $V+2W-1$ windows of $2W-1$ consecutive integers that meet J, and every integer of J lies in exactly $2W-1$ of them; summing the counts gives $(2W-1)L$, so some window contains at least $(2W-1)L/(V+2W-1)$ of the $d(t)$. Its centre t_0 is nonzero: $s(t) \ge t/(2\sqrt{r})$ gives $d(t) > \sqrt{N/r} - 2$, while $W \le W_{\mathrm{real}} + 1$ and $\sqrt{N/r}/W_{\mathrm{real}} = 2\sqrt{2}\,r$, so $d(t) > W-1$ for N large and no window containing a $d(t)$ is centred at 0. For the squarefree shell apply Lemma 20.6 on $I_j(r)$ with $z = \ln T_2$, which leaves $\kappa_j L - O(L/\ln L)$ values of t with both members squarefree, and repeat the count. □

Corollary 21.6. *(a) If $r = o(N^{3/11})$ then, with $y = x^{-1/2}$, the bound is $(\sqrt{2}-1)r^{1/3}y/(1+y^3/8)(1-o(1))$, maximal at $x^{3/2} = 1/4$, where $D_{\max}(\mathcal{F}_1(r), W) \ge (0.438 - o(1))r^{1/3}$ and $D_{\max}(\mathcal{F}_1^{\square}(r), W) \ge (0.137 - o(1))r^{1/3}$ (0.200 with $15 \mid j$) – Theorem 20.8 extended past the transfer range with a weaker constant. (b) At $r = \beta N^{3/11}$, taking $x = 0.4$ gives full-shell bounds $\approx (0.296, 0.271, 0.224)N^{1/11}$ for $\beta = 0.5, 1, 1.44$ and κ_0 times these for the squarefree shell; in particular $D_{\max}(\mathcal{F}_1) \ge (0.27 - o(1))N^{1/11}$ and $D_{\max}(\mathcal{F}_1^{\square}) \ge (0.0848 - o(1))N^{1/11}$ at $r = N^{3/11}$. (c) For $N^{3/11} \ll r$ with $W_{\mathrm{real}} \to \infty$: the bound is $(2W-1)(1-O(W/L))$, so the symmetric family attains its cap up to $1 - o(1)$.*

21.6. **Corollary Y.** For a sub-family $\mathcal{F}$ with n' tested integers in total write

$$F_{\mathcal{F}}(N, r) := \max\Big(2^{-1/2}N^{1/4}r^{-1/4},\ \min\big(n'/(2W),\ 2(n'^2/(4W\,D_{\max}(\mathcal{F}, W)))^{1/3}\big)\Big)$$

for the lower bound that Theorem 4.4 delivers at parameter r through $\mathcal{F}$.

Corollary 21.7 (Y). *If $\mathcal{F}$ is the full or the squarefree $a = 1$ shell, then $\inf_{r \in \mathbb{Z}_{\ge 1}} F_{\mathcal{F}}(N, r) \le (0.92 + o(1))N^{2/11}$.*

Proof. At $r = \lfloor N^{3/11} \rfloor$ the counting term is $(2^{-1/2} + o(1))N^{2/11}$ and the Sidon branch is $\le 2(R^2W/(4D))^{1/3}$ since $n' \le RW$. For the full shell $R = (\frac{1}{2} + o(1))r$, $W = (2^{-3/2} + o(1))N^{1/11}$ and $D \ge (0.27 - o(1))N^{1/11}$ by Corollary 21.6(b), so the branch is $\le 2(\frac{1}{4} \cdot 2^{-3/2}/(4 \cdot 0.27))^{1/3}N^{2/11} = (0.87 + o(1))N^{2/11}$; for the squarefree shell $R = (3/\pi^2 + o(1))r$ and $D \ge (0.0848 - o(1))N^{1/11}$ give $(0.92 + o(1))N^{2/11}$. □

Theorem 4.4 proves cost $\ge F_{\mathcal{F}}(N, r)$ pointwise, so an algorithm free to choose r is bounded below by $\inf_r F_{\mathcal{F}}$; Corollary 21.7 says that this certificate is at most $0.92\,N^{2/11}$ through the two whole-shell sub-families. It is a ceiling on what the method certifies, not an upper bound on

the cost of any representation, and it says nothing about thinner $a = 1$ sub-families or stronger lemmas; Corollary 21.8 below extends it to the shells with $a \geq 2$ and to the entire family. Within that scope, any in-model bound above $N^{2/11}$ through Lemma 4.3 must come from a thinner sub-family. Theorems 20.8, 21.2, 21.4 and 21.5 together give proven lower bounds for the full $a = 1$ shell along every sequence $r(N) \in [N^{1/5}, N^{1/3}]$ (after passage to a subsequence, W_{real} is either bounded or tends to infinity): $\gg r^{1/3}$ up to the crossing, $\gg W$ beyond it, $\gg N^{1/12}$ where W is bounded. For the squarefree shell, Theorems 20.8 and 21.5 cover the transfer range and the regime $W_{\text{real}} \to \infty$; no squarefree analogue of Theorems 21.2 or 21.4 is proved for every individual N where W is bounded (Theorem 21.3 gives one for almost all N). The matching upper bounds are the conjectures of Section 22.

21.7. **The whole family and the shells with $a \geq 2$.** Corollary 21.7 concerns the two $a = 1$ shells. The ceiling extends to the entire family and to every fixed-a shell. For a sub-family $\mathcal{F}$ write $n'_{\mathcal{F}}$ for its total number of tested integers, and take the counting term of $F_{\mathcal{F}}$ to be $2\sqrt{n'_{\mathcal{F}}}$, the bound $|B||G| \geq n'_{\mathcal{F}}$ gives.

Corollary 21.8 (Y′). *(a) Let $\mathcal{F}_{\text{all}}(r) := \{(a, b) : ab \leq r\}$ be the entire Lehman–Harvey family, with $W := \max_{ab \leq r} W_{ab} = \lceil \sqrt{N}/(4r) \rceil$ and $D := D_{\max}(\mathcal{F}_{\text{all}}(r), W)$. At $r = \lfloor N^{3/11} \rfloor$ the first branch of Lemma 4.3 is $O(N^{3/22} \log N)$ and the second is $O(N^{3/22} (\log N)^{2/3})$, so their minimum is $o(N^{2/11})$, the lemma certifies nothing beyond the counting bound, and*

$$\inf_{r \in \mathbb{Z}_{\geq 1}} F_{\mathcal{F}_{\text{all}}}(N, r) \;\leq\; \big(2\sqrt{3/22} + o(1)\big) N^{2/11} (\log N)^{1/2} \;\leq\; (0.74 + o(1)) N^{2/11} (\log N)^{1/2}.$$

(b) For every fixed $a \geq 1$ let $\mathcal{F}_a(r) := \{(a, b) : r/(2a) < b \leq r/a\}$. There is an explicit C_a such that $\inf_{r \in \mathbb{Z}_{\geq 1}} F_{\mathcal{F}_a}(N, r) \leq C_a N^{2/11}$ for all large N.

Proof. (a) The windows of distinct cells are disjoint: for different a the exponents $z = n - aN - b$ (n tested) lie in disjoint ranges of length $< N/2$, and for fixed a the starts increase with b in steps $\geq \sqrt{N/r} - 3 > W$. Since $W_{ab} \leq \sqrt{N}/(4r\sqrt{ab}) + 1$,

$$n'_{\mathcal{F}_{\text{all}}} \;\leq\; \frac{\sqrt{N}}{4r} \sum_{n \leq r} \frac{d(n)}{\sqrt{n}} + \sum_{n \leq r} d(n) \;=\; \Big(\frac{1}{2} + o(1)\Big) \frac{\sqrt{N} \log r}{\sqrt{r}} + O(r \log r),$$

because $\sum_{n \leq r} d(n) = r \log r + O(r)$ gives $\sum_{n \leq r} d(n) n^{-1/2} = (2 + o(1)) \sqrt{r} \log r$ by partial summation. At $r = \lfloor N^{3/11} \rfloor$ this is $n' \leq (\frac{3}{22} + o(1)) N^{4/11} \log N$, while $W = (\frac{1}{4} + o(1)) N^{5/22}$ and, by monotonicity of $D_{\max}$ in the family and in the window together with Corollary 21.6(b), $D \geq D_{\max}(\mathcal{F}_1(r), W(N, r)) \geq (0.27 - o(1)) N^{1/11}$. Hence $n'/(2W) \leq (\frac{3}{11} + o(1)) N^{3/22} \log N$ and $2(n'^2/(4WD))^{1/3} \leq (0.82 + o(1)) N^{3/22} (\log N)^{2/3}$, both $o(N^{2/11})$, so $F_{\mathcal{F}_{\text{all}}}(N, r) = 2\sqrt{n'} \leq (2\sqrt{3/22} + o(1)) N^{2/11} (\log N)^{1/2}$ at this r. (b) The start of the cell (a, b) is $\lceil 2\sqrt{b \cdot aN} \rceil - aN - b$, which is the start of the cell $(1, b)$ for the modulus $N' := aN$; so the starts of $\mathcal{F}_a(r)$ are those of the shell $\mathcal{F}_1(\rho)$ of N' at $\rho := r/a$, and the theorems of this section apply to them, since they use only $u' = 2\sqrt{N'}$ and not the primality of N'. The widths $W_{ab} = \lceil \sqrt{N}/(4r\sqrt{ab}) \rceil$ are all at most $W(N, r)$, since $ab > r/2$ on the shell, and we use $W := W(N, r)$ as the common window bound in Lemma 4.3; $\mathcal{F}_a(r)$ has $R = r/(2a) + O(1)$ windows. The proof of Theorem 21.5 for (N', ρ) places the $L = (\sqrt{2} - 1) x^{-1/2} \rho^{1/3} (1 + o(1))$ start differences of the symmetric family in at most $V' + 1$ consecutive integers, $V' = W_{\text{real}}(N', \rho)(1 + o(1))/(4x^{3/2}) + L + 1$ with $W_{\text{real}}(N', \rho) = a^2 W_{\text{real}}(N, r)$, and the sliding-window count with the window $2W - 1$, $W = W(N, r)$, gives $D_{\max}(\mathcal{F}_a(r), W) \geq L(2W - 1)/(V' + 2W - 1)$; the centre is nonzero since $d(t) > \sqrt{N'/\rho} - 2 > W - 1$. At $r = \lfloor N^{3/11} \rfloor$ every quantity here is $\asymp N^{1/11}$ with constants

depending on a and x, so $D_{\max}(\mathcal{F}_a(r), W) \geq (c_a - o(1))N^{1/11}$ with, for the choice $x = 1$,

$$w := 2^{-3/2}, \qquad \ell_a := (\sqrt{2}-1)a^{-1/3}, \qquad c_a := \frac{2w\ell_a}{a^2 w/4 + \ell_a + 2w} > 0.$$

Since $R = (1/(2a) + o(1))N^{3/11}$ and $W = (w + o(1))N^{1/11}$, the Sidon branch $2(R^2W/(4D))^{1/3}$ is at most $(2(w/(16a^2c_a))^{1/3} + o(1))N^{2/11}$ and the counting term $2\sqrt{RW}$ at most $(2\sqrt{w/(2a)} + o(1))N^{2/11}$, so $C_a := 1 + \max\{2\sqrt{w/(2a)},\ 2(w/(16a^2c_a))^{1/3}\}$ serves. □

Part (a) says that Lemma 4.3 with a uniform window is weak on the whole family, whose widths range over a factor $\sqrt{r}$; a variable-width refinement of the lemma is the natural next step, and it would be bounded in turn by the shells it contains. Part (b) says the phenomenon is not special to $a = 1$: every a-shell is an $a = 1$ shell in disguise.

21.8. **Resonance at general W (first order).** The following is a leading-order computation, not a theorem. For $N^{1/4} \ll r = o(N^{1/3})$ (so $W_{\text{real}} \to \infty$ and $W_{\text{real}} \ll \sqrt{r}$) let $m = \lceil 7.3\sqrt{W_{\text{real}}} \rceil$, choose $\delta \equiv -m^2 \pmod 4$ in the middle third of $\big(m^2/(8W_{\text{real}} - 1),\ 1.41\,m^2/(4W_{\text{real}} - 1.41)\big]$ (that third has length $> 0.0758\,m^2/W_{\text{real}} > 4$, so such δ exists) and put $n = (m^2 + \delta)/4$. The family $k_{\mp}(t) = nt^2 \mp mt + 1$ then has its stationary point $t_*^3 = u\delta/(8n^{5/2})$ interior to the shell $(nt_*^2/r \in (0.63, 0.90))$, phase curvature $f''(t_*) = -6m/t_*$, and the members with $|t - t_*| \leq \sqrt{(W-2)t_*/(3m)}$ share one window – about $2\sqrt{Wt_*/(3m)} \approx 0.2\,r^{1/4}$ of them, subject to the family's $\asymp \sqrt{r/W_{\text{real}}}$ members in the shell, so the safe first-order profile is $\min(\sqrt{r/W_{\text{real}}}, c\,r^{1/4})$. This exceeds the symmetric family's $2W$ once $r \gtrsim N^{2/7}$; the estimates are those of Theorem 21.4 with (n, m) growing with W_{real}, and a proof would require uniform control of the expansions in that regime.

21.9. **The planar two-progression envelope.** The machinery of Section 20 also bounds the exact-start cluster of a single family, in exactly the functional form of Conjecture (P′). Throughout, $W = W(N, r)$ is Lemma 4.3's window on the shell and $W_{\text{real}} = \sqrt{N}/(2\sqrt{2}\,r^{3/2})$, so that $2^{-3/2} \leq W_{\text{real}}$ for $r \leq N^{1/3}$, $W \leq W_{\text{real}} + 1$ for N large, and hence $W/W_{\text{real}} \leq 2.83$.

Lemma 21.9 (log-derivative of the speed)**.** *With the notation of Lemma 20.13 and $\sigma := \sqrt{k_+} + \sqrt{k_-}$, put $\varphi(d) := \sigma/(w(2Aw + P))$, so that $v' = -2\Delta_d\,\varphi$ wherever $2Aw + P \neq 0$. Then*

$$\frac{\varphi'}{\varphi} = -\frac{d}{w}\Big(A + \frac{P}{w} + \frac{1}{2(S+w)}\Big).$$

On the shell $(r/2, r]$ with $A \geq 1/(2r)$ one has $P \geq 0$, φ is strictly decreasing, $Ad/w \leq |\varphi'/\varphi| \leq 5.5\,Ad/w$, and

$$\frac{1}{8\sqrt{2}\,Ar^{3/2}} \ \leq\ \varphi \ \leq\ \frac{1}{Ar^{3/2}}.$$

Proof. Since $v = d/\sigma$ and $(v^2)' = -4\Delta_d d/(w(2Aw + P))$, $v' = (v^2)'/(2v) = -2\Delta_d\varphi$. From $\sigma^2 = S + w$, $w' = dP/w$ and $(2Aw+P)' = 2Ad(2Aw+P)/w$ one gets $\sigma'/\sigma = d(2Aw+P)/(2w(S+w))$, $w'/w = dP/w^2$ and $(2Aw + P)'/(2Aw + P) = 2Ad/w$, and $\varphi'/\varphi = \sigma'/\sigma - w'/w - (2Aw + P)'/(2Aw + P) = (d/w)\big[(2Aw + P)/(2(S + w)) - P/w - 2A\big]$, where $(2Aw + P)/(2(S + w)) - 2A = -(2AS + 2Aw + 1)/(2(S + w)) = -A - 1/(2(S + w))$. On the shell $S, w \in (r, 2r)$ and $\sigma \in (\sqrt{2r}, 2\sqrt{r})$; for $A \geq 1/(2r)$, $P = 2AS - 1 \geq 2Ar - 1 \geq 0$, so every term of the bracket is non-negative and at least A, while $P/w \leq 2AS/w < 4A$ and $1/(2(S + w)) < 1/(4r) \leq A/2$. Finally $2Aw+P = 2A(w+S)-1$ lies in $(2Ar, 8Ar)$, since $4Ar-1 \geq 2Ar$, so $\varphi > \sqrt{2r}/(2r\cdot 8Ar) = 1/(8\sqrt{2}\,Ar^{3/2})$ and $\varphi < 2\sqrt{r}/(r \cdot 2Ar) = 1/(Ar^{3/2})$. If the set of real $d > 0$ with $(k_-(d), k_+(d))$ in the shell has two components, then $C > r$, and for d between them $k_+(d) \leq r$, $d < r/2$ and $S - r = C - r + Ad^2 > d^2/(2r)$, so $w^2 - r^2 = (S - r)(S + r) - d^2 > 0$: thus $S, w \in (r, 2r)$ and

$\sigma \in (\sqrt{2r}, 2\sqrt{r})$ across the gap as well, and every estimate above, in particular $\varphi' < 0$, holds on the convex hull of the shell parameters. □

Theorem 21.10 ($\mathrm{V_P}$, the planar two-progression envelope)**.** *Let $N^{1/5} \le r \le N^{1/3}$. For all sufficiently large N, every drift-free two-progression family with $\Delta_d \neq 0$, of any least positive period q, has at most*

$$2^{\omega(q)}\Big(4 + 26\min\{W, r^{1/3}\} + 20\,N^{1/8}r^{-1/8}q^{-1/2}\Big)$$

pairs of cells of the shell $(r/2, r]$ whose start differences lie in one window of $2W-1$ consecutive integers.

Proof. If $q > r$, Lemma 20.15 bounds the family's entire shell population by $2^{\omega(q)}$, below the right-hand side; so assume $q \le r$. Write $n \le 2^{\omega(q)}$ for the number of classes (Lemma 20.14). The start difference of the pair $(k_-(d), k_+(d))$ is $D(d) = \lceil u\sqrt{k_+}\rceil - \lceil u\sqrt{k_-}\rceil - d = f(d) + \vartheta(d)$ with $f(d) := u\,v(d) - d$ and $|\vartheta| < 1$, so $D(d) \in [t - W + 1, t + W - 1]$ forces $f(d) \in I_t := (t - W, t + W)$, an interval of length $2W$. We bound the number of admissible d in the shell with $f(d) \in I_t$; by (a) of Theorem 20.16 the admissible d of any real interval J number at most $\min\{|J| + 1,\ n(|J|/q + 1)\}$, and the shell parameters lie in an interval of length $\le \sqrt{r/A}$ (the *population* bound). Let H be the convex hull of the real shell parameters. If $A < 1/(2r)$, then on $0 < d < r/2$ both $S - r - d$ and $2r - S - d$ decrease, so the shell parameters form an interval; if $A \ge 1/(2r)$, Lemma 21.9 (including its final paragraph) applies throughout H. Recall $u/(4\sqrt{2}\,r^{3/2}) = W_{\mathrm{real}}$, and put $K := 32 \cdot 210^{-1/3} = 5.3836\ldots$, an upper bound for $n\,q^{-1/3}$ by Lemma 20.14 (the sharp value $2.6918\ldots$ from the proof of Theorem 20.16 would improve the constants below and is not needed). (0) $A < 1/(2r)$. As in (b) of Theorem 20.16, $d\sigma'/\sigma = d^2(2Aw + P)/(2w\sigma^2)$ is ≤ 0 or $< 3d^2/(2w\sigma^2) < 3/16$, so $v' > 13/(32\sqrt{r})$ and $f' > 0.81\sqrt{N/r} - 1 \ge 0.81\,N^{1/3} - 1 > 2W$ (as $2W < 2W_{\mathrm{real}} + 2 \le N^{1/5}/\sqrt{2} + 2$ for $r \ge N^{1/5}$). On the interval of shell parameters f is increasing with $f' > 2W$, so the d with $f(d) \in I_t$ form an interval of length < 1, which contains at most one integer: the window holds at most one pair. Assume $A \ge 1/(2r)$ from now on, so that $f' = -2\Delta_d u\varphi - 1$ on H. (1) $\Delta_d > 0$. Then $|f'| = 1 + 2\Delta_d u\varphi \ge 1 + \Delta_d W_{\mathrm{real}}/A$, f is monotone on H, and $\{f \in I_t\}$ is an interval of length at most $2W/(1 + \Delta_d W_{\mathrm{real}}/A) \le \min\{2W,\ 2WA/(\Delta_d W_{\mathrm{real}})\} \le \min\{2W,\ 16\sqrt{2}\,Aq^2/|M|\}$, because $W/W_{\mathrm{real}} \le 2\sqrt{2}$ and $\Delta_d = M/(4q^2)$. With the population bound the count is at most $n\big(1 + \min\{2W/q,\ 16\sqrt{2}\,Aq/|M|,\ \sqrt{r/A}/q\}\big)$; the last two terms agree at $A^{3/2} = \sqrt{r}|M|/(16\sqrt{2}\,q^2)$, where both equal $(16\sqrt{2})^{1/3}r^{1/3}q^{-1/3}|M|^{-1/3} = 2\sqrt{2}\,r^{1/3}q^{-1/3}|M|^{-1/3}$, so the count is at most $n + \min\{2Wn,\ 2\sqrt{2}\,K\,r^{1/3}\} \le n(1 + 15.3\min\{W, r^{1/3}\})$, since $2\sqrt{2}\,K = 15.23\ldots$. (2) $\Delta_d < 0$. Put $g := 2|\Delta_d|u\varphi$, decreasing on H with $g \ge g_0 := |\Delta_d|W_{\mathrm{real}}/A$, so that $f' = g - 1$ is decreasing and f is concave on H, which splits into three consecutive intervals $D_1 = \{g \ge 3/2\}$, $R = \{1/2 < g < 3/2\}$ and $D_2 = \{g \le 1/2\}$. (2a) On D_1, $f' \ge \max\{1/2, g/3\} \ge \max\{1/2, g_0/3\}$, so $\{f \in I_t\} \cap D_1$ is an interval of length at most $\min\{4W, 6W/g_0\}$, and $6W/g_0 = 6WA/(|\Delta_d|W_{\mathrm{real}}) \le 48\sqrt{2}\,Aq^2/|M| < 68\,Aq^2/|M|$. As in (1) the count is at most $n(1 + \min\{4W/q,\ 68Aq/|M|,\ \sqrt{r/A}/q\}) \le n + \min\{4Wn,\ 68^{1/3}K\,r^{1/3}\} \le n(1 + 22\min\{W, r^{1/3}\})$, since $68^{1/3}K = 21.97\ldots$. (2b) On D_2, $-1 \le f' \le -1/2$, so $\{f \in I_t\} \cap D_2$ is an interval of length at most $4W$. Since $g \le 1/2$ at a point of H and $g \ge g_0$, $|\Delta_d| \le A/(2W_{\mathrm{real}})$, i.e. $A \ge W_{\mathrm{real}}|M|/(2q^2)$, and the population bound gives at most $n(\sqrt{2r/(W_{\mathrm{real}}|M|)} + 1)$ pairs. Hence the count is at most $n(1 + \min\{4W, \sqrt{2r/W_{\mathrm{real}}}\}) \le n(1 + 4\min\{W, r^{1/3}\})$: if $W_{\mathrm{real}} \ge r^{1/3}/8$ then $\sqrt{2r/W_{\mathrm{real}}} \le \sqrt{16\,r^{2/3}} = 4r^{1/3}$, and otherwise $W < W_{\mathrm{real}} + 1 < r^{1/3}/8 + 1 \le r^{1/3}$ for r large. (2c) On R, $f'' = g\,\varphi'/\varphi \le -\frac{1}{2} \cdot Ad/w \le -Ad/(4r)$ by Lemma 21.9. Since $g < 3/2$ on R and $g \ge g_0$, $|\Delta_d| < 1.5\,A/W_{\mathrm{real}} \le 3\sqrt{2}\,A < 4.25\,A$, so $C = (1/4 + |\Delta_d|)/A <$

$1/(4A) + 4.25 \le r/2 + 4.25$ and $Ad^2 = S - C > r/2 - 4.25$, whence $d > 0.7\sqrt{r/A}$ for $r \ge 425$ and $m := \min_R |f''| \ge 0.175\sqrt{A/r}$. A concave function with $f'' \le -m$ takes values in an interval of length $2W$ on a set of at most two intervals of total length at most $4\sqrt{W/m} \le 9.57\sqrt{W}\,(r/A)^{1/4}$. Moreover $|\Delta_d| < 1.5\,A/W_{\mathrm{real}}$ gives $A > W_{\mathrm{real}}|M|/(6q^2)$, so $(r/A)^{1/4} < 1.57\,q^{1/2}r^{1/4}(W_{\mathrm{real}}|M|)^{-1/4}$, and $\sqrt{W}\,W_{\mathrm{real}}^{-1/4} \le 1.69\,W_{\mathrm{real}}^{1/4}$ (for $W_{\mathrm{real}} \ge 1$ use $W \le 2W_{\mathrm{real}}$; for $W_{\mathrm{real}} < 1$, $W = 1$ and $W_{\mathrm{real}}^{-1/2} \le 2^{3/4} = 1.682$). Hence the total length is at most $25.4\,(W_{\mathrm{real}}r)^{1/4}q^{1/2}|M|^{-1/4} = 25.4 \cdot 2^{-3/8}N^{1/8}r^{-1/8}q^{1/2}|M|^{-1/4} \le 19.6\,N^{1/8}r^{-1/8}q^{1/2}|M|^{-1/4}$, and the two intervals contain at most $n(19.6\,N^{1/8}r^{-1/8}q^{-1/2}+2)$ admissible d ($|M| \ge 1$). Adding (2a)–(2c), the count for $\Delta_d < 0$ is at most $n(4 + 26\min\{W, r^{1/3}\} + 19.6\,N^{1/8}r^{-1/8}q^{-1/2})$, and (0) and (1) are smaller. □

The theorem does not cover $\Delta_d = 0$; in the squarefree application below this is harmless, since the Beatty identity makes $4k_-k_+$ a square, impossible for two distinct squarefree cells. Least periods $q > r$ are covered by Lemma 20.15: such a family has at most $2^{\omega(q)} = r^{o(1)}$ shell pairs altogether.

The proof gives slightly more than the statement: the resonant term may be replaced by $20\,N^{1/8}r^{-1/8}q^{-1/2}|M|^{-1/4}$, since the level set of (2c) is the union of at most two intervals of total length $L \le 19.6\,N^{1/8}r^{-1/8}q^{1/2}|M|^{-1/4}$ and a residue class modulo q meets it in at most $L/q + 2$ points. Regime (2c), which occurs only when $\Delta_d < 0$ (equivalently $M < 0$), can be non-empty only if $|M|W_{\mathrm{real}}/(6q^2) < A < 4\sqrt{2}\,|M|W_{\mathrm{real}}/q^2$, since $g = 2|\Delta_d|u\varphi \in [|\Delta_d|W_{\mathrm{real}}/A,\ 8\sqrt{2}\,|\Delta_d|W_{\mathrm{real}}/A]$ by Lemma 21.9. In particular families with $C = 0$, which have $M = q^2$ and $\Delta_d = 1/4 > 0$, are never in the resonant regime. The refinement counts the families that can reach a given fraction of the resonant scale. Here, as throughout, a family is a pair (A, C) together with all the residue classes on which both members are integral; sub-families on fewer classes are not counted separately.

Corollary 21.11 (thin resonant hierarchy)**.** *Let $N^{1/5} \le r \le N^{1/3}$, N large, $\varepsilon > 0$ and $0 < \lambda \le N^{1/8}r^{-1/8}$. The number of drift-free two-progression families with $\Delta_d \ne 0$, $C \ne 0$ and least period $q \le r$ that have, in some window of $2W - 1$ consecutive start differences, at least $\lambda + 2^{\omega(q)}(4 + 26\min\{W, r^{1/3}\})$ pairs of shell cells is at most $C_\varepsilon N^\varepsilon (N^{1/8}r^{-1/8}/\lambda)^4$. In particular at $r = \lfloor N^{1/3} \rfloor$, for every $K \ge 1$ at most $C_\varepsilon N^\varepsilon K^4$ such families have a cluster of at least $N^{1/12}/K + 30 \cdot 2^{\omega(q)}$.*

Proof. By the refined bound, such a family has $20 \cdot 2^{\omega(q)}N^{1/8}r^{-1/8}q^{-1/2}|M|^{-1/4} \ge \lambda$, i.e. $q^2|M| \le X_q := 16^{\omega(q)}(20\,Y)^4$ with $Y := N^{1/8}r^{-1/8}/\lambda \ge 1$. A family with $C \ne 0$ is determined by its least period q, its M and its $\alpha = Aq^2$, a divisor of $|q^2(q^2 - M)/4| = |\alpha\gamma|$ (Lemma 20.4). A qualifying family has at least two shell pairs, with parameters $d_1 \ne d_2 < r/2$ and $s_i = Ad_i^2 + C \in (r, 2r)$, so $A = (s_1 - s_2)/(d_1^2 - d_2^2) < r$, $|C| < 2r + r^3/4$, and hence $|\alpha\gamma| = A|C|q^4 \ll r^8 \le N^{8/3}$: every divisor count is $\le C_\varepsilon N^\varepsilon$, whatever λ. Hence the number of families is at most $C_\varepsilon N^\varepsilon \sum_{q:\, q^2 \le X_q}(2X_q/q^2 + 1) \le 3C_\varepsilon N^\varepsilon \sum_{q \ge 1} X_q/q^2$, since $1 \le X_q/q^2$ on the range of summation, and $\sum_q X_q/q^2 = (20Y)^4 \sum_q 16^{\omega(q)}q^{-2} = (20Y)^4 \prod_p (1 + 16/(p^2 - 1)) = O(Y^4)$. At $r = \lfloor N^{1/3} \rfloor$, $W = 1$, $N^{1/8}r^{-1/8} = N^{1/12}(1 + o(1))$ and $\lambda = N^{1/12}/K$ gives $Y = K(1 + o(1))$, while $2^{\omega(q)}(4 + 26) = 30 \cdot 2^{\omega(q)}$. □

Section 23.3 counts 42–68 enumerated families with cluster $\ge \frac{1}{2}N^{1/12}$ and 490–767 with cluster $\ge \frac{1}{4}N^{1/12}$ at $r = \lfloor N^{1/3} \rfloor$ across 40–96 bits. These are not a test of the corollary – its thresholds carry the additional term $30 \cdot 2^{\omega(q)}$, and separate $O(K^4)$ upper bounds do not bound the ratio of two counts – but they are consistent with a hierarchy in which the number of families reaching a K-th of the resonant scale grows polynomially in K.

The bound has the shape of Conjecture (P′): the offset and the speed drift confine a single family to $O(q^{o(1)}\min\{W, r^{1/3}\})$ pairs, and in the proof the term $N^{1/8}r^{-1/8}$ arises only from the stationary regime (2c) of a resonant family ($\Delta_d < 0$, $A \asymp W_{\mathrm{real}}|M|/q^2$) reaches $N^{1/8}r^{-1/8}$; the class factor $2^{\omega(q)} = q^{o(1)}$ is Lemma 20.14, no longer an assumption. Theorem 21.5 matches the $\min\{r^{1/3}, W\}$ term from below when $W_{\mathrm{real}} \to \infty$, and Theorems 21.2, 21.4 match the resonant scale $N^{1/8}r^{-1/8}$ for the full shell when W_{real} is bounded, so in these regimes the single-family profile is matched up to $r^{o(1)}$ factors (up to constants for fixed period); for the squarefree shell at bounded W Theorem 21.3 gives the lower scale for almost all N, and for $W_{\mathrm{real}} \to \infty$ above the crossing the resonance term is matched only by the first-order profile of Section 21.8. Conjecture (P‴) of Section 22 asserts that the whole statistic is carried by such a family; with Theorem 21.10 it implies (P′).

21.10. **An unconditional upper bound at the top of the range.** Everything so far bounds the statistic from below. At $r = \lfloor N^{1/3} \rfloor$ it can also be bounded from above, by reducing it to a count of integer points close to an explicit convex curve and applying the second-difference argument of Filaseta and Trifonov [24]; see [43] for the general theory. Write $\|x\| := \min_{n\in\mathbb{Z}} |x - n|$.

Lemma 21.12 (near-points on a convex curve)**.** *Let $f \in C^2[a, b]$ with $b - a = H \geq 1$ and $0 < \lambda \leq f'' \leq \Lambda$ on $[a, b]$, let $0 < \delta \leq 1/4$, and put $h_0 := \min\{(3\Lambda)^{-1/3}, (8\delta)^{-1}\}$. Then the number R of integers $x \in [a, b]$ with $\|f(x)\| < \delta$ satisfies*

$$R \;\leq\; \Big(\frac{H}{h_0} + 1\Big)\Big(4\sqrt{\delta/\lambda} + 2\Big).$$

Proof. For each such x let $y(x)$ be the integer nearest to $f(x)$. (i) Let $x < x' < x''$ be three such integers with $h := x' - x \leq h_0$ and $h' := x'' - x' \leq h_0$, and let $\Delta_y := [(y(x'') - y(x'))/h' - (y(x') - y(x))/h]/(h + h')$, so that $hh'(h + h')\Delta_y \in \mathbb{Z}$. The same expression formed with the values of f equals $f''(\xi)/2$ for some $\xi \in (x, x'')$, and differs from Δ_y by at most $2\delta/(hh')$. Hence $|hh'(h+h')\Delta_y| \leq \Lambda hh'(h+h')/2 + 2\delta(h+h') \leq \Lambda h_0^3 + 4\delta h_0 \leq 1/3 + 1/2 < 1$, so $\Delta_y = 0$: the three points $(x, y(x))$, $(x', y(x'))$, $(x'', y(x''))$ are collinear. (ii) Partition the near-points into maximal chains in which consecutive members differ by at most h_0. By (i) every three consecutive members of a chain are collinear, so a chain lies on one line $y = \alpha x + \beta$. Put $g := f - \alpha x - \beta$; then $g'' \geq \lambda$ and $|g| < \delta$ at the chain's points. If $\min g > -\delta$, then $g(x) > -\delta + \lambda(x - x_0)^2/2$ about the minimiser x_0, so the chain lies in an interval of length at most $4\sqrt{\delta/\lambda}$. Otherwise the chain lies in the two components of $\{-\delta < g < \delta\}$; on the right component, of length ℓ, g increases by at most 2δ while $g'(x) \geq \lambda \cdot (x - x_{\mathrm{start}})$, so $2\delta \geq \lambda\ell^2/2$ and $\ell \leq 2\sqrt{\delta/\lambda}$, and symmetrically on the left. Either way the chain has at most $4\sqrt{\delta/\lambda} + 2$ points. (iii) Distinct chains are separated by gaps exceeding h_0, so there are at most $H/h_0 + 1$ of them. □

Theorem 21.13 (U, upper envelope at $r = \lfloor N^{1/3} \rfloor$)**.** *For all sufficiently large N,*

$$D_{\max}(\mathcal{F}_1(\lfloor N^{1/3} \rfloor), 1) \;\leq\; 6N^{2/9},$$

and the same bound holds for the squarefree shell.

Proof. Let $r = \lfloor N^{1/3} \rfloor$ and $u = 2\sqrt{N} \in [2r^{3/2}, 2(r + 1)^{3/2})$. The starts $s_k = \lceil u\sqrt{k} \rceil - N - k$ increase with k on the shell, so $D_{\max} = \max_{t\geq 1} D(t)$ with $D(t) := \#\{(k, d) : k,\ k + d \in (r/2, r],\ d \geq 1,\ s_{k+d} - s_k = t\}$. (a) *One k per d.* Fix d and put $\varphi_d(k) := u(\sqrt{k + d} - \sqrt{k}) - d$. Since $\lceil x \rceil - \lceil y \rceil \in (x - y - 1, x - y + 1)$, $s_{k+d} - s_k = t$ forces $\varphi_d(k) \in (t - 1, t + 1)$. Now φ_d is decreasing with $|\varphi_d'(k)| = ud/(2\sqrt{k}\sqrt{k + d}(\sqrt{k + d} + \sqrt{k})) \geq ud/(4r^{3/2}) \geq d/2$ on the shell, so the real k in the shell with $\varphi_d(k) \in (t - 1, t + 1)$ form an interval of length $< 4/d$, which for $d \geq 5$ contains at most one integer. For an actual pair put $T := t + d$, $q := T/u$ and

$q_0 := \sqrt{k+d} - \sqrt{k} = d/(\sqrt{k+d}+\sqrt{k})$, so that $|q-q_0| < 1/u$ and $d/(2\sqrt{r}) \le q_0 \le d/\sqrt{2r}$, whence $|q-q_0| \le q_0/(dr)$ and $q^2 \le d/2$ for r large. The equation $\sqrt{\kappa+d} - \sqrt{\kappa} = q$ has the solution $K(q) := \frac{1}{4}(d/q - q)^2$, with $K(q_0) = k$, and we set $\kappa_t(d) := K(T/u)$, so that

$$\kappa_t(d) \;=\; \frac{1}{4}\Big(\frac{ud}{t+d} - \frac{t+d}{u}\Big)^2 \;=\; \frac{Nd^2}{(t+d)^2} - \frac{d}{2} + \frac{(t+d)^2}{16N}.$$

Since $K'(q) = -\sqrt{K(q)}\,(d+q^2)/q^2$ and $\sqrt{K(q)} \le d/(2q)$, for q_* between q_0 and q we have $|K'(q_*)| \le d(d+q_*^2)/(2q_*^3) \le 0.75\,d^2/q_*^3 \le 6.1\,r^{3/2}/d$ (using $q_*^2 \le d/2$ and $q_* \ge q_0(1-1/(dr)) \ge (d/(2\sqrt{r}))(1-1/(dr))$), hence $|\kappa_t(d) - k| \le |K'(q_*)|/u \le 3.1/d < 4/d$. Thus $\|\kappa_t(d)\| < 4/d$ and $\kappa_t(d) \in (r/2-1, r+1]$. The values $d \le 4$ contribute at most $\sum_{d\le 4}(4/d+1) < 13$ pairs, so $D(t) \le 13 + \#\{d \ge 5 : \|\kappa_t(d)\| < 4/d,\ \kappa_t(d) \in (r/2-1, r+1]\}$. (b) *The range of d.* Write $c_0 := 1-2^{-1/2} = 0.2928\ldots$ and $T := t+d = \lceil u\sqrt{k+d}\rceil - \lceil u\sqrt{k}\rceil \le u(\sqrt{r}-\sqrt{r/2})+1 = c_0 u\sqrt{r}+1$, so $t < 0.3\,u\sqrt{r} \le 0.61\,r^2$ for r large and $T/u < c_0\sqrt{r} + 1/u$. Since $\kappa_t(d) \in (r/2-1, r+1]$ and $ud/T = 2\sqrt{\kappa_t(d)} + T/u$,

$$\frac{d}{T} \;>\; \frac{\sqrt{2r-4}}{u} \;\ge\; \frac{\sqrt{1/2-1/r}}{r(1+1/r)^{3/2}},$$
$$\frac{d}{T} \;\le\; \frac{2\sqrt{r+1} + c_0\sqrt{r} + 1/u}{2r^{3/2}} \;\le\; \frac{1}{r}\Big(\sqrt{1+1/r} + \frac{c_0}{2} + \frac{1}{2u\sqrt{r}}\Big);$$

the limiting constants are $2^{-1/2} = 0.7071\ldots$ and $1 + c_0/2 = 1.1464\ldots$, so $0.7/r < d/T \le 1.15/r$ for r large. As $d/t = (d/T)/(1-d/T)$, this gives $d \in [X_1, X_2]$ with $X_1 := 0.7\,t/r$ and $X_2 := 1.16\,t/r$; put $H := X_2 - X_1 \le 0.46\,t/r$. If $t < 23r$ then $X_2 < 27$ and $D(t) < 40$; assume $t \ge 23r$. (c) *The curve.* On $[X_1, X_2]$ we have $d \le 1.16\,t/r$ and

$$\kappa_t''(d) \;=\; \frac{2Nt(t-2d)}{(t+d)^4} + \frac{1}{8N} \;\in\; \Big[\frac{2N}{t^2}\Big(1-\frac{7}{r}\Big),\ \frac{2N}{t^2}\Big(1+\frac{1}{r}\Big)\Big] \;=:\; [\lambda, \Lambda].$$

(d) *The count.* Apply Lemma 21.12 to $f = \kappa_t$ on $[X_1, X_2]$ with $\delta := 4/X_1 = 40r/(7t) < 5.72\,r/t$, which is $\le 1/4$ for $t \ge 23r$. Then $H/h_0 \le H(3\Lambda)^{1/3} + 8\delta H \le 0.46\,(t/r)(6N/t^2)^{1/3}(1+1/r)^{1/3} + 21.1 \le 0.85\,t^{1/3} + 22$ for r large (using $N^{1/3} \le r+1$), and $\sqrt{\delta/\lambda} \le \sqrt{2.86\,rt/N}\,(1+4/r) \le 1.33$ (using $t \le 0.61r^2$ and $N \ge r^3$). Hence $D(t) \le 13 + (0.85\,t^{1/3}+23)(7.4) \le 6.3\,t^{1/3} + 184 \le 5.4\,r^{2/3} + 184$, since $t^{1/3} \le (0.61)^{1/3}r^{2/3}$. For N large this is at most $6N^{2/9}$. Monotonicity of $D_{\max}$ in the family gives the squarefree case. □

With Theorem 21.2, $(0.633 - o(1))N^{1/12} \le D_{\max}(\mathcal{F}_1(\lfloor N^{1/3}\rfloor), 1) \le 6N^{2/9}$: for the full shell the exponent at the top of the range lies in $[1/12, 2/9]$. Conjecture (P′) concerns the squarefree shell, for which it predicts the exponent $1/12$ (matched from below by Theorem 21.3 for almost all N); the analogous full-shell conjecture would assert $1/12$ for $\mathcal{F}_1$ as well. The bound is dominated by the number of chains, $H/h_0 \asymp t^{1/3}$, which is $\asymp N^{1/6}$ at the resonant scale $t \asymp N^{1/2}$ of Theorem 21.2 and $\asymp N^{2/9}$ for the largest $t \asymp N^{2/3}$, where the two cells of a pair lie at opposite ends of the shell. Only the second derivative of κ_t was used. Exact differentiation gives $\kappa_t'''(d) = 12Nt(d-t)/(t+d)^5 = -(12N/t^3)(1+O(1/r))$ on the relevant interval, so higher-order near-point estimates [43] are expected to improve the exponent for $t = o(N^{2/3})$, although none is proved here; at $t \asymp N^{2/3}$ the derivative-only third-difference spacing scale is $N^{1/6}$, and the standard higher-derivative tests do not by themselves evidently reach the conjectural $N^{1/12}$. One route to that scale is the arithmetic of the rational parabolas carrying many near-points – the two-progression families of Section 20; another is a direct near-curve or exponential-sum estimate controlling points spread over many parabolas. Nothing here distinguishes them. In the terms of Theorem 4.4, a uniform power envelope $D_x \le x^\eta$ of the form used there improves

the in-model bound beyond $N^{1/6}$ only once $\eta < 1/2$; Theorem 21.13's $2/3$ (as a power of the shell size $r/2$) does not, and the gap from $2/3$ to the conjectured $1/4$ is the upper-envelope problem in its sharpest form.

The same argument bounds the modulus-free statistic of Section 20, where the curve is exactly quadratic.

Theorem 21.14 (U*, upper envelope of the modulus-free statistic)**.** *For every $\theta \in (0, 1.2]$ and all sufficiently large r, $D^*_\theta(r) \le 2.1\, r^{2/3}$, and the same bound holds for the modulus-free statistic of every subset of the cells of the shell. Consequently, for every fixed $\varepsilon > 0$ there are r_0 and $N_0(\varepsilon)$ such that $r_0 \le r \le N^{1/5-\varepsilon}$ and $N \ge N_0(\varepsilon)$ imply $D_{\max}(\mathcal{F}_1^{\square}(r), W) \le 2.1\, r^{2/3}$.*

Proof. Reversing the ordered pairs replaces τ by $-\tau$, and at $\tau = 0$ the window contains no shell speed for r large (every speed is at least $1/(2\sqrt{r})$ while $\theta\rho_r = O(r^{-3/2})$), so fix $\tau > 0$ and the window $(\tau-\theta\rho_r, \tau+\theta\rho_r)$; put $\eta := \theta\rho_r$. For a pair $(k, k+d)$ of the shell put $v_d(k) := \sqrt{k+d}-\sqrt{k}$, which decreases in k with $|v_d'(k)| = d/(2\sqrt{k}\sqrt{k+d}(\sqrt{k}+\sqrt{k+d})) \ge d/(4r^{3/2})$; the real k with $v_d(k)$ in the window form an interval of length at most $2\eta \cdot 4r^{3/2}/d = \sqrt{2}\theta/d \le 1.7/d$, so at most two pairs have $d = 1$ and at most one pair has each $d \ge 2$; the pairs with $d \le 3$ number at most 4. For $d \ge 4$ put $K(q) := \frac{1}{4}(d/q - q)^2$, so that $K(q_0) = k$ for $q_0 := v_d(k)$, and $\kappa_\tau(d) := K(\tau) = d^2/(4\tau^2) - d/2 + \tau^2/4$. On the segment between q_0 and τ we have $|q - q_0| < \eta$ and $q_0 \ge d/(2\sqrt{r})$, so $|q - q_0|/q_0 \le \theta/(2\sqrt{2}\, dr) = O(r^{-1})$ uniformly for $d \ge 4$, $\theta \le 1.2$; since $|K'(q)| = (d^2 - q^4)/(2q^3)$ and $q_0^2 < d/4$, this gives $|K'(q)| = |K'(q_0)|(1 + O(r^{-1}))$, and $|K'(q_0)| = 1/|v_d'(k)| \le 4r^{3/2}/d$ because K inverts v_d. Hence $|K'(q)| \le 4.5\, r^{3/2}/d$ for r large, and $|k - \kappa_\tau(d)| \le 4.5\, r^{3/2}\eta/d < 0.8\,\theta/d \le 0.96/d$: every $d \ge 4$ carrying a pair of the window is a near-point $\|\kappa_\tau(d)\| < 1/d$. Since $d = q_0(\sqrt{k} + \sqrt{k+d})$ with $\sqrt{2r} < \sqrt{k} + \sqrt{k+d} \le 2\sqrt{r}$, every such d lies in $[X_1, X_2]$ with $X_1 := \max\{4, \sqrt{2r}(\tau - \eta)\}$ and $X_2 := 2\sqrt{r}(\tau + \eta)$. If $H := X_2 - X_1 < 1$ there is at most one such d and the window holds at most 5 pairs, so assume $H \ge 1$; a carrying $d \ge 4$ forces $\tau\sqrt{r} \ge 2 - O(r^{-1})$, hence $\eta/\tau = O(r^{-1})$, $X_1 \ge \tau\sqrt{2r}\,(1 + O(r^{-1}))$ and $H \le (2 - \sqrt{2})\tau\sqrt{r}\,(1 + O(r^{-1}))$, uniformly in τ; also $\tau \le (1 - 2^{-1/2})\sqrt{r} + \eta$ for a non-empty window. Apply Lemma 21.12 to $f = \kappa_\tau$ on $[X_1, X_2]$ with $f'' = 1/(2\tau^2) =: \lambda = \Lambda$ and $\delta := 1/X_1 \le 1/4$. Then $\delta/\lambda = 2\tau^2/X_1 \le \sqrt{2}\,\tau/\sqrt{r}\,(1 + O(r^{-1})) \le \sqrt{2} - 1 + O(r^{-1}) < 0.42$, so a chain has fewer than $4\sqrt{0.42} + 2 < 4.6$ points; and $h_0 = \min\{(2\tau^2/3)^{1/3}, X_1/8\}$ gives

$$\frac{H}{h_0} \le \max\Big\{\frac{2 - \sqrt{2}}{(2/3)^{1/3}}\,\tau^{1/3}\sqrt{r},\ \frac{8(2 - \sqrt{2})}{\sqrt{2}}\Big\}(1 + O(r^{-1})),$$

where $(2 - \sqrt{2})/(2/3)^{1/3} = 0.670558\ldots$ and $8(2 - \sqrt{2})/\sqrt{2} = 3.3137\ldots$, and moreover $\tau^{1/3} \le (1 - 2^{-1/2})^{1/3} r^{1/6}(1 + o(1))$ with $(1 - 2^{-1/2})^{1/3} = 0.6641045\ldots$; the product of the two constants is $0.44532\ldots$, so $H/h_0 \le 0.446\, r^{2/3}$ for r large, uniformly in τ and θ. Hence the window contains at most $4 + 4.6\,(0.446\, r^{2/3} + 1) = 2.0516\, r^{2/3} + 8.6 \le 2.1\, r^{2/3}$ pairs for r large. A subset of the cells has fewer pairs. For the last statement apply Lemma 20.9 with $\theta_0 = 0.9 < 1 < 1.1 = \theta'$: $D_{\max}(\mathcal{F}_1^{\square}(r), W) \le D^*_{1.1}(r) \le 2.1\, r^{2/3}$ once $r \ge r_0$. □

With Theorem 20.8, for every fixed $0 < \theta \le 1$ and as $r \to \infty$ the exponent of $D^*_\theta(r)$ lies in $[1/3, 2/3]$, and along every sequence $r \to \infty$ in the transfer range Lemma 20.9 gives the same interval for $D_{\max}(\mathcal{F}_1^{\square}(r), W)$; Conjecture 20.12 attributes the conjectural $r^{1/3+o(1)}$ envelope to a single drift-free two-progression family, and Theorem 20.11 gives the proven lower bound $D^*_{2P,\theta}(r) \ge (c_\theta - o(1))r^{1/3}$. The bound is again carried by the chain count $H/h_0 \asymp \tau^{1/3}\sqrt{r}$, largest for the fastest pairs. The near-points contributed by a fixed near-chain or two-progression family lie on that family's parabola in the (d, k)-plane, so a non-degenerate family contributes

at most two points to any collinear chain of Lemma 21.12: for these mechanisms a sharper argument must control the number of chains rather than their length.

21.11. **The additive energy of the speeds.** Theorems 20.11 and 21.14 concern the maximum of the modulus-free statistic. Its second moment is controlled unconditionally. Let $\mathcal{K}$ be the squarefree cells in $(r/2, r]$ (or any subset of them), $R := |\mathcal{K}|$, $\mathcal{P} := \{(k, k') : k, k' \in \mathcal{K},\ k < k'\}$ the pairs, and for $\theta > 0$

$$E_\theta(r) := \#\{(P, Q) \in \mathcal{P}^2 : |v_P - v_Q| < \theta\rho_r\},$$

the number of ordered pairs of pairs whose speeds differ by less than $\theta\rho_r$; $P = Q$ is allowed, so $E_\theta(r) \geq |\mathcal{P}| = \binom{R}{2}$.

Proposition 21.15 (E, near-minimal additive energy)**.** *For fixed $\theta > 0$ and $\varepsilon > 0$, $E_\theta(r) \ll_{\theta,\varepsilon} r^{2+\varepsilon}$, for every subset $\mathcal{K}$ of the shell. Consequently the number of pairwise disjoint windows of half-width $\theta\rho_r/2$ containing at least λ pairs is $\ll r^{2+\varepsilon}/\lambda^2$; with $C(P) := \#\{Q : |v_Q - v_P| < \theta\rho_r\}$, $\#\{P : C(P) \geq r^\eta\} \ll r^{2+\varepsilon-\eta}$ for every $\eta > 0$; and for the full squarefree shell the mean of $C(P)$ over $\mathcal{P}$ is $\ll r^\varepsilon$.*

Proof. $|v_P - v_Q| < \theta\rho_r$ reads $|\sqrt{k'} + \sqrt{l} - \sqrt{k} - \sqrt{l'}| < \theta\rho_r$ with all four cell labels in $(r/2, r] \subset [M, 2M]$, $M := r/2$. Robert and Sargos [73] proved that for fixed real $\alpha \neq 0, 1$ the number of $(m_1, m_2, m_3, m_4) \in [M, 2M]^4$ with $|m_1^\alpha + m_2^\alpha - m_3^\alpha - m_4^\alpha| \leq \delta M^\alpha$ is $\ll_{\alpha,\varepsilon} M^{2+\varepsilon}(1 + \delta M^2)$. With $\alpha = 1/2$ and

$$\delta := \theta\rho_r M^{-1/2} = \frac{\theta}{4r^2}, \qquad \delta M^2 = \frac{\theta}{16},$$

this gives $E_\theta(r) \ll r^{2+\varepsilon}$. If λ pairs lie in one window of half-width $\theta\rho_r/2$, every two of them are within $\theta\rho_r$ and contribute λ^2 to E_θ, and disjoint windows contribute disjoint sets of ordered pairs. $\sum_P C(P) = E_\theta(r)$ gives the tail bound by Markov's inequality and, since $|\mathcal{P}| = \binom{R}{2} \asymp r^2$ for the full squarefree shell, the mean. □

Thus, for the full squarefree shell, the energy has the near-minimal order $r^{2+o(1)}$ unconditionally: polynomially large clusters occupy a vanishing proportion of the pairs, although the maximal cluster is at least $(c_\theta - o(1))r^{1/3}$ (Theorem 20.11). This is a tail statement, not a Poisson pair-correlation limit. Numerically (Section 23) the energy of the squarefree shell at $r = 2^{12}, \dots, 2^{15}$ agrees with that of a phase-randomised null to within one unit in the third decimal of $E_1/\binom{R}{2}$ ($1.074, 1.074, 1.074, 1.075$ against $1.073, 1.074, 1.074, 1.074$; $E_1 = 0.0496\, r^2$ in both), while the maximal speed-centred clusters are $9, 9, 11, 15$ against $4, 5, 5, 5$, and the pairs whose cluster exceeds the null's maximum number $61, 97, 332, 1420$ – between $2.7 \cdot 10^{-5}$ and $7.9 \cdot 10^{-5}$ of all pairs.

The energy form of Lemma D, and a transfer-range limitation. The proof of Lemma 4.3 bounds the cross-window overlaps at the babies by $W \sum_{\gamma \neq \gamma'} D(\gamma' - \gamma) \leq W K_2^2 D_{\max}$. The $K_2(K_2 - 1)$ differences $\gamma' - \gamma$ have multiplicities at most K_2, so the sum is at most $K_2 \Sigma_{K_2}$, where Σ_m denotes the sum of the m largest values of $D(h)$ over $h \neq 0$ and $\bar{\Sigma}_m := \Sigma_m/m$ is non-increasing in m. The baby count gives $K_1 K_2^2 \bar{\Sigma}_{K_2} \geq n'^2/(4W)$ when $K_1 \leq n'/(2W)$, and the giant count gives $K_2 K_1^2 \bar{\Sigma}_{K_1} \geq n'^2/(4W)$ when $K_2 \leq n'/(2W)$. If either $K_i > n'/(2W)$ then $|B| + |G| \geq n'/(2W)$. Otherwise put $m := \min(K_1, K_2)$; multiplying the two inequalities and using the monotonicity $\bar{\Sigma}_{K_i} \leq \bar{\Sigma}_m$ gives $K_1 K_2 \geq (n'^2/(4W\bar{\Sigma}_m))^{2/3}$, and since the larger of K_1, K_2 equals $K_1 K_2/m$,

$$|B| + |G| \ \geq\ m + \frac{1}{m}\Big(\frac{n'^2}{4W\bar{\Sigma}_m}\Big)^{2/3}.$$

The integer $m = \min(K_1, K_2)$ depends on the unknown cover, so the cover-uniform energy certificate is

$$|B| + |G| \;\ge\; \min\Big\{ \frac{n'}{2W},\; \min_{m \ge 1} \Big[m + \frac{1}{m} \Big(\frac{n'^2}{4W \bar{\Sigma}_m} \Big)^{2/3} \Big] \Big\},$$

where Σ_m is the sum of the m largest values of $D(h)$ over all signed $h \neq 0$ (with zeros appended if necessary) and $\bar{\Sigma}_m = \Sigma_m / m$; the weaker cover-dependent consequence $|B| + |G| \ge 2(n'^2/(4W\bar{\Sigma}_m))^{1/3}$ follows by AM–GM, and is Lemma 4.3 with $D_{\max}$ replaced by the mean of the m largest values of D. One might hope that $\bar{\Sigma}_m \ll D_{\max}$ at the relevant scale $m \asymp N^{1/5}$, since the maximum is attained by a few resonant families. In the transfer range it is not so. Write $U := 2\sqrt{N}$. For $r \le N^{1/5-\epsilon}$ and every odd $j \in [0.4\, r^{1/3}, r^{1/3}]$ (note $16^{-1/3} = 0.3968\ldots < 0.4$), Lemma 20.7 puts all $|I_j(r)| \ge (\sqrt{2}-1) r^{1/3} - 3$ members of the symmetric family j in one window of half-width $\rho_r/2$; Lemma 20.6 with $z = \log r$ keeps a proportion $\ge \kappa_0 - o(1)$ of them squarefree (its error terms $3L/z$, $4^{\pi(z)}$, $\pi(\sqrt{T_2})$, $2\pi(\sqrt{jT_2+1})$ are all $o(r^{1/3})$ here); inside the window the start differences $\Delta_P = U v_P - d_P + \vartheta_P$ lie in an interval $[a_j, b_j]$ of length $\le U\rho_r + |I_j(r)| + 2 = W_{\mathrm{real}} + O(r^{1/3}) = (1+o(1))W$, so every integer h with $b_j - W < h < a_j + W$ – at least $(1-o(1))W \ge W/2$ of them – has $D(h) \ge (\kappa_0(\sqrt{2}-1) - o(1)) r^{1/3} \ge (0.1299 - o(1)) r^{1/3}$; and the intervals for distinct j, centred near $\sqrt{2N/j}$, are $\asymp \sqrt{N/r} \asymp rW$ apart, hence disjoint. The $(0.3 + o(1)) r^{1/3}$ odd j thus give at least $(0.15 - o(1)) W r^{1/3} \gg N^{1/5}$ integers h with $D(h) \ge (0.1299 - o(1)) r^{1/3}$, and $\bar{\Sigma}_m \ge (0.1299 - o(1)) r^{1/3}$ for every $m \le (0.15 - o(1)) W r^{1/3}$. Under (S″) therefore $\bar{\Sigma}_m = r^{1/3 + o(1)}$ throughout the transfer range for every $m \le N^{1/5}$: the rearrangement estimate cannot replace the power $r^{1/3}$ by a smaller one there. This does not by itself determine the optimised exponent of the energy-form certificate – the $N^{2/11}$ of Theorem 4.4 arises near $r = N^{3/11}$, outside the range treated – and it does not say that the difference set of a particular baby set must meet these values of h; whether the planar regime allows $\bar{\Sigma}_m \ll D_{\max}$ at $m \asymp N^{2/11}$ is left open. At 40–48 bits the measured ratios $D_{\max}/\bar{\Sigma}_m$ at $m = N^{1/5}$ are 1.3–2.2 (Section 23).

22. Conjectures and the refined-prime programme

Conjecture 22.1 (S″, envelope)**.** *$D_{\max}(\mathcal{F}_1^{\square}(r), W) \le r^{1/3 + o(1)}$ for $r \le N^{1/3}$.*

For the modulus-free statistic, Theorem 20.11 gives the lower-bound coefficient c_θ and Theorem 21.14 the upper bound $2.1\, r^{2/3}$; Conjecture 20.12 locates the maximum among the two-progression families, and Lemma 20.9 carries these statements into the transfer range.

Conjecture 22.2 (P′, planar envelope)**.** *For $N^{1/5} \le r \le N^{1/3}$,*

$$D_{\max}(\mathcal{F}_1^{\square}(r), W) \;\le\; r^{o(1)} \max\big(\min(r^{1/3}, 2W),\; N^{1/8} r^{-1/8} \big).$$

The first term combines the symmetric law (measured coefficient ≈ 0.38) with the cap of Proposition 21.1; the second is the first-order resonance range (coefficient 1.94); neither constant is asserted as an upper constant; the class count $n \le 2^{\omega(q)} = q^{o(1)}$ is now Lemma 20.14, and the unproved content of (P′) is the control of pile-ups and of clusters not contained in any drift-free family. Theorem 21.13 proves the unconditional upper exponent $2/9$, for both shells, where (P′) predicts $1/12$ at the top of the range.

Conjecture 22.3 (P‴, planar two-progression extremality)**.** *For $N^{1/5} \le r \le N^{1/3}$,*

$$D_{\max}(\mathcal{F}_1^{\square}(r), W) \;\le\; r^{o(1)}\, D_{2P}^{\mathrm{pl}}(N, r),$$

where $D_{2P}^{\mathrm{pl}}(N, r)$ is the largest number of pairs of squarefree cells of one drift-free two-progression family of any least positive period q (those with $q > r$ contribute $r^{o(1)}$ by Lemma 20.15) whose start differences lie in one window of $2W - 1$ consecutive integers.

By Theorem 21.10 for $q \le r$ and Lemma 20.15 for $q > r$, $2^{\omega(q)} \le \exp(O(\log r/\log\log r)) = r^{o(1)}$ and $q^{-1/2} \le 1$, while the Beatty case $\Delta_d = 0$ contributes no pair of distinct squarefree cells; hence $D^{\mathrm{pl}}_{2P}(N,r) \le r^{o(1)} \max\{\min\{r^{1/3}, 2W\}, N^{1/8} r^{-1/8}\}$, and (P‴) implies (P′). (P‴) is the precise form of the assumption, made informally in (P′), that several families do not pile into one window; its evidence is that at eleven of twelve exact points through 64 bits the maximum inside the census box equals the all-pairs maximum (Section 23) – no omitted mechanism is larger at those points, though ties, submaximal residuals and a later asymptotic crossover are not tested – and that the pooled statistic over all enumerated families (Section 23) shows no pile-up at the sampled radii. Under (S″) below the crossing and (P′) beyond it, Theorem 4.4 gives $\ge c\, N^{2/11-o(1)}$ for every r, its exponent minimised near $r \approx N^{3/11}$ (no representation of that cost is exhibited), and by Corollary 21.7 this is also what the whole-shell certificate can at most give.

Mechanisms outside the covered class. Theorems 20.16 and 21.10 bound one drift-free family at a time. They do not bound: (i) non-drift-free families $s = Ad^2 + Bd + C$, $B \ne 0$ (four appeared among the maximisers of the planar census at small d). Integrality on a progression of period q forces $B \in q^{-2}\mathbb{Z}$ (by the first difference of the corresponding integer-valued quadratic). To first order the speed drift is $B/(2\sqrt{2A}\, w)$ with $w = 2\sqrt{k_+ k_-} \asymp r$, so each integral residue class contributes at most $\min\{2\theta q\sqrt{A/r}+1, \sqrt{r/A}/q+1\} \le \sqrt{2\theta}+1$ pairs away from stationary points of the speed. No analogue of Lemma 20.14 is proved here for $B \ne 0$, so this per-class estimate does not yet give an $O(2^{\omega(q)})$ bound for the whole family; nor does it exclude degenerate stationary configurations – simultaneous vanishing of successive speed derivatives in the modulus-free problem, $f' = f'' = 0$ for the phase $f = uv - d$ in the planar one – which impose Diophantine conditions on (A,B,C) that no argument here excludes. Numerically the mechanism is inert: a census of such families places the first-order stationary point $d_* = 2(B^2 + \Delta_d)/(AB)$, $\Delta_d = 1/4 - AC$, inside the shell by construction (for $q \le 40$, $B = \beta/q$ with β a multiple of $\gcd(2\alpha, q)$ so that many residue classes are admissible, A on a logarithmic grid in $[2^{-10}, 4]$, C near its target), and among 71 520 such families at each of $r = 2^{16}, 2^{18}, 2^{20}$ (7 006–10 340 with two or more shell pairs) the largest window holds $3, 4, 4$ squarefree and $4, 4, 4$ pairs, against the drift-free $16, 31, 45$ and $r^{1/3} = 40, 64, 102$; a scan of ± 300 in γ around the target for twelve (A,B) at 2^{18} reaches 5. (ii) Pile-ups of several families in one window, which the pooled test of Section 23 examines inside the census box. (iii) Higher-degree one-parameter families $d = d(n)$, $s = s(n)$ with $\deg s = 2\deg d$, whose shell population is $O(r^{1/(2\deg d)})$ – at most $r^{1/4}$ beyond the quadratic case, too small for the modulus-free $r^{1/3}$ but of the planar $N^{1/12}$ scale. (iv) Near-integral points of κ_τ or κ_t (Theorems 21.14, 21.13) spread over many parabolas with no exact family carrying most of them – the major-arc alternative, which a direct near-curve or exponential-sum estimate would have to address without any classification. (S‴) and (P‴) assert that these alternatives do not raise the whole-shell maximum beyond the drift-free benchmark, up to the stated $o(r^{1/3})$ and $r^{o(1)}$ losses – not that each is individually smaller; the exact sweeps are finite-size evidence for that extremality assertion, not a proof, and (iv) is the alternative that no argument given here excludes.

22.1. **The refined prime sub-family.** The prime shell $\mathcal{F}_1^P(r)$, of size $(1+o(1))r/(2\ln r)$, does not carry the symmetric pairs (their members $t(jt \mp 1)/2$ are composite), and by Lemma 20.4(iii) every family with infinitely many prime members has non-square M, hence first-order occupancy $O(r^{2/7})$ in the transfer range. In the computations of Section 23, its modulus-free statistic sits at the phase-randomised null through $r = 2^{18}$, and the certified exact values at $r = 2^{19}$ and 2^{20} are 4 (no null was simulated at these two points). We state two completeness conjectures sufficient for one refined-subfamily route; a direct analytic envelope, or a different construction of a low-$D_{\max}$ subfamily, would not require them.

Conjecture 22.4 (T, two-progression completeness)**.** *Let $\mathfrak{G}(r)$ be the collection of drift-free two-progression families of Proposition 20.3 and Lemma 20.4, each restricted to its pairs of squarefree cells in $(r/2, r]$, and for $\mathcal{G} \in \mathfrak{G}(r)$ put*

$$O_\theta(\mathcal{G}; r) := \max_{\tau \in \mathbb{R}} \#\{(k, k') \in \mathcal{G} : |\sqrt{k} - \sqrt{k'} - \tau| < \theta\rho_r\}.$$

Let $E_{\eta,\theta}(r)$ be the set of squarefree cells occurring in a pair of some family with $O_\theta(\mathcal{G}; r) \geq r^\eta$. For every fixed $\eta > 0$ and $\theta \in (0, 1]$, the squarefree cells of $(r/2, r]$ outside $E_{\eta,\theta}(r)$ have modulus-free statistic $D^*_\theta \leq r^{\eta+o(1)}$.

Conjecture 22.5 (T_P, planar refined-prime envelope)**.** *For every fixed $\eta > 0$, uniformly for $N^{1/5} \leq r \leq N^{1/3}$ and also on the $W = 1$ shell at $r = \lfloor N^{1/3} \rfloor$, there is a disjoint sub-family of the $a = 1$ cells with $R' \geq r^{1-o(1)}$ windows and $D_{\max}(\cdot, W) \leq r^{\eta+o(1)}$.*

Under (T) and (T_P), Theorem 4.4 gives $\geq N^{1/5-\delta}$ for every $\delta > 0$ – Harvey's exponent as the in-model optimum. In the transfer range the counting branch alone already gives $N^{1/5}$, so the exponent is decided in the planar regime, where the exact-start statistic enters.

The cost side (first order). The resonance criterion depends on (N, r) only through W_{real}, so at $r = N^{1/3}$ the list of candidate resonant families is N-independent and only the first-order scales change (range $N^{1/12}$, population $N^{1/6}$). A family with n classes and range $\geq c\,N^{1/12}$ has $q^2|M| \leq (1.94n/c)^4$, and summing populations $1.18\,n|M|^{-1/2}N^{1/6}$ over $q^2|M| \leq X$ gives $\ll N^{1/6}X^{1/2}\ln X$. Hence removing every enumerated resonant family above the threshold r^η, $0 < \eta < 1/4$, costs $r^{1-2\eta+o(1)}$ cells (with $n = q^{o(1)}$ assumed), a vanishing fraction; for $\eta \geq 1/4$ nothing exceeds the threshold up to constants. The count controls only the enumerated drift-free resonant families; it gives no bound on the residual statistic after their removal, and that residual bound is the unproved content of (T_P).

The prime onset. On the family of Theorem 21.2, with the window $|d - d_*| \leq 0.8\sqrt{d_*}$, the numbers of d with both $k_\mp(d)$ prime are $7, 12, 60, 334, 4716$ at $96, 128, 160, 200, 256$ bits, of which $5, 11, 39, 197, 2372$ share the most frequent start difference. These counts are lower bounds for $D_{\max}(\mathcal{F}_1^P(r), 1)$: in these computations the plain prime shell leaves its null level of 2–3 between 128 and 160 bits, and under the Bateman–Horn heuristic [5] the number of prime pairs on the window grows like $N^{1/12}/\ln^2 r$, so that the prime shell has no $r^{o(1)}$ planar envelope under that heuristic. The family is admissible on $d \equiv 0 \pmod{12}$, and more generally every normal form $nt^2 \mp mt + 1$ is admissible on $t \equiv 0 \pmod 6$: both members are $\equiv 1$ modulo 2 and 3 there, and for $p \geq 5$ the product $(nt^2+1)^2 - m^2t^2$ is a quartic with leading coefficient n^2, or $1 - m^2t^2$ when $p \mid n$, neither vanishing identically modulo p. Heuristically – under Bateman–Horn, and given an admissible resonant family of the crossing scale – restricting to primes therefore does not by itself lift the ceiling of Corollary 21.7: the logarithms cancel in the Sidon branch at the crossing, and the thinning must remove the resonant families themselves.

22.2. **Where conjecture (T) sits.** The modulus-free statistic $D^*(r)$ is an extreme-value statistic of the pair correlation of the real sequence $\{\sqrt{k} : r/2 < k \leq r\}$ ($\approx r/2$ points, raw mean spacing $\asymp r^{-1/2}$) at a window $\asymp r^{-3/2}$, i.e. r^{-1} times the raw spacing, where the expected count is $O(1)$. Reduction modulo 1 only merges pairs, so D^* is at most the corresponding extreme count for the fractional parts $\{\sqrt{k}\}$ at the same absolute scale ($r^{-1/2}$ times their mean spacing $1/r$). For $\sqrt{n} \bmod 1$ over the full range the two-point correlation is Poissonian once squares are removed [21] – an average over offsets at the mean-spacing scale – while the gap distribution exists and is not Poisson [20], the non-Poisson gaps coming from the lattice structure that appears here as the Beatty chains. For the sequences $\{\alpha n^\theta\}$ the pair correlation is Poissonian for every $\alpha > 0$ when $0 < \theta < 43/117$ [55, 71], and the m-point correlations for $m \geq 3$ when

$0 < \theta < 1/(m^2 + m - 1)$ [56], ranges that exclude $\theta = 1/2$; for almost every α the k-level correlations of $\{n^\alpha\}$ are Poissonian once $\alpha > 4k^2 - 4k - 1$ [79], far from $\alpha = 1/2$. No published result on triple or higher correlations of $\sqrt{n} \bmod 1$, on maximal pair counts at a fixed offset at any scale, or on dyadic ranges, arithmetic progressions or $\sqrt{kN}$ is known to the authors. Conjecture (T) is therefore an extreme-value statement about pair correlations of $\sqrt{k}$ far below the mean spacing, with the arithmetic families removed: related to the known two-point law but not implied by it.

23. Numerical evidence

Unless stated otherwise, all statistics below are exact integer computations; where a statistic depends on real speeds computed in double precision, the error is bounded and the count is bracketed or recomputed in 200-bit precision, as indicated at each place. Balanced semiprimes ("RSA moduli") have $q/p < \sqrt{2}$ and are drawn from a seeded generator; every computation is archived with its parameters and seeds and is regenerated deterministically from them.

23.1. **The modulus-free statistic.** Table 5 gives $D^*(r)$ of the squarefree shell at Harvey's resolution ρ_r against a phase-randomised null (each δ-class shifted by an independent random fraction of its spacing) and the full shell, together with the speed at the returned maximiser.

r	R	D^*	$\eta = \ln D^*/\ln r$	τ^2 at the maximum	null	full shell $(Q-1)$
2^8	79	3	0.198	3.70 (generic)	3–4	5 (4)
2^9	157	4	0.222	3.5 ($j = 7$)	3–4	8 (5)
2^{10}	310	4	0.200	18.77 (generic)	4	10 (9)
2^{11}	621	6	0.235	8.5 ($j = 17$)	4–5	15 (12)
2^{12}	1246	9	0.264	7.5 ($j = 15$)	4–6	18 (18)
2^{13}	2491	9	0.244	52.5 ($j = 105$)	5	25 (25)
2^{14}	4980	11	0.247	52.5 ($j = 105$)	5–6	–
2^{15}	9958	17	0.272	52.5 ($j = 105$)	6	–

Table 5. Squarefree shell, modulus-free statistic at $\rho_r = 1/(4\sqrt{2}\, r^{3/2})$ (floating-point sweep over sorted speeds; numerically verified). The maximiser returned by the sweep for every $r \geq 2^{11}$ is a run of consecutive links of an $e = 1$ near-chain at $\tau^2 = j/2$ (ties are not enumerated).

A memory-lean exact sweep (all pair speeds streamed into a coarse histogram, candidate bins refined exactly, the maximum certified by a level descent or a collective refinement, the speed range split into parts when the occupancy is high) extends the table: on the squarefree shell $D_1^* = 18, 24, 31$ at $r = 2^{16}, 2^{17}, 2^{18}$ (certified exact, also against the float64 speed error: with $\epsilon_r := 2 \cdot 4u_{\mathrm{mach}} \cdot 0.3\sqrt{r}$, $u_{\mathrm{mach}} := 2^{-53}$ the binary64 unit roundoff, a rigorous bound on the error of $(k'-k)/(\sqrt{k'}+\sqrt{k})$, the sweeps at half-widths $\rho_r \mp \epsilon_r$ bracket the true count and both return the value stated; $D^*/r^{1/3} \approx 0.45, 0.47, 0.48$), with maximisers the near-chains $j = 255$ and $j = 1155$ (identified exactly as the drift-free families $A = 1/j$) and, at 2^{18}, the four-class family $A = 7/15$, $C = 8/15$ ($q = 15$, $M = 1$; all 82 members in one window, 31 squarefree); on the prime shell $D_1^* = 4$ at both $r = 2^{19}$ and 2^{20} (certified for the computed speeds; these two values carry no boundary bracket), against $r^{1/3} = 80.6$ and 101.6. The modulus-free census enumerates the drift-free families with $q \leq 120$, $0 < |M| \leq 64$ and α a divisor of $q^2(q^2 - M)/4$ within a factor 6 of $A^* q^2$, together with every symmetric family in the same window, and computes each family's cluster by a float64 speed sweep (Table 6; the pooled boundary brackets and 200-bit rechecks of the leading families are described below); its maximum is a lower bound for $D_1^*(r)$ and equals

the exact value at 2^{18} (at 2^{16} the exact maximiser is a chain with period $q = 510$, outside the census).

$\log_2 r$	census max	family (A, C)	q	n	D_{sym}	ratio
16	16	$(7/15, 8/15)$	15	4	14	1.14
18	31	$(7/15, 8/15)$	15	4	25	1.24
20	45	$(7/15, 8/15)$	15	4	36	1.25
22	78	$(13/105, 212/105)$	105	8	59	1.32
24	145	$(13/105, 212/105)$	105	8	89	1.63
26	172	$(13/105, 212/105)$	105	8	157	1.10
28	280	$(53/105, 52/105)$	105	8	245	1.14
30	521	$(53/105, 52/105)$	105	8	388	1.34

TABLE 6. The largest squarefree cluster among the drift-free two-progression families enumerated (those with $q \le 120$, $0 < |M| \le 64$ and admissible α within a factor 6 of $A^* q^2$), against the symmetric family's; every maximiser has $|M| = 1$ and $n = 2^{\omega(q)}$ classes. The maximum equals the certified exact $D_1^*(2^{18}) = 31$. Counts are float64 census values confirmed at 200-bit precision, except at 2^{26} and 2^{30}, where the 200-bit values are shown.

Four further checks. (a) The boundary certification of the three exact modulus-free maxima just described. (b) A pile-up test inside the census box: the distinct pairs of every enumerated family are pooled (73 855 to 1 625 553 pairs) and the largest window over the union is compared with the largest single-family cluster – they coincide at all eight radii, squarefree $16, 31, 45, 78, 145, 172, 280, 521$ and all pairs $41, 82, 113, 183, 353, 439, 755, 1424$. For squarefree pairs the equality is boundary-certified by tight float64 two-tolerance brackets at $2^{16}, 2^{18}, 2^{20}, 2^{22}, 2^{24}, 2^{28}$ and by 200-bit recomputation of the pooled statistic at 2^{26} and 2^{30}; for all pairs, by tight brackets through 2^{24} and by the 200-bit recomputation at 2^{26} and 2^{30}, while at 2^{28} the all-pairs bracket $[736, 774]$ leaves the value 755 uncertified. Distinct drift-free families share at most one pair, two parabolas $s = Ad^2 + C$ meeting in at most one $d > 0$. (c) At the non-dyadic radii $r = 150\,000$ and $200\,000$ the census box returns 26 and 27, and independent full sweeps (exact for the computed speeds) return the same values. (d) The top three census families and the symmetric family at each of the eight radii, and the eleven symmetric clusters of Table 7, were recomputed at 200 bits; the float64 and 200-bit counts agree except at $r = 2^{26}$ and 2^{30}, where they differ by at most three units and the 200-bit values are the ones quoted throughout: near speed 1, four binary64 unit roundoffs are a tenth of ρ_r at $r = 2^{30}$.

An energy computation enumerates all pairs of the squarefree shell at $r = 2^{12}, \dots, 2^{15}$ ($7.8 \cdot 10^5$ to $5.0 \cdot 10^7$ pairs) with the stable speed formula in float64 and computes the additive energy $E_1(r)$ of Proposition 21.15 and the speed-centred clusters $C(P)$, for the shell and for a phase-randomised null (each d-class shifted by an independent uniform fraction of its local spacing; no boundary certification is claimed for these counts): $E_1/\binom{R}{2} = 1.074, 1.074, 1.074, 1.075$ against $1.073, 1.074, 1.074, 1.074$, i.e. $E_1 = 0.0496\, r^2$ in both, while $\max_P C(P) = 9, 9, 11, 15$ against $4, 5, 5, 5$ and the pairs with $C(P)$ above the null's maximum number $61, 97, 332, 1420$ – between $2.7 \cdot 10^{-5}$ and $7.9 \cdot 10^{-5}$ of all pairs. On the exact starts of the 40- and 48-bit moduli of the planar census (Table 8) the mean of the $N^{1/5}$ largest values of $D(h)$ is $4.2, 3.7, 3.1$ (40 bits, $r = N^{1/3}, 1.44N^{3/11}, N^{1/4}$) and $4.6, 4.1$ (48 bits, $1.44N^{3/11}, N^{1/4}$) against $D_{\max} = 8, 6, 4$ and $10, 6$, ratios 1.3–2.2. Finally, the census of Table 6 with the enlarged box $q \le 1200$, $0 < |M| \le 16$ (846 117 and 874 173 families) returns the same maxima 78 and 145 at $r = 2^{22}$ and 2^{24}, attained by the same $q = 105$ family; the runners-up are the eight-class families of periods 231, 195

and 210 (64, 57, 54 at 2^{22}), and at 2^{24} the four-class family of period 55 (94) followed by the eight-class family of period 210 and the four-class family of period 33 (both 88), so the original box $q \le 120$ lost nothing at these radii.

The symmetric family, enumerated directly with stable arithmetic (speeds as $(k'-k)/(\sqrt{k'}+\sqrt{k})$), has clusters $D_{\mathrm{sym}} = 8, 14, 14, 25, 36, 59, 89, 157, 245, 388, 616$ at $r = 2^{14}, 2^{15}, 2^{16}, 2^{18}, \dots, 2^{32}$, with $D_{\mathrm{sym}}/r^{1/3} \in [0.347, 0.391]$ for $r \ge 2^{18}$, maximising $j = (0.25 \pm 0.04)r^{1/3}$ and squarefree density 0.45–0.48; the full sweep at 2^{16} ($3.96 \cdot 10^8$ pairs) certifies $D^* = 18$, attained by a consecutive chain (the certification does not classify tied or smaller clusters). The density κ_j of Lemma 20.6 matches the measured squarefree-pair density of 200 families at $r = 2^{24}$ (pooled 0.354 against 0.358; $15 \mid j$: 0.456 against 0.460), which is consistent with the observed preference of the maximisers for j divisible by 3 (all eleven) and by 15 (eight of eleven) (the maximising j are $9, 9, 15, 15, 21, 45, 75, 105, 165, 255, 405$). Table 7 compares each cluster with the limiting prediction $2(\sqrt{2}-1)\kappa_{j^*}r^{1/3}$ of Theorem 20.11 at its balance parameter $x = 1$ (the value of $8\theta j^{3/2}/\sqrt{r}$ at which the two bounds in that proof agree): the ratio lies in $[0.979, 1.003]$ at the four largest r and within 10% at every $r \ge 2^{16}$; the outlier at 2^{15} is a family of 25 members, for which the observed 14 squarefree pairs against the expected 10.4 is a 1.4σ binomial fluctuation.

$\log_2 r$	j^*	κ_{j^*}	$2(\sqrt{2}-1)\kappa_{j^*}r^{1/3}$	D_{sym}	ratio
14	9	0.4183	8.8	8	0.91
15	9	0.4183	11.1	14	1.26
16	15	0.4563	15.2	14	0.92
18	15	0.4563	24.2	25	1.03
20	21	0.4365	36.7	36	0.98
22	45	0.4563	61.0	59	0.97
24	75	0.4563	96.8	89	0.92
26	105	0.4762	160.3	157	0.98
28	165	0.4641	248.0	245	0.99
30	255	0.4595	389.8	388	0.995
32	405	0.4563	614.5	616	1.002

TABLE 7. The symmetric family's largest single-window cluster D_{sym} (direct enumeration with stable arithmetic) against the limiting prediction at the balance parameter $x = 1$ with the Legendre density κ_{j^*} of the maximising j^*. The limiting constant is $2(\sqrt{2}-1)\cdot 5/\pi^2 = 0.419686\ldots$, approached as j^* acquires more small prime factors ($\kappa_{105} = 0.476$ for $j^* = 3 \cdot 5 \cdot 7$).

The construction of Theorem 20.8 at $r = 2^{16}, \dots, 2^{28}$ gives 0.20–$0.33\, r^{1/3}$ squarefree pairs in one window against the proven 0.2598.

23.2. **The planar census with real moduli.** Table 8 gives the exact statistic $D_{\max}(\mathcal{F}_1^{\square}(r), W)$ for the whole squarefree shell of the 40-bit moduli $N = 964754165423$ (#0) and 949472385131 (#1) and the 48-bit moduli $N = 173976010050053$ (#0) and 219499552204567 (#1), with, in parentheses, the value for a random-start null (the same number of integer starts drawn uniformly from the range of the shell's starts, with the same window) and the family of the returned maximiser identified by the exact parabola fit of Proposition 20.3 (support ≥ 4).

r	40-bit #0	40-bit #1	48-bit #0	48-bit #1
$N^{1/4}$	4 (4) generic	4 (4) chain $j = 15$	6 (4) chain $j = 15$	5 (5) symmetric $j = 39$
$N^{3/11}$	4 (5) symmetric $j = 17$	4 (5) symmetric $j = 17$	7 (5) $A = 17/6$	8 (5) $A = 1/29$
$N^{0.3}$	7 (5) $A = 1/10$	7 (5) $A = 1/10$	9 (5) $A = 1/14$	7 (5) $A = 27/4$
$N^{1/3}$	8 (5) $A = 1/2$	9 (5) $A = 1/2$	11 (7) $A = 1/2$	11 (6) $A = 1/2$

TABLE 8. Exact-start statistic of the squarefree $a = 1$ shell. Below the crossing the maximisers are cap-regime families (detuning $u\,v' - 1 \approx -1.1$ to -1.7 across their support); from the crossing upward every above-null maximiser is resonant, with the detuning changing sign inside its support and of absolute value at most 0.03 at the median in 19 of 26 instances. At $N^{1/3}$ the maximiser is the family $(m^2 - m + 1, m^2 + m + 1)$ on all four moduli, at median parameter $d_{\mathrm{med}} = 182,182,432,444$ of the cluster against $d_* = (3u)^{1/3} = 180,180,432,444$.

Peeling (removing the identified family and recomputing) reaches the null in 1–6 rounds after removing 0.16–0.67% of the cells; all 26 above-null maximisers were identified with full support; of these, 22 were drift-free with $\Delta_d < 0$ and 4 were non-drift-free small-d families. The prime shell is at the null (2–3) at 40–60 bits except one resonant instance at 56 bits, $r = N^{0.3}$, with six prime pairs. The certified exact sweep extends the exact statistic to 56 and 64 bits (Table 9, the moduli of the resonant census): at eleven of the twelve completed points it equals the resonant census's value (Section 23.3), so at those finite points the census maximum equals the exact all-pairs maximum – which excludes a larger omitted cluster or pile-up at those points only, not tied or smaller ones, nor a mechanism with a larger asymptotic exponent; at the twelfth, the 64-bit squarefree shell at the crossing, the exact value 17 exceeds the census's 16 by a drift-free family $A = 1/62$, $C = 482/31$ whose period $q = 62$ lies outside the census's $q \le 48$ box. The prime shell stays at its null level (3, or 6 at the known resonant instance) through 64 bits.

bits	r	W	full: exact / census	squarefree: exact / census	prime: exact / census
56	$55095 \approx 1.44N^{3/11}$	7	18 / 18	13 / 13	3 / 3
56	$109907 \approx N^{0.3}$	3	26 / 26	20 / 20	6 / 6
56	$399153 = \lfloor N^{1/3} \rfloor$	1	– / 37	18 / 18	3 / 3
64	$220033 \approx 1.44N^{3/11}$	11	29 / 29	17 / 16	3 / 3
64	$504125 \approx N^{0.3}$	4	– / 49	26 / 26	– / 2

TABLE 9. Certified exact values of $D_{\max}$ against the resonant census's maxima on the moduli $N = 63594180015099841$ (56 bits) and $N = 10196702678238263821$ (64 bits). The shells marked – were not swept; the census's values there are lower bounds.

23.3. **The resonant-family census.** Enumerating the drift-free resonant families from Lemma 20.4 ($q \le 48$, $|M| \le 400$, α a divisor of $q^2(q^2 - M)/4$ in the resonance range) reproduces the exact maxima of Table 8 at every parameter point covered by both computations and scales to 96 bits. At $r = N^{1/3}$ the list has 4643–4656 candidate families at every size; the largest full / squarefree / prime clusters found are 10/8/2 and 10/9/2 (the two 40-bit moduli), 18/11/2 and 22/11/2 (48 bits), 37/18/3 (56), 55/27/3 (64), 89/45/4 (72), 128/63/4 (80), 299/145/6 (96) – squarefree $(0.59–0.90)N^{1/12}$, full $(1.0–1.5)N^{1/12}$, lower bounds for the full statistics beyond 48 bits. At the crossing $r = 1.44\,N^{3/11}$ the census gives full / squarefree clusters from 8/6 at 40 bits to 254/171 at 96 bits, i.e. $(0.50–0.68)N^{1/11}$ and $(0.30–0.59)N^{1/11}$, with prime clusters 2–4

through 80 bits and 8 at 96; at $r = N^{0.3}$ the full clusters run from 8 to 230 and the prime ones reach 7. At every tested size from 48 bits onwards the top full-pair family is $(A, C) = (1, 2)$, the family of Theorem 21.2, whose first-order range law $1.19\,N^{1/12}$ predicts $18, 30, 46, 75, 120, 294$ against $18, 37, 55, 89, 128, 299$. At every size there are 42–68 families with cluster $\geq \frac{1}{2}N^{1/12}$ and 490–767 with cluster $\geq \frac{1}{4}N^{1/12}$; the former together comprise $(15 \pm 4)N^{1/6}$ members – a shell fraction falling from 0.30 to 0.00061.

23.4. **The theorems.** Two of the theorems were checked directly on balanced semiprimes of the sizes named, with exact start differences and with the resonance centres and windows evaluated at 256-bit precision.

- Theorem 21.2 at $\lambda = 0.8$ on moduli of $48, 64, 96, 128, 160, 200$ bits: the start differences on the window take 3 distinct values at every size, and the most frequent is shared by 16, 54, 241, 1803, 11674, 106590 of the 28, 71, 460, 3008, 18870, 188112 members – 52–76% against the guaranteed third.
- Theorem 21.4: at $r = \lfloor N^{1/3} \rfloor / C$, $C = 1, 2, 3, 4$ (the $C = 1$ row is an endpoint check of the same mechanism; it lies in the theorem's interval only when N is a cube), the families $(n, m) = (1, 1)$, $(3, 3)$, $(7, 5)$, $(13, 7)$ have $\rho = 3,\ 1,\ 3/7,\ 3/13$ inside $J_w = (1.41, 4)$, $(0.5, 1.41)$, $(0.27, 0.77)$, $(0.18, 0.5)$, stationary points at $z/r = 0.83,\ 0.79,\ 0.68,\ 0.60$, at most three start differences on the window $|t - t_*| \leq 0.8\sqrt{t_*/(3m)}$, and a most frequent value shared by $(0.70, 0.23, 0.18, 0.13)\,N^{1/12}$, $(0.63, 0.17, 0.13, 0.06)\,N^{1/12}$ and $(0.49, 0.27, 0.13, 0.07)\,N^{1/12}$ members at 64, 96 and 128 bits.

23.5. **Achievable explicit covers.** Lemma 4.3 and Theorem 4.4 bound explicit representations $Z' \subseteq B - G$ from below through one mechanism, the approximate-Sidon statistic. They say nothing about which covers are achievable, and structure that the lemma does not see could in principle make small covers possible; covers were therefore also sought directly. For the shifted starts $S = \{s_k = \lceil 2\sqrt{kN} \rceil - N - k : r/2 < k \leq r\}$ of the $a = 1$ shell, minimise $W|B'| + |G|$ over offset sets $B' \ni 0$ and giant sets G with $S \subseteq B' + G$ (so that $Z' \subseteq (B' + [0, W)) - (-G)$), by greedy set cover for G given B' and a local search over B' among the most frequent start differences. The search returns upper bounds on the minimal cover. Two comparisons are built in: the trivial cover ($B' = \{0\}$, $G = S$, cost $W + |S|$), and the same search on a random set of $|S|$ distinct integers in the same range.

bits	r	$\lvert S\rvert$	$D_{\max}$	trivial	best	ratio	control ratio	Lemma 4.3 bound
22	148	74	3	75	53	0.707	0.760	15.4
24	216	108	4	109	84	0.771	0.780	18.0
26	359	180	4	181	146	0.807	0.785	25.3
28	547	274	4	275	231	0.840	0.822	33.5
30	893	447	5	448	392	0.875	0.868	43.1
32	1447	724	7	725	660	0.910	0.912	53.1
34	2229	1115	8	1116	1044	0.935	0.944	67.7
36	3483	1742	8	1743	1662	0.954	0.954	91.2

TABLE 10. Explicit covers of the shifted start set of the $a = 1$ shell at $r = \lfloor N^{1/3} \rfloor$ ($W = 1$), one balanced semiprime at each size ($N = 3\,302\,401$, $10\,142\,569$, $46\,450\,753$, $164\,541\,959$, $713\,555\,207$, $3\,032\,841\,269$, $11\,085\,944\,759$, $42\,281\,311\,687$ at 22 to 36 bits), search budget 12 offsets from the 25 most frequent differences. "Ratio" is best cost over trivial cost; "control ratio" the same for a random set of the same size in the same range. With a budget of 80 offsets from 120 candidates the ratios at 22, 26, 30 bits become 0.680, 0.657, 0.621 for S and 0.693, 0.641, 0.643 for the control. In the window regime $r = \lfloor N^{1/4} \rfloor$ ($W = 2, 2, 3, 3, 4, 4, 5, 6$ at 22 to 36 bits) no offset pays for itself and the best cover found is the trivial one, for S and for the control alike.

The outcome is the same at every size and budget: the shifted start set is as compressible as a random set of the same size and range, to within a few per cent, in both directions. The savings that exist – up to 38% of the trivial cost at the largest budget – are the generic savings of any set whose pairs number a constant multiple of its range, bought by many offsets each saving one or two giants; the clusters that Lemma 4.3 tracks contribute nothing visible beyond that. With windows of length $W \geq 2$ an offset costs W and saves at most $D_{\max} - 1 < W$, and the trivial cover is optimal within the search. The achievable covers sit 3.4–18 times above the bound of Lemma 4.3 and within a constant factor of the trivial cover, and the gap between the two does not visibly close over this sample. This is a statement about what a greedy search finds, not a theorem about all covers; it points the same way as Proposition 21.15: the additive structure of the starts is a tail phenomenon, and the bulk, which is what a cover has to pay for, is random-like.

23.6. **The balanced sub-family.** The balanced census computes exactly, for two RSA moduli at each even size from 30 to 46 bits and one at 48 bits (the moduli of Table 3), at $r = \lfloor N^{1/3} \rfloor$ and $C = 2$, the statistic $D_{\max}(\mathcal{F}, W)$ of Lemma 4.3 over the cell starts of four families – the balanced family $a \leq b \leq 2a$, $ab \leq r$; its primitive cells $\gcd(a, b) = 1$; the primitive cells with $ab > r/4$ (window length 1); and the $a = 1$ shell – each against three random sets of the same size and range, with the lemma's tolerance W equal to the family's maximal window; the values quoted in this paragraph are those of the two moduli at each of $30, 34, 38, 42$ bits unless another size is named. The balanced family has exact statistic $29, 26$; $42, 44$; $73, 67$; $115, 115$ and tolerance statistic $37, 37$; $62, 61$; $96, 94$; $159, 154$, the maximising tolerance centres lying within W of $N + 1 - \lceil 2\sqrt{N} \rceil$ (in the implementation's negated starts $aN + b - \lceil 2\sqrt{abN} \rceil$; simultaneous negation leaves $D_{\max}$ unchanged). The diagonal ray of Lemma 19.1 supplies witness clusters of sizes $29, 30$; $49, 48$; $75, 74$; $124, 119$, that is 77–81% of the reported tolerance maxima: the diagonal controls and, at these samples, dominates the scale of the tolerance statistic, but does not by itself attain the full maximum, and no such attribution is made for the exact-difference statistic. The family's windows reach $W = 8, 13, 20, 32$, and Lemma 4.3 gives $10, 14, 18, 25$, below the trivial $2\sqrt{n'} = 40, 64, 97, 157$. The primitive short-window sub-family has statistic

$3,4$; $5,6$; $5,5$; $9,8$ against 1–2 for random sets – between 0.28 and 0.45 times $r^{1/3}$ on these samples; for comparison the shell values are $6,6$; $7,6$; $12,9$; $14,13$, between 0.45 and 0.67 times $r^{1/3}$, against 5–8 for random sets of the shell's size – and Lemma 4.3 gives $24.8, 22.8$; $40.2, 36.4$; $69.6, 68.9$; $111.2, 109.8$; the singleton starts of all primitive cells ($3,4$; $5,6$; $5,5$; $10,9$ on $204, 209$; $538, 509$; $1226, 1208$; $3315, 3066$ cells) give $30.3, 28.0$; $48.7, 44.2$; $84.4, 83.6$; $130.0, 127.8$, and the energy form of the lemma (with $\bar{\Sigma}_m$ the mean of the m largest exact clusters over all $h \neq 0$ and the bound minimised over $m = \min(|B|, |G|)$, so that it holds for every cover) gives $30.3, 32.1$; $55.6, 52.7$; $92.8, 91.0$; $161.5, 156.0$, the mean of the top $m \approx 15, 26, 43, 73$ clusters being 2.7–5.3; for the shell the energy form gives $44.3, 46.7$; $83.8, 81.6$; $137.7, 136.5$; $252.0, 239.8$. Compare $N^{1/6} = 31, 31$; $50, 49$; $76, 76$; $125, 121$ and $N^{1/5} = 62, 63$; $111, 107$; $182, 180$; $330, 314$. The primitive family's maximiser from 44 bits on is the equal-product structure of Lemma 19.2 ($t = 12N - 18$; $13, 15$; $19, 19$; 25 pairs at $44, 46, 48$ bits). The unique-product family $\mathcal{F}_{\text{uniq}}$ ($201, 206$; $525, 497$; $1180, 1162$; $3149, 2916$ cells at $30, 34, 38, 42$ bits, and $4492, 5035$; $7556, 7120$; 11009 at $44, 46, 48$) has exact statistic 3–9 throughout, tolerance statistic $14, 14$; $24, 22$; $34, 34$; $52, 50$ at its own W, and singleton Lemma-4.3 values $30, 28$; $48, 44$; $82, 81$; $142, 142$ at 30–42 bits, rising to 300 at 48 bits; in every case a maximal exact cluster consists of consecutive cells on the line $2b = 3a + 1$, the unique maximiser at thirteen of the nineteen moduli. This is not in tension with Section 23.5, which concerns direct covers of the $a = 1$ shell; moreover one high-multiplicity difference can control $D_{\max}$ without furnishing a low-cost cover of the primitive remainder. Table 11 collects the unique-product family's sizes and statistics.

bits	r	$\#\mathcal{F}_{\text{bal}}^{\text{prim}}$	$\#\mathcal{F}_{\text{uniq}}$	$D_{\max}(\mathcal{F}_{\text{uniq}}, 1)$	Lemma 4.3	$N^{1/6}$	$N^{1/5}$
30	958; 988	204; 209	201; 206	3; 4	30.0; 27.7	31.0; 31.4	61.5; 62.6
32	1549; 1532	328; 324	322; 319	4; 5	37.3; 34.4	39.4; 39.2	82.1; 81.5
34	2549; 2412	538; 509	525; 497	5; 6	48.0; 43.5	50.5; 49.1	110.6; 107.0
36	3597; 3719	758; 785	736; 763	6; 6	56.5; 57.9	60.0; 61.0	136.0; 138.8
38	5826; 5723	1226; 1208	1180; 1162	5; 5	82.3; 81.4	76.3; 75.7	181.6; 179.7
40	9568; 8794	2018; 1854	1930; 1777	7; 6	102.1; 101.7	97.8; 93.8	244.6; 232.5
42	15728; 14538	3315; 3066	3149; 2916	7; 6	141.5; 141.5	125.4; 120.6	329.6; 314.4
44	22524; 25327	4745; 5338	4492; 5035	7; 7	179.3; 193.5	150.1; 159.1	408.9; 438.7
46	38227; 36007	8059; 7588	7556; 7120	9; 6	233.2; 256.6	195.5; 189.8	561.6; 541.8
48	56099	11820	11009	9	299.8	236.9	706.9

TABLE 11. The unique-product family $\mathcal{F}_{\text{uniq}}(r)$ at $r = \lfloor N^{1/3} \rfloor$, $C = 2$, for the moduli of Table 3: its size (between $0.196\,r$ and $0.210\,r$), its exact statistic, and the bound of Lemma 4.3 for its singleton starts ($W = 1$), against $N^{1/6}$ and $N^{1/5}$. Over the range $\sqrt{r}$ grows from 31 to 237 and $r^{1/3}$ from 9.9 to 38; in every case a maximal exact cluster consists of consecutive cells $(a, b), (a, b-1)$ on the line $2b = 3a + 1$, verified by enumerating every difference attaining the maximum on all nineteen moduli: at thirteen moduli that difference is the unique maximiser, and at the other six between one and seven further differences attain the same maximum with pairs of other shapes.

24. WHAT IS PROVED, WHAT IS CONDITIONAL, WHAT IS COMPUTED, WHAT IS OPEN

This section is a summary, not an exhaustive inventory.

Proved. Lemma 4.3 (the difference-cover statistic bounds every explicit representation) and Theorem 4.4 (its consequence in the covering model, with the conversion $\eta \mapsto (1 - \eta)/(5 - 4\eta)$ of an envelope into a floor). For the $a = 1$ shell: the classification of exact repeated speeds

(Proposition 20.1); the Beatty chains, forcing $D_{\max} \gg \sqrt{r}$ on the full shell for $r \lesssim N^{1/4}$; the near-chains (Proposition 20.2); the two-progression normal form and its integrality conditions (Proposition 20.3, Lemma 20.4); the modulus-free lower bound $D^*_\theta(r) \geq (c_\theta - o(1))r^{1/3}$ with the sharp constant of the symmetric mechanism (Theorems 20.8, 20.11) and the transfer between the exact-start and modulus-free statistics (Lemma 20.9); the class count $n \leq 2^{\omega(q)}$, the long-period bound and the envelope $7\,r^{1/3}$ for every drift-free family with $\Delta_d \neq 0$ (Lemmas 20.14, 20.15, Theorem 20.16); the planar cap (Proposition 21.1); the resonant clusters $\gg N^{1/12}$ at $r = N^{1/3}$ for the full shell at every N (Theorems 21.2, 21.4) and for the squarefree shell for almost all N (Theorem 21.3); the sliding-window bound and its consequences (Theorem 21.5, Corollary 21.6); the planar two-progression envelope and the thin resonant hierarchy (Theorem 21.10, Corollary 21.11); the unconditional upper bounds $6N^{2/9}$ at $r = N^{1/3}$ and $2.1\,r^{2/3}$ for the modulus-free statistic (Theorems 21.13, 21.14); the near-minimal additive energy (Proposition 21.15) and the energy form of Lemma 4.3; and the ceilings of Corollaries 21.7, 21.8 on what Lemma 4.3 can certify through whole shells. For the balanced restriction: the rays (Lemma 19.1), the equal-product pairs (Lemma 19.2) and the half-offset families (Lemma 19.3), which give $D_{\max} \geq rN^{-1/4}/15 - 2$ for every family retaining the constructed cells, in particular for the primitive balanced family.

Conditional. Under Conjecture 19.4, Proposition 19.5 gives $|B| + |G| \gg_C N^{5/28-o(1)}$; under Conjecture 19.6, Corollary 19.7 gives $N^{2/11-o(1)}$; under the envelopes (S″) below the crossing and (P′) beyond it, Theorem 4.4 gives $N^{2/11-o(1)}$, which by Corollary 21.7 is also the most the whole-shell certificate can give; under (T) and (T_P), Theorem 4.4 gives $N^{1/5-\delta}$ for every $\delta > 0$. The application of Lemma 19.3 to the unique-product family is conditional on the constructed cells being its representatives. The extremality conjectures (S‴), (P‴) and the first-order computations – offset resonance, the general-W profile, the cost side of the refined-prime programme – are marked as such where they occur.

Computed. The censuses of Section 23: the modulus-free sweeps and the drift-free family census, the symmetric family, the planar and resonant censuses to 96 bits, the energy and non-drift-free computations, the direct cover search and the balanced censuses, each with its exactness status and its controls.

Open. The upper envelope of $D_{\max}$ – between the proved exponent $2/3$ and the conjectured $1/4$ (as powers of the shell size) at the top of the range, and between $2/3$ and $1/3$ for the modulus-free statistic; whether a single drift-free family carries the maximum up to $o(r^{1/3})$; the excision estimate that Conjecture 19.6 needs (an $O_C(r/M)$ incidence bound for the cells in high-window half-offset configurations, against the provable $O(r^{4/3}/M^2)$); the representative property of the half-offset cells in the unique-product family; and the positive-proportion squarefree problem in short intervals that separates Theorem 21.2 from its squarefree analogue for every individual N. Nothing in this part bounds the cost of factoring: every floor is a statement about the explicit-difference representation model, and inside that model the results leave open representations at the $N^{2/11+o(1)}$ scale, none of which is constructed here.

Part 4. Factoring $N = pq$ from N alone when $\gcd(p-1, q-1)$ is large

Summary. Let $N = pq$ with primes $p < q$ be balanced, $\sqrt{N/C} \leq p < \sqrt{N}$ for a fixed $C > 1$, and let $g := \gcd(p-1, q-1)$. For any unit α with $\alpha^{N-1} \not\equiv 1 \pmod N$, a babystep–giantstep search on $\beta := \alpha^{N-1}$ factors N deterministically in $O(N^{1/4}/\sqrt{g})(\log N)^{O(1)}$ bit operations, using nothing but N. Some $\alpha \leq (\log N)^4$ qualifies unless $g > (p-1)^{3/4}/(\log N)^3$, by a count of smooth numbers inside the subgroup of Fermat liars; in the remaining class Burgess's character-sum bound supplies a base below $N^{1/(8\sqrt{\mathrm{e}})+\varepsilon}$ for every $\varepsilon > 0$. Consequently every balanced semiprime with $g \geq N^{5/24}/(2\sqrt{C}(\log N)^3)$ is

factored from N alone with exponent below $1/5$; writing $g = N^\gamma$, the exponent $1/4 - \gamma/2$ is below $1/5$ for every $\gamma > 1/10$. Composing the McKee–Pinch progression search, made deterministic, with Pollard–Strassen on $N-1$ and Harvey–Hittmeir's order selection gives a curve $E(\gamma)$ that is below $1/5$ for every $\gamma > 1/20$. The exponent $1/4 - \gamma/2$ from N alone is that of McKee and Pinch's Pollard-rho variant with iteration $x \mapsto x^{N-1} + 3$, an expected-time attack restated by Hinek; the present argument is its deterministic form – a babystep–giantstep search with the base chosen deterministically – and needs no knowledge of g.

25. Introduction and context

Throughout, $N = pq$ with distinct primes $p < q$, and for a fixed $C > 1$ the smaller factor lies in the balance range $J_C := [\sqrt{N/C}, \sqrt{N}]$, so that $p, q = \Theta_C(N^{1/2})$; "N large" means large in terms of C. Write

$$g := \gcd(p-1, q-1), \qquad k := \frac{p-1}{g}, \qquad l := \frac{q-1}{g}, \qquad \gcd(k, l) = 1.$$

Moduli with a large g were proposed by Wiener [81] as a defence against his small-exponent attack, and are studied under the name "common prime RSA" by Hinek [37, 38]. Attacks from N alone are part of that literature: with $g = N^\gamma$, McKee and Pinch [58], as restated in [38, Attack 6.1], factor N in $O(N^{1/4-\gamma/2})$ expected operations by a Pollard-rho variant whose iteration $x \mapsto x^{N-1} + 3$ takes at most $(p-1)/g$ values modulo p; given g, their babystep–giantstep search costs $O(N^{1/4-\gamma})$ expected [38, Theorem 6.4], and for known $g \geq N^{1/4}$ the congruence $(N-1)/(2g) \equiv a + b \pmod g$ factors N in polynomial time [38, Theorem 6.3]; g itself is sought by factoring $(N-1)/2$ [38, §6.3.4]. The small-exponent attacks [37] use the public exponent. Here only N is used, g is not assumed known, and the search is deterministic, so that the exponent $1/4 - \gamma/2$ becomes a proven worst-case bound; Section 28 then combines the progression search behind McKee and Pinch's known-g exponent with a deterministic recovery of g from $N-1$ into the curve $E(\gamma)$ of Proposition 28.2. The starting point is elementary: since $N - 1 = (p-1)q + (q-1) = (q-1)p + (p-1)$,

$$\gcd(N-1, p-1) = \gcd(q-1, p-1) = g, \qquad \gcd(N-1, q-1) = g,$$

so for any unit α the element $\beta := \alpha^{N-1}$ has order dividing $(p-1)/g = k$ modulo p and dividing l modulo q, and these two orders are coprime. A collision search on β therefore has to cover only $\min(k, l) \leq \sqrt{kl} < \sqrt{N}/g$ exponents, a Pollard–Strassen or babystep–giantstep cost of $\tilde{O}(N^{1/4}/\sqrt{g})$. The two questions are how to find a base α that is not a Fermat liar ($\beta \not\equiv 1$) deterministically, and what the resulting bound is worth against the deterministic state of the art, $N^{1/5+o(1)}$ [31, 32].

The search step uses two tools from the deterministic-factoring literature: Hittmeir's lemma that an element of exact order m modulo N has the same order modulo every prime factor iff $\gcd(\beta^{m/r} - 1, N) = 1$ for every prime $r \mid m$ [39, Lemma 2.3], and Harvey's collision algorithm [31, Algorithm 4.1], which finds a collision modulo exactly one prime among n giants and m babies in $O(n \log^3 N + m \log^2 N)$ bit operations under the preconditions that $\mathrm{ord}_N(\alpha) > m$ and no giant equals a baby in $\mathbb{Z}/N\mathbb{Z}$. The base-selection step uses a count of smooth numbers in a subgroup of $\mathbb{F}_p^\times$; the same count drives the order-selection theorem of Harvey and Hittmeir [33] and its balanced-semiprime variant in Section 16. Notation: $\Psi(x, y)$ is the number of y-smooth integers in $[1, x]$, $\pi(y)$ the number of primes up to y, ρ the Dickman function, and $\tilde{O}$ hides factors $(\log N)^{O(1)}$.

26. The collision search

Lemma 26.1 (smooth numbers below p)**.** *For N sufficiently large in terms of C and $y := (\log N)^4$,*

$$\Psi(p-1,y) \;>\; \frac{(p-1)^{3/4}}{(\log N)^3}.$$

Proof. Put $u := \lfloor \log(p-1)/\log y \rfloor$; for N large $u \le \pi(y)$, and the products of u distinct primes $\le y$ are distinct y-smooth integers below p, so $\Psi(p-1,y) \ge \binom{\pi(y)}{u} \ge (\pi(y)/u)^u$. Since $\pi(y) > y/\log y$ for $y \ge 17$ and $u \log y \le \log(p-1)$, $\pi(y)/u > (\log N)^4/\log(p-1) \ge (\log N)^3$, while $u \ge \log(p-1)/(4\log\log N) - 1$; therefore $\Psi(p-1,y) \ge (\log N)^{3u} \ge (p-1)^{3/4}/(\log N)^3$. □

Lemma 26.2 (a large common factor of $p-1$ and $q-1$)**.** *Fix $C > 1$. Let $N = pq$ with distinct primes $p < q$, $p \in J_C$, and N sufficiently large in terms of C. For every unit α with $\alpha^{N-1} \not\equiv 1 \pmod N$, a collision search on $\beta := \alpha^{N-1}$ factors N in $O(\sqrt{\min(k,l)})(\log N)^{O(1)} \le O(N^{1/4}/\sqrt{g})(\log N)^{O(1)}$ bit operations. Moreover some integer $2 \le \alpha \le (\log N)^4$ satisfies $\alpha^{N-1} \not\equiv 1 \pmod N$ unless $g > (p-1)^{3/4}/(\log N)^3$. Consequently every N with*

$$\frac{N^{5/24}}{2\sqrt{C}\,(\log N)^3} \;\le\; g \;\le\; \frac{(p-1)^{3/4}}{(\log N)^3}$$

is factored in $N^{7/48+o(1)}$ bit operations from N alone.

Proof. The image of the power map $x \mapsto x^{N-1}$ on the cyclic group $\mathbb{F}_p^\times$ has order $(p-1)/\gcd(p-1, N-1) = k$, so $\operatorname{ord}_p(\beta) \mid k$; similarly $\operatorname{ord}_q(\beta) \mid l$, and the two component orders are coprime. If exactly one of them is 1, then $\gcd(\beta - 1, N)$ is a proper factor; both are 1 exactly when $\beta \equiv 1 \pmod N$, which is excluded. So assume both exceed 1; being coprime they are distinct. Put $e := \min\{\operatorname{ord}_p \beta, \operatorname{ord}_q \beta\}$: the prime of order e divides $\beta^e - 1$ and the other does not, so $\gcd(\beta^e - 1, N)$ is proper, and $e \le \min(k,l) \le \sqrt{kl} = \sqrt{(p-1)(q-1)}/g < \sqrt{N}/g$.

For $D' = 4, 16, 64, \ldots$ put $H := \sqrt{D'}$ and compare the babies β^i, $0 \le i < H$, with the giants β^{jH}, $1 \le j \le H$; the differences $jH - i$ cover $[1, D']$. An exact equality modulo N gives a positive relation exponent $M = jH - i \le D'$; factoring M by trial division up to $\sqrt{M} \le \sqrt{D'}$, within the budget of this round, and stripping its prime divisors gives the exact order $E = \operatorname{ord}_N(\beta)$, and since the component orders differ, [39, Lemma 2.3] gives a prime $r \mid E$ with $\gcd(\beta^{E/r} - 1, N) > 1$, which is proper because E is exact. If there is no exact equality then $\operatorname{ord}_N(\beta) > D' \ge H$ and the separation preconditions of [31, Algorithm 4.1] hold; once $D' \ge e$ the representation $e = jH - i$ is a collision modulo exactly one prime and Harvey's algorithm returns a factor (directly when the gcd of the evaluated product with N is proper, and by scanning the babies once when it is N, a giant agreeing with one baby modulo p and with another modulo q). The first such D' is below $4e$, and the geometric sum of the costs is $O(\sqrt{e})(\log N)^{O(1)} \le O(N^{1/4}/\sqrt{g})(\log N)^{O(1)}$.

Put $y := (\log N)^4$; for N large, $y < p$, so every $2 \le \alpha \le y$ is a unit. If all of them satisfy $\alpha^{N-1} \equiv 1 \pmod N$, every y-smooth integer in $[1, p-1]$ lies modulo p in $H_p := \{x \in \mathbb{F}_p^\times : x^{N-1} = 1\}$, of order $\gcd(N-1, p-1) = g$, so $\Psi(p-1,y) \le g$; Lemma 26.1 then gives $g > (p-1)^{3/4}/(\log N)^3$. Finally, for $g \ge N^{5/24}/(2\sqrt{C}(\log N)^3)$, $N^{1/4}/\sqrt{g} \le (2\sqrt{C})^{1/2} N^{7/48} (\log N)^{3/2}$, and trying the $O((\log N)^4)$ bases adds a polylogarithmic factor. □

The bound $N^{1/4}/\sqrt{g}$ therefore holds for every sufficiently large balanced N with $g \le (p-1)^{3/4}/(\log N)^3$. Its exponent is McKee and Pinch's [58, 38]: their rho variant reaches $N^{1/4-\gamma/2}$ in expected time, and what is new here is the deterministic form, with a proven bound and

no knowledge of g. Knowledge of g would do better – $O(N^{1/4-\gamma})$ expected by their babystep–giantstep search [38, Theorem 6.4], polynomial time for $g \geq N^{1/4}$ [38, Theorem 6.3] – but g is a divisor of $N-1$ that the present method never has to find.

27. The Burgess completion

The remaining class $g > (p-1)^{3/4}/(\log N)^3$ is where the smooth-number count fails to supply a non-liar base; a character-sum bound supplies one instead. We use Burgess's bound in the form: for every prime P, every nonprincipal character χ modulo P, every integer $r \geq 1$ and every $x \geq 1$,

$$\left|\sum_{n \leq x} \chi(n)\right| \ll_r \; x^{1-1/r}\, P^{(r+1)/(4r^2)} \log P$$

[11, Theorem 2], with an implied constant depending on r alone; hence for every fixed $\varepsilon > 0$ and $x \geq P^{1/4+\varepsilon}$ (taking $r > 1/(4\varepsilon)$) the sum is $o(x)$ uniformly in χ. The least k-th power non-residue is treated in this uniform way by Norton [66]; we give the short argument in the form needed.

Corollary 27.1 (the large common-factor class)**.** *Fix $C > 1$ and $\varepsilon > 0$. Let $N = pq$ with $p \in J_C$, $p < q$, and $g > (p-1)^{3/4}/(\log N)^3$. Then N is factored deterministically in $O_{C,\varepsilon}(N^{1/(8\sqrt{\mathrm{e}})+\varepsilon})$ bit operations. Consequently, for every sufficiently large balanced semiprime with*

$$g \;\geq\; \frac{N^{5/24}}{2\sqrt{C}\,(\log N)^3},$$

factoring costs $N^{7/48+o(1)}$ or $O_{C,\varepsilon}(N^{1/(8\sqrt{\mathrm{e}})+\varepsilon})$ bit operations; since $1/(8\sqrt{\mathrm{e}}) = 0.0758\ldots < 7/48 < 1/5$, choosing $\varepsilon < 1/5 - 1/(8\sqrt{\mathrm{e}})$ puts both below $N^{1/5}$.

Proof. The subgroup $H_p := \{x \in \mathbb{F}_p^\times : x^{N-1} = 1\}$ has order $\gcd(N-1, p-1) = g$ and index k; since $\mathbb{F}_p^\times$ is cyclic and $k \mid p-1$, it is the subgroup of k-th powers, and likewise H_q has index l. As $p \neq q$, not both indices are 1; let $P \in \{p, q\}$ be a prime whose subgroup H_P has index $h \geq 2$. The indicator of H_P on $\mathbb{F}_P^\times$ is $\frac{1}{h}\sum_\chi \chi$ over the h characters with $\chi^h = 1$, so for $x \geq P^{1/4+\varepsilon}$ the number of integers in $[1, x]$ lying in H_P modulo P is $x/h + o(x)$ by Burgess's bound, uniformly in h (the $h-1$ nonprincipal sums are each $o(x)$ and are divided by h). Fix u with $1 < u < \sqrt{\mathrm{e}}$ and put $y := x^{1/u}$. If every prime $\leq y$ lay in H_P, so would every y-smooth integer, and then $\Psi(x, y) \leq x/h + o(x) \leq x/2 + o(x)$; but $\Psi(x, x^{1/u}) = (\rho(u) + o(1))x$ [80, Chapter III.5] with $\rho(u) = 1 - \log u > 1/2$ for $1 < u < \sqrt{\mathrm{e}}$, a contradiction for x large. Hence some prime $\alpha \leq y = P^{(1/4+\varepsilon)/u}$ lies outside H_P, i.e. $\alpha^{N-1} \not\equiv 1 \pmod P$ and so $\beta := \alpha^{N-1} \not\equiv 1 \pmod N$; letting $u \uparrow \sqrt{\mathrm{e}}$ and using $P = \Theta_C(N^{1/2})$, some $\alpha \leq N^{1/(8\sqrt{\mathrm{e}})+\varepsilon'}$ qualifies for every $\varepsilon' > 0$ and N large in terms of C and ε'. Enumerate the bases up to this bound, computing $\gcd(\alpha, N)$ and β, and skip those with $\beta \equiv 1$; for the first remaining base the collision search of Lemma 26.2 factors N in $O(\sqrt{\min(k,l)})(\log N)^{O(1)}$ operations, and here $\min(k,l) = k < (p-1)^{1/4}(\log N)^3$ because $k < l$, so this is $N^{1/16+o(1)}$, dominated by the enumeration since $1/(8\sqrt{\mathrm{e}}) > 1/16$. The second statement combines this with Lemma 26.2. □

Remark 27.2 (direct enumeration). In the class of Corollary 27.1 a direct enumeration is also available but slower: $l/k = (q-1)/(p-1) \leq C(1+o(1))$, so $l \leq 2Ck$ for N large, and from $N - 1 = (k+l)g + klg^2$ one has $4kl(N-1) + (k+l)^2 = (2klg + k + l)^2$. Whenever $(k+l)^2 < 4klg + 2(k+l) - 1$ this gives $2klg + k + l = \lfloor\sqrt{4kl(N-1)}\rfloor + 1$, and the inequality holds throughout the class for N large, since $(k+l)^2 \leq (1+2C)^2k^2 \leq 4klg$ as soon as $g \geq (1+2C)^2/4$. So one square test per coprime pair (k, l) with $k < (p-1)^{1/4}(\log N)^3$ recovers g, p and q in $N^{1/4+o(1)}$ operations.

28. A deterministic curve for the class

Lemma 26.2 and Corollary 27.1 give the exponent $1/4-\gamma/2$ for $g=N^\gamma$ without knowledge of g; it is below $1/5$ for $\gamma>1/10$. Two further routes, both elementary once a common divisor of $p-1$ and $q-1$ is known and both known in that form, combine with a deterministic recovery of that divisor from $N-1$ into a curve below $1/5$ for every $\gamma>1/20$.

The Fermat progression. Let g' be any common divisor of $p-1$ and $q-1$ and write $p=1+g'x$, $q=1+g'y$. Then $N-1=g'(x+y)+g'^2xy$, so $M':=(N-1)/g'$ is an integer with $x+y\equiv M'\pmod{g'}$; with $c:=M' \bmod g'$ and $x+y=c+g't$,

$$p+q=r+g'^2t,\qquad r:=2+g'c,\qquad t\ge 0. \tag{1}$$

For balanced p,q with $q/p<C$ the sum $S:=p+q$ lies in $[\lceil 2\sqrt{N}\rceil,\lfloor K_C\sqrt{N}\rfloor]$, $K_C:=\sqrt{C}+1/\sqrt{C}$, so t lies in $[t_-,t_+]$ with $t_-:=\max\{0,\lceil(\lceil 2\sqrt{N}\rceil-r)/g'^2\rceil\}$ and $t_+:=\lfloor(\lfloor K_C\sqrt{N}\rfloor-r)/g'^2\rfloor$, and the number of candidates is $n(g'):=t_+-t_-+1\le (K_C-2)\sqrt{N}/g'^2+2$. Neither $g'=\gcd(p-1,q-1)$ nor $\gcd(x,y)=1$ is used.

Lemma 28.1 (the Fermat cell)**.** *For every integer α, $\alpha^{p+q}\equiv\alpha^{N+1}\pmod N$.*

Proof. Modulo p, $\alpha^p\equiv\alpha$ gives $\alpha^{p+q}\equiv\alpha^{q+1}$ and $\alpha^{pq+1}=(\alpha^p)^q\alpha\equiv\alpha^{q+1}$; symmetrically modulo q; combine by the Chinese remainder theorem. □

This is Harvey's cell $(a,b)=(1,1)$ [31], for which the congruence holds modulo both primes at once.

Proposition 28.2 (a deterministic curve for the common-factor class)**.** *Fix $C>1$. Every balanced $N=pq$ with distinct odd primes and $g:=\gcd(p-1,q-1)=N^\gamma$ is factored deterministically from N alone in $N^{E(\gamma)+o(1)}$ bit operations, where*

$$E(\gamma)=\begin{cases}1/4-\gamma & 0\le\gamma\le 1/12,\\ 1/8+\gamma/2 & 1/12\le\gamma\le 1/8,\\ 1/4-\gamma/2 & 1/8\le\gamma\le 1/6,\\ 1/2-2\gamma & 1/6\le\gamma\le 1/5,\\ \gamma/2 & 1/5\le\gamma\le 1/4,\\ 1/4-\gamma/2 & 1/4\le\gamma<3/8,\\ 1/(8\sqrt{\mathrm{e}})+\varepsilon & 3/8\le\gamma<1/2,\end{cases}$$

for every fixed $\varepsilon>0$ in the last branch, with implied constants depending on C and ε; the $N^{o(1)}$ is the divisor count $d(N-1)$. In particular $E(\gamma)<1/5$ for every $\gamma>1/20$, while $1/4-\gamma/2$ is below $1/5$ only for $\gamma>1/10$; the deterministic record for the class is $\min\{1/5,E(\gamma)\}$ [31]. The curve is not monotone: its maximum over $\gamma\ge\gamma_0$ is $1/4-\gamma_0$ for $\gamma_0\le 1/16$ and $3/16$ for $1/16\le\gamma_0\le 1/8$.

Proof. Three routes are run together with the dovetailing described at the end, so the algorithm needs neither g nor γ.

Route A is Lemma 26.2 with Corollary 27.1, whose cost is

$$O(N^{1/4-\gamma/2})(\log N)^{O(1)} \quad \text{when } g\le(p-1)^{3/4}/(\log N)^3, \text{ i.e. for every fixed } \gamma<3/8,$$
$$O_{C,\varepsilon}\big(N^{\max\{1/4-\gamma/2,\ 1/(8\sqrt{\mathrm{e}})\}+\varepsilon}\big) \quad \text{in general,}$$

the second exponent being the Burgess base search; for $\gamma\ge 3/8$ that is the last branch (at $\gamma=3/8$ itself $g=N^{3/8}$ exceeds $(p-1)^{3/4}/(\log N)^3$ by the logarithmic factor, so the general bound is the one that applies).

Route B (given a common divisor $g' \le N^{1/4}$). Let $n := n(g')$, $m := \lceil\sqrt{n}\rceil$ and $S_0 := r + g'^2 t_-$, and let α be a unit modulo N with $\mathrm{ord}_N(\alpha) > g'^2 m$: Harvey and Hittmeir [33, Theorem 1.1] supply it, or factor N, for $D := g'^2 m$, which satisfies $1 \le D < N-1$ for N large since $D = O_C(g'N^{1/4}) = O_C(N^{1/2})$, at cost $O(\sqrt{D})(\log N)^{O(1)} = O_C(N^{1/8}\sqrt{g'})(\log N)^{O(1)}$. Put $\beta := \alpha^{g'^2}$ and $\eta := \alpha^{N+1-S_0} \bmod N$. By Lemma 28.1 the true shifted index $u := t - t_- \in [0, n)$ satisfies $\beta^u \equiv \eta \pmod N$. Compute the babies β^i, $0 \le i < m$, and the giants $\eta\beta^{-mk}$, $0 \le k < \lceil n/m\rceil$, sort them and match modulo N; a match (i, k) with $u = i + mk < n$ gives $S = S_0 + g'^2 u$, accepted if and only if $S^2 - 4N$ is a perfect square, in which case $p, q = (S \mp \sqrt{S^2 - 4N})/2$. Since $\mathrm{ord}_N(\beta) \ge \mathrm{ord}_N(\alpha)/g'^2 > m$, the babies are pairwise distinct modulo N, each giant matches at most one baby, and the matched u form the residue class of the true one modulo $\mathrm{ord}_N(\beta)$ intersected with $[0, n)$ – at most $n/m + 1 \le m + 1$ of them, each tested in $O(1)$ operations. The cost is $O(m + n/m) = O_C(N^{1/4}/g')$ group operations plus the selection, i.e. $N^{\max\{1/4-\gamma',\ 1/8+\gamma'/2\}+o(1)}$ for $g' = N^{\gamma'}$.

Route C (given g'). For $t = t_-, \dots, t_+$ put $s := c + g't$ and $xy := (M' - s)/g'$, and accept t if and only if $z^2 - sz + xy$ has integer roots x, y with $1 + g'x$ dividing N; $O(1)$ operations per trial, $O_C(\sqrt{N}/g'^2 + 1)$ in all. For $g' \ge N^{1/4}$ there are $O_C(1)$ candidates, and when $g' > x + y$ there is one: Hinek's Theorem 6.3 [38] in this notation.

Finding g. Every prime factor of g is at most g. For a bound B, Pollard–Strassen [69, 78] computes the B-smooth part of $N - 1$ with its factorisation in $O(\sqrt{B})(\log N)^{O(1)}$ operations, and g divides it once $B \ge g$; its divisors number at most $d(N-1) = N^{o(1)}$. Dovetail on a budget 2^j, $j = 1, 2, \dots$: for every dyadic range $[2^s, 2^{s+1})$ whose cost bound,

$$\begin{aligned} &\max\{N^{1/4}/2^s,\ N^{1/8}2^{(s+1)/2},\ 2^{(s+1)/2}\} && \text{for Route B, when } 2^s \le N^{1/4},\\ &\max\{\sqrt{N}/4^s + 1,\ 2^{(s+1)/2}\} && \text{for Route C,} \end{aligned}$$

is at most 2^j (candidates above $N^{1/4}$ in the range straddling $N^{1/4}$ are passed to Route C), compute the 2^{s+1}-smooth part of $N - 1$ and run the route on every divisor g' of it in the range, interleaving the steps of Route A. Every output is verified, so a false candidate cannot return a false factor. The work at budget j is $O(\log N)\, d(N-1)\, 2^j (\log N)^{O(1)} = N^{o(1)}2^j$, the budgets double, and the range containing g is run as soon as 2^j exceeds its cost bound; the total is therefore $N^{o(1)}$ times the least of the three routes' costs at γ. That minimum is $E(\gamma)$: Route A gives $1/4 - \gamma/2$ for $\gamma < 3/8$ and the plateau from $3/8$ on; Route B gives $\max\{1/4 - \gamma,\ 1/8 + \gamma/2\}$ for $\gamma \le 1/4$; Route C gives $\max\{\gamma/2,\ 1/2 - 2\gamma\}$ with $1/2 - 2\gamma$ read as 0 for $\gamma \ge 1/4$, the $\gamma/2$ being the smooth-part cost; the breakpoints are $1/12$, $1/8$, $1/6$, $1/5$ and $1/4$, where consecutive branches agree, and $3/8$, where the curve steps up from $1/16$ to the plateau $1/(8\sqrt{\mathrm{e}}) = 0.0758\dots$. □

What is new here. The known-g exponent $1/4 - \gamma$ of Route B is McKee and Pinch's [58], as restated in [38, Theorem 6.4], and the progression (1) is elementary; Route C for $g \ge N^{1/4}$ is [38, Theorem 6.3]; the order selection is Harvey and Hittmeir's; Pollard–Strassen and divisor enumeration are standard. What we have not found stated is the composition: the progression search realised as an exact babystep–giantstep collision, the recovery of g from the smooth part of $N - 1$ inside the same budget – Hinek [38, §6.3.4] seeks g by the elliptic curve method or the number field sieve – and the resulting N-only curve, obtained with no knowledge of γ. Route B does not realise the known-g exponent for every γ: the order selection dominates for $\gamma > 1/12$.

29. An implementation check

The collision search of Lemma 26.2 was implemented as stated – the doubling bounds $D' = 4, 16, \dots$ with $H = \lfloor\sqrt{D'}\rfloor + 1$ babies and $\lceil D'/H\rceil$ giants, exact-order recovery from an exact

match, and Harvey's Algorithm 4.1 realised by a product tree and multipoint evaluation modulo N with the baby scan when the aggregate gcd is N – and run on 240 moduli $N = pq$ constructed as $p = 1 + kg$, $q = 1 + lg$ with $g = 2\ell$ for a prime ℓ, $\gcd(k, l) = 1$, $k < l < 2k$ (so $q/p < 2$), k and g each of about a quarter of the bits of N, and N of 36 to 44 bits, by rejection sampling from a seeded generator; the values of k ranged over $[259, 2014]$ and those of g over $[334, 2042]$. The size of k matters for the check: with k or l equal to 1 every non-liar base factors N at the immediate step $\gcd(\beta - 1, N)$ and the search never runs. All 240 moduli were factored, 239 of them by the collision search and one at the immediate gcd (a component order equal to 1); the base $\alpha = 2$ was a non-liar in every case, as expected since the liars form a subgroup of density $1/k$ modulo p. For the successful base the two component orders were recomputed from the known factors: the least exponent $e = \min(\mathrm{ord}_p\,\beta, \mathrm{ord}_q\,\beta)$ was at most 2014, and the ratio $e/(\sqrt{N}/g)$, which Lemma 26.2 bounds by 1, had maximum 0.9955 and mean 0.50 over the 239 collision cases, exceeding 0.9 in 21 of them. This checks the implementation and the inequality $e < \sqrt{N}/g$ on the tested population, not the asymptotics; the moduli, seed, per-modulus orders, exponents and search horizons are archived with the report.

Routes B and C of Section 28 were run, with g supplied, on 48 constructed balanced moduli of 33 to 56 bits with $\gcd(p - 1, q - 1) = N^\gamma$, $0.024 \le \gamma \le 0.312$, and the base $\alpha = 2$ (no baby collision occurred, so no other base was needed). Route B factored all 48, with babies plus giants between 1.01 and 1.92 times $\sqrt{n}$ on the 20 moduli with $n \ge 100$ candidates – an object count, not a count of bit operations – and Route C factored all 48 within n trials; 24 reruns with a proper common divisor in place of the gcd also factored their moduli. The recovery of g, the dovetailing and the order selection were not run; at these sizes they do not arise.

30. Where the class arises

The class of moduli with a large common factor is the exceptional set of a conditional procedure in Part 1, whose Proposition 7.3 states: for a balanced semiprime $N = pq$ that survives a Fermat preprocessing of $N^{1/6+o(1)}$ square tests, and under a uniform evaluation hypothesis of exponent γ for the interval sub-products of a certain product of $N^{1/3+o(1)}$ factors, a breadth-first bisection of that product factors N in $(1 + Z_*)N^{\gamma/3+o(1)}$ ring operations by trying the bases $\alpha = 2, 3, \ldots$ up to $(\log N)^4$, *unless* $\gcd(p - 1, q - 1) \ge N^{5/24}/(2\sqrt{C}(\log N)^3)$; here Z_* is a width parameter of that product. Lemma 26.2 and Corollary 27.1 dispose of the exception from N alone: every balanced semiprime with $g \ge N^{5/24}/(2\sqrt{C}(\log N)^3)$ is factored in $N^{7/48+o(1)}$ operations when $g \le (p-1)^{3/4}/(\log N)^3$ and in $O_{C,\varepsilon}(N^{1/(8\sqrt{\mathrm{e}})+\varepsilon})$ operations otherwise. Consequently a conditional algorithm of exponent t for the non-exceptional moduli extends to all balanced semiprimes with exponent $\max(t, 7/48)$ (up to $o(1)$), which equals t whenever $t \ge 7/48$; for a target exponent below $7/48$ the intermediate range $N^{5/24-o(1)} \le g \le (p-1)^{3/4}/(\log N)^3$ would contribute its $N^{7/48+o(1)}$ term, and the class needing separate treatment would be exactly that range.

Part 5. Exposure labels and finite separating certificates for a deterministic fixed-list elliptic curve method

Summary. For $N = pq$ with distinct odd primes p and q, a fixed list of elliptic curves, run with a stage-1 bound B_1 and an exact stage 2 to B_2, factors N as soon as one curve behaves differently modulo p and modulo q. We make this precise through an *exposure label* – for each curve and prime, the residual order of the stage-1 point if it is at most B_2, or one of three degenerate construction states, or no exposure – and prove that a curve factors N if and only if the two labels differ, so that distinct label vectors over a set of primes form a finite certificate that the list factors every product of two distinct primes from the set. Such certificates are computed for every prime of $[2^{18}, 2^{19})$, $[2^{20}, 2^{21})$ and

$[2^{22}, 2^{23})$ with ten, eleven and twelve Suyama curves. Under the conjecture that an explicit family separates in logarithmically many curves, the method with stage-1 bound $B_1 = x^{1/u}$ and $B_2 = B_1^2$ is a deterministic $N^{1/(2u)+o(1)}$ algorithm for products of two primes with bounded ratio; heuristically the route costs $L_N[1/2, 1] = \exp((1+o(1))\sqrt{\log N \log\log N})$ and cannot compete with the number field sieve. Unconditionally, Pollard's $p-1$ method with base 2 succeeds on a positive proportion of the semiprimes with both prime factors in $(x^{1-\varepsilon}, x]$, in $O_\theta(N^{\theta/2})$ multiplications for every $\theta > 1/(2\sqrt{\mathrm{e}})$ and small enough ε: a success theorem for a partial algorithm, not an improvement of any worst-case guarantee.

31. Introduction

The elliptic curve method [50] factors $N = pq$ when a random curve has a smooth order modulo one of the primes. It is analysed as a randomised algorithm for each target prime. Here we ask a different question: for a *fixed* list of curves, chosen without knowledge of N, which pairs (p, q) does the list factor, and how long must the list be to factor every pair of primes in a range? To our knowledge the best proven deterministic exponent is Harvey's $N^{1/5+o(1)}$ [31]; a fixed curve list with a logarithmic separating length would give, for semiprimes with bounded prime ratio, $N^{1/6+o(1)}$ at $u = 3$ by a mechanism unrelated to Harvey's, and N^ε for every fixed ε under the conjecture stated below. We are not aware of any asymptotic covering or separating theorem for an explicit fixed curve family of this kind, unconditionally or under GRH; the rigorous analyses of the elliptic curve method [50, 54] average over a random curve.

The observation that organises this part is that coverage is the wrong criterion. A curve that succeeds modulo both primes with the same residual order gives gcd $= N$; what factors N is a *difference* between the two primes' behaviour, and the right bookkeeping is a label per curve and prime fine enough that equal labels mean exactly "no one-sided relation".

The results are as follows. The label criterion (Proposition 32.4): a curve factors N if and only if the two labels differ, and a deterministic scan of the stage-2 atoms in key order may stop at the first two-sided atom. The finite certificates (Table 12): the first ten Suyama curves at $(B_1, B_2) = (64, 4096)$ separate every prime of $[2^{18}, 2^{19})$, the first eleven at $(102, 10321)$ every prime of $[2^{20}, 2^{21})$, and the first twelve at $(161, 26008)$ all 268 216 primes of $[2^{22}, 2^{23})$, while the collapsed binary signatures of the same predicate need 46–79 curves on the four smaller intervals and had not separated the largest after 40; at $u = 4$ ($B_1 = x^{1/4}$, $B_2 = x^{1/2}$) the separating lengths on all primes of $[x, 2x)$ are $31, 32, 36, 53$ at $x = 2^{16}, 2^{18}, 2^{20}, 2^{22}$, with coverage lengths $33, 35, 41, 53$. The implementation factors all 616 test pairs in at most three curves, and a counter of its ladder, differential-addition, inversion and product multiplications (the polynomial arithmetic excluded) has fitted exponent 0.175 against the predicted 1/6 up to logarithmic factors. Conjecture 33.1 (logarithmic separating length for an explicit family) gives the deterministic $N^{1/(2u)+o(1)}$ algorithm for bounded prime ratio and every fixed u; the route costs $L_N[1/2, 1] := \exp((1 + o(1))\sqrt{\log N \log\log N})$ heuristically; a non-uniform deterministic $L_N[1/2, 1]$ algorithm exists, and the two-sided degeneracy of the $p-1$ method is recovered by the binary Miller chain iff the 2-adic valuations of the two orders differ (Proposition 35.2). For a random list, pairwise separation in $O(\log x)$ curves follows from two-sided exposure-probability bounds with CRT independence and a Deuring–Lenstra anti-concentration estimate (Proposition 33.2); the content of the conjecture is explicitness. Finally, Proposition 35.3: for every $\theta \in (1/(2\sqrt{\mathrm{e}}), 1/2)$ and every $\varepsilon \in (0, 1/4)$ with $\theta(1-\varepsilon) > 1/(2\sqrt{\mathrm{e}})$, Pollard's base-2 $p-1$ method succeeds on a positive relative proportion of the semiprimes $N = pq$ with $x^{1-\varepsilon} < p, q \le x$ in $O_\theta(N^{\theta/2})$ modular multiplications, from the theorems of Friedlander and Fouvry on shifted primes; letting ε and the exponent slack tend to zero jointly gives exponents arbitrarily close to $1/(4\sqrt{\mathrm{e}}) = 0.151632\ldots < 1/5$.

32. Exposure labels

Fix $B_1 < B_2$ and let $L = \operatorname{lcm}\{\ell^e \le B_1\}$. For a Suyama parameter σ [10], the Montgomery curve [60] $By^2 = x^3 + Ax^2 + x$ with $a_{24} = (A+2)/4$ and its starting point are constructed from $U = \sigma^2 - 5$ and $V = 4\sigma$ through the denominator $16U^3V$; stage 1 computes $Q = [L]P$ on the x-line.

Definition 32.1 (exposure label)**.** For a prime p, set $\lambda_\sigma(p) = -1$ if $16U^3V \equiv 0 \pmod p$; otherwise set $\lambda_\sigma(p) = -2$ if $a_{24} \equiv 0 \pmod p$ and $\lambda_\sigma(p) = -3$ if $a_{24} \equiv 1 \pmod p$ (the curve $A = \mp 2$ is singular); otherwise set $\lambda_\sigma(p) = d$, the order of Q modulo p, if that order satisfies $d \le B_2$ (so $d = 1$ is a stage-1 success), and $\lambda_\sigma(p) = 0$ if the order exceeds B_2.

Proposition 32.2 (signatures factor pairs)**.** *Let $\mathcal{P}$ be a finite set of primes and $A_1, \dots, A_K$ deterministic ring computations such that if A_i succeeds modulo exactly one of p, q then its gcd step factors $N = pq$; write $s_i(p) = 1$ if A_i succeeds modulo p and $s_i(p) = 0$ otherwise, and $s(p) := (s_i(p))_{i \le K} \in \{0,1\}^K$. If the signatures $s(p)$ are distinct for every pair of distinct primes $p, q \in \mathcal{P}$, then the K computations factor every pq with distinct $p, q \in \mathcal{P}$ (a perfect-square test handles $p = q$; a computation that succeeds modulo both primes returns* $\gcd = N$ *and is passed over). Coverage is neither necessary nor sufficient for this.*

Proof. Distinct signatures supply an i with $s_i(p) \ne s_i(q)$, hence a one-sided gcd. A set in which one prime has the all-zero signature and all others are distinct is separated but not covered; two primes with the same nonzero signature are covered but not separated. □

Remark 32.3. Every binary separating family has $K \ge \lceil \log_2 |\mathcal{P}| \rceil$, so logarithmic length is optimal up to a constant factor (labels with more than two outcomes can beat the binary bound). For ECM the conservative binary success bit collapses the construction-denominator and final projective-Z paths into one coordinate; analogous collapses define the bit for the $p \pm 1$ methods.

Proposition 32.4 (the label criterion)**.** *Let p and q be distinct odd primes, $N = pq$, and let the deterministic computation on N take the gcds of the denominator, of a_{24} and of $a_{24} - 1$, of the stage-1 Z-coordinate, and run a stage 2 that tests exactly the relations $[v]Q = O$ for $1 \le v \le B_2$ through the x-coordinate atoms described in Section 34: identity atoms $[v]Q = O$, taken first, and collision atoms $x([tm]Q) = x([r]Q)$, $1 \le r < m$, $1 \le t$, formed only from points that are the identity modulo neither prime. In the projective form $X_1Z_2 - X_2Z_1$ a point with $Z \equiv 0$ modulo a prime is the identity there whatever its X – the ladder returns the identity as $(X : 0)$, and as $(0 : 0)$ when it is reached as $[j]Q = O$ with $x([j-1]Q) = x(Q) = \pm 1$, i.e. from a point of order 4 with $[2]Q = (0,0)$ – so the identity atoms detect every such point: one that is the identity modulo one prime only is a proper factor, and one that is the identity modulo both is excluded from all subsequent collision atoms. For the remaining points, whose Z is a unit modulo both primes, a collision atom vanishes modulo a prime exactly when the affine x-coordinates agree there, i.e. when $[tm]Q = \pm[r]Q$, i.e. when $[tm + r]Q = O$ or $[tm - r]Q = O$ modulo that prime, since on a Montgomery curve the x-coordinate determines a point up to sign. Then the curve factors N if and only if $\lambda_\sigma(p) \ne \lambda_\sigma(q)$; moreover, if each identity atom $[v]Q = O$ is assigned the key v and each collision atom $x([tm]Q) = x([r]Q)$ the key $tm + r$, an implementation scanning the atoms in nondecreasing key order may stop at the first vanishing two-sided atom.*

Proof. If the negative labels differ, or if exactly one of the two labels is negative, one of the three construction gcds is proper. If both labels are non-negative and differ, let D_p, D_q be the actual orders of the two stage-1 points. Since the labels differ, at least one of D_p, D_q is at most B_2; let $a \le B_2$ be the smaller exposed order and $e > a$ the other actual order (either both are at most B_2 and unequal, or the label on the e-side is 0 and $e > B_2$). Write $a = tm + r$ with

$m = \lfloor\sqrt{B_2}\rfloor + 1$ and $0 \le r < m$. The atom at key a – the baby identity $[a]Q = O$ if $t = 0$, the giant identity $[tm]Q = O$ if $r = 0$, or the collision $x([tm]Q) = x([r]Q)$ – vanishes modulo the a-side; on the other side it tests only the relation indices a and $tm - r < a$, neither a multiple of $e > a$, so it does not vanish: the atom is one-sided. Every atom of key smaller than a tests only relation indices smaller than a, hence vanishes on neither side, so the first vanishing atom in key order is one-sided and the scan may stop at the first vanishing two-sided atom. Conversely, equal labels yield only two-sided outcomes: the same construction degeneracy occurs on both sides; no relation occurs when both orders exceed B_2; or, when $d_p = d_q$, each tested relation has the same truth value on the two sides ($[v]Q = O$ holds on both exactly when $d_p \mid v$; every vanishing identity atom is then two-sided and its point excluded, and every remaining point has a unit Z on both sides, so each collision atom vanishes on both sides or on neither). □

Distinct label vectors $(\lambda_{\sigma_i}(p))_i$ over $\mathcal{P}$ therefore certify that the list factors every product of distinct primes of $\mathcal{P}$. The collapsed bit $[\lambda \neq 0]$ of Proposition 32.2 is sufficient but discards most of the information. The stopping clause is for a scan in nondecreasing key order; a preparatory scan of the identity atoms in a different order (Section 34) must not stop at a two-sided identity, and must exclude its point from the collision atoms.

33. How long must the list be?

33.1. **The random-signature law.** If independent random curves succeed with probability θ on each prime, one binary coordinate fails to distinguish a fixed pair with probability $R(\theta) = \theta^2 + (1-\theta)^2$, so the expected number of unresolved pairs after K curves is $\binom{n}{2}R(\theta)^K$ and the separation threshold is $K \approx 2\ln n/(-\ln R(\theta))$. Since $R(\theta) \ge 1/2$ with equality at $\theta = 1/2$, a binary coordinate removes the largest fraction of unresolved pairs when its success rate is one half – for a *separating* family the per-member success rate should be tuned to $1/2$, not maximised, and near-fair signatures separate n primes in about $2\log_2 n$ members. For a label with distribution (π_s) the collision probability is $R_{\text{label}} = \sum_s \pi_s^2$, and a larger alphabet can separate faster without changing the group operations. These are heuristics, checked below.

Conjecture 33.1 (E, explicit ECM separation)**.** *For fixed $u > 1$ and $C > 1$ put $B_1 = x^{1/u}$, $B_2 = B_1^2$. There is $C_{u,C}$ such that the first $C_{u,C}\log x$ Suyama parameters, run with stage 1 and the exact stage 2, have distinct exposure-label signatures on all primes in $[x, Cx]$ for all sufficiently large x.*

If true, the list factors every $N = pq$ with distinct primes and $q/p \le C$ deterministically in $N^{1/(2u)+o(1)}$ ring operations: choose $x = \sqrt{N/C}$, so both factors lie in $[x, Cx]$; stage 1 costs $B_1^{1+o(1)}$ and a stage 2 that evaluates the product of the baby polynomial at all giants by one product tree and one quasi-linear multipoint evaluation costs $B_2^{1/2+o(1)} = B_1^{1+o(1)}$. At $u = 3$ this is $N^{1/6+o(1)}$. Taken for every fixed u, the conjecture gives, in this bounded-ratio setting, a deterministic N^{ε} algorithm for every fixed $\varepsilon > 0$ (one algorithm per ε), so it is deliberately strong; finite tests can expose an exceptional-prime tail but cannot establish it. The conjecture concerns the unbounded list; the implementation of Section 34 caps its list and evaluates stage 2 in fixed blocks, and neither restriction belongs to the asymptotic statement.

33.2. **Randomised lists.** Conjecture 33.1 asks for an *explicit* list. For a random list, separation follows from bounds on the exposure probabilities alone. Let μ be the distribution that draws a, x_0, y_0 uniformly from $[1, M]$ with $M \ge x^3$ and forms the curve $y^2 = x^3 + ax + b$, $b := y_0^2 - x_0^3 - ax_0$, with the point $P = (x_0, y_0)$ (Lenstra's ensemble [50]); its labels are defined as in Section 32 with the singular curves as the degenerate state. For a prime p write θ_p for the probability that a curve drawn from μ has a non-zero label at p.

Proposition 33.2 (randomised separation)**.** *Fix $u > 1$ and a constant $\theta \in (0, 1/2]$, let $B_1 = x^{1/u}$, $B_2 = B_1^2$, and suppose that every prime $p \in [x, 2x)$ satisfies $\theta \le \theta_p \le 1 - \theta$. Then for $x \ge x_0(u, \theta)$, $K := \lceil 3 \log x/(\theta(1-\theta)) \rceil$ independent draws from μ have pairwise distinct exposure-label vectors on all primes of $[x, 2x)$ with probability at least $1 - 1/x$.*

Proof. Fix distinct primes $p, q \in [x, 2x)$. The label at p depends only on $(a, x_0, y_0) \bmod p$ and the label at q only on the reduction modulo q; since $M \ge x^3 \ge pq \cdot x/4$, the joint distribution of the two reductions is within total variation $O(1/x)$ of the product of uniform distributions (Chinese remainder theorem), so the two labels are independent up to that error. Write $\pi_p(s)$ for the probability of the label s at p. Then

$$\Pr[\lambda_p = \lambda_q] \;\le\; (1-\theta_p)(1-\theta_q) + \pi_p(1)\pi_q(1) + \sum_{s \ge 2} \pi_p(s)\pi_q(s) + \Pr[\text{both degenerate}] + O(1/x).$$

The degenerate term is $O(1/x^2)$. If the label at p is $s \ge 2$ then, for every prime $\ell \mid s$, $\ell^{v_\ell(L)+v_\ell(s)}$ divides the order of P and hence $\#E(\mathbb{F}_p)$, so $m_s := s \prod_{\ell \mid s} \ell^{v_\ell(L)}$ divides $\#E(\mathbb{F}_p)$, and $m_s > B_1$ because $\ell^{v_\ell(L)+1} > B_1$ by the definition of L. Let $\mathcal{S}_m(p) := \{n : |n - (p+1)| \le 2\sqrt{p},\ m \mid n\}$, so that $|\mathcal{S}_m(p)| \le 4\sqrt{p}/m + 1$. We use the following estimate, which follows from the ingredients of Lenstra's Proposition (1.16) [50] – Deuring's theorem, which expresses the number of pairs (a, b) with a given order through a Kronecker class number, and the class-number bound $H(\Delta) \ll |\Delta|^{1/2}(\log|\Delta|)(\log\log|\Delta|)^2$ – applied to every order in the Hasse interval (Lenstra states the proposition for $|n-(p+1)| < \sqrt{p}$, a restriction his lower bound needs and the upper bound does not), and transferred from uniformly random curves to the curve-with-point ensemble μ, whose weighting by the chosen point changes each probability by a factor $1 + O(p^{-1/2})$ and whose finite box adds a discrepancy $O(p/M)$: the probability under μ that m divides $\#E(\mathbb{F}_p)$ is at most $C(\log p)(\log\log p)^2(1/m + p^{-1/2})$. Hence $\pi_p(s) \le C(\log p)(\log\log p)^2(B_1^{-1} + p^{-1/2})$ uniformly in $s \ge 2$, and $\sum_{s \ge 2} \pi_p(s)\pi_q(s) \le \max_{s \ge 2} \pi_p(s) = o(1)$. Since $\pi_p(1) \le \theta_p$, the remaining terms are at most $(1-\theta_p)(1-\theta_q) + \theta_p\theta_q = 1 - \theta_p - \theta_q + 2\theta_p\theta_q$, a bilinear function whose maximum over $\theta_p, \theta_q \in [\theta, 1-\theta]$ is $1 - 2\theta(1-\theta)$. So one draw fails to separate p from q with probability at most $1 - \theta(1-\theta)$ for $x \ge x_0(u, \theta)$, chosen so that the $o(1)$ terms total at most $\theta(1-\theta)$; K independent draws fail with probability at most $\exp(-K\theta(1-\theta)) \le x^{-3}$, and a union bound over the fewer than x^2 pairs gives the claim. □

The hypothesis is two-sided: a lower bound on the exposure probability is the hypothesis on smooth numbers in the Hasse interval under which the running time of ECM is analysed [50], and an upper bound could come, for example, from a positive proportion of point orders with a prime factor above B_2 (other residual orders exceeding B_2 are possible too), which is open for the same short-interval reasons; heuristically $\theta_p \approx 0.45$ at $u = 3$ for this ensemble [3], and the measured Suyama rates are 0.69–0.81 (Table 12). Conditionally on the hypothesis, the resulting non-uniform curve-list advice has $O(\log x)$ curves and $O(\log x \log M)$ bits – $O(\log^2 x)$ bits for $M = x^{O(1)}$; what Conjecture 33.1 adds is explicitness (and the passage to the Suyama family, for which the same argument needs the analogue of the class-number bound).

33.3. **Full-population computations.** For every prime of the dyadic intervals $[x, 2x)$, $x = 2^{14}, 2^{16}, 2^{18}, 2^{20}, 2^{22}$ ($n = 1612, 5709, 20390, 73586, 268216$ primes), the labels of the Suyama curves $\sigma = 6, 7, \ldots$ were computed exactly at $B_1 = \mathrm{round}(x^{1/u})$, $B_2 = \mathrm{round}(x^{2/u})$ by a vectorised Montgomery ladder. The computation refines the partition of the primes by label prefix curve by curve, and computes a later curve's label only for the primes whose class in the exact partition or in the binary partition is not yet a singleton; since classes only refine, this determines the coverage and separating lengths exactly without computing the whole label matrix. The label routine is checked in the accompanying test suite against scalar point orders and, end to

end, against the algorithm (one curve factors pq iff the labels differ, on every pair of consecutive primes of a small range).

u	x	n	B_1, B_2	exposed rate (curve 1)	K_{cov}	$K_{\mathrm{sep}}^{\mathrm{label}}$	K_{sep} binary
3	2^{14}	1612	$25,645$	0.806	5	6	79
3	2^{16}	5709	$40,1625$	0.769	7	7	51
3	2^{18}	20390	$64,4096$	0.746	10	10	47
3	2^{20}	73586	$102,10321$	0.715	12	11	46
3	2^{22}	268216	$161,26008$	0.691	13	12	> 40
4	2^{14}	1612	$11,128$	0.407	20	19	27
4	2^{16}	5709	$16,256$	0.324	33	31	34
4	2^{18}	20390	$23,512$	0.311	35	32	34
4	2^{20}	73586	$32,1024$	0.291	41	36	46
4	2^{22}	268216	$45,2048$	0.253	53	53	55

TABLE 12. Covering length (every prime exposed by some curve), label-separating length (all label vectors distinct) and the separating length of the binary projection $[\lambda \neq 0]$, on every prime of $[x, 2x)$. At $u = 3$, $x = 2^{22}$ the binary projection had not separated after the 40 curves computed (18 pairs remained unresolved).

The exact labels separate in 6–12 curves at $u = 3$ – $6, 7, 10, 11, 12$ at $x = 2^{14}, 2^{16}, 2^{18}, 2^{20}, 2^{22}$, i.e. increments of $1, 3, 1, 1$ across successive quadruplings of x – consistent with slow, at most logarithmic, growth. The binary projection of the same predicate needs 46–79 on the four smaller intervals, because its success rate 0.69–0.81 is far from the optimal $1/2$, as the design rule predicts. A coarser bit does better: call the *one-large-prime bit* of a curve at p the event that the construction denominator vanishes, the stage-1 point is the identity, or the residual order is a prime in $(B_1, B_2]$ – the classical stage-2 success condition, which omits the singular states and the composite residual orders; it separates the four smaller populations in $24, 32, 39, 37$ curves. Hence *for* $(B_1, B_2) = (64, 4096)$ *the first ten Suyama curves are a finite certificate for every product of distinct primes in* $[2^{18}, 2^{19})$*, for* $(102, 10321)$ *the first eleven certify* $[2^{20}, 2^{21})$*, and for* $(161, 26008)$ *the first twelve certify* $[2^{22}, 2^{23})$ – computer-verified separating lists whose certificate is the parameter pair and the list of curves, verified by recomputing the labels, which the accompanying code does in minutes; the archived record is the class counts after each curve. At 2^{22} the uncovered fraction after k curves is $0.309, 0.097, 0.030, 0.010, 0.0033, \ldots$, close to 0.31^k, and the last pair is separated by the twelfth curve. (With parameters derived from N the certificates do not apply literally.) At $u = 4$ the exact labels separate every pair on all five intervals; the lengths $31, 32, 36, 53$ at $x = 2^{16}, \ldots, 2^{22}$ have increments $+1, +4, +17$, the last far larger than the earlier ones.

The separation tail is structured. For the six one-large-prime runs on the three smaller intervals (stage 1 alone and with the stage-2 continuation), the pairs still unresolved at the last recorded snapshot before separation number 311 in all, and $p \equiv q \pmod 3$ holds for 195 of them (0.63), compared with 0.50 for random pairs of the same populations. This is a finite verified enrichment; its explanation through the residue dependence of curve orders on $p \bmod 3$ is heuristic. *Plain* Montgomery curves – $A = 7, 8, \ldots$ with starting x-coordinate 3, without Suyama's forced 3-torsion – have one-large-prime success rate near $1/2$ at $u = 3$ and separate the three smaller populations in $22, 26, 30$ curves by that bit, compared with Suyama's $24, 32, 39$ and the birthday thresholds $2 \log_2 n = 21.3, 24.9, 28.6$.

34. The algorithm

The implementation tests N for a perfect square, derives $x = \sqrt{N/C}$, $B_1 = \mathrm{round}(x^{1/u})$, $B_2 = \mathrm{round}(x^{2/u})$, and runs $\sigma = 6, 7, \ldots$, capped at 400 curves. For each curve, the implementation takes the construction gcd and the singularity gcd $\gcd(a_{24}(a_{24}-1), N)$, backtracking to $\gcd(a_{24}, N)$ and $\gcd(a_{24}-1, N)$ when that gcd equals N; runs stage 1 and takes the gcd of its projective Z-coordinate; and runs a stage 2 over $\mathbb{Z}/N$ with $m = \lfloor\sqrt{B_2}\rfloor + 1$, $T = \lfloor B_2/m \rfloor$, $R = B_2 - Tm$: baby points jQ, $1 \le j < m$, giant points $i(mQ)$, $1 \le i \le T$, preparatory identity gcds (a proper gcd is returned at once; a point that is the identity modulo both primes, whose gcd is N and carries no factor, is removed from the baby and giant lists and the search continued), batch inversion, $F(t) = \prod_j (t - x(jQ))$ by a product tree evaluated at the giants $i < T$, the partial product over $j \le R$ at giant T, the gcd of the product, and, when backtracking is required, gcds of the $O(\sqrt{B_2})$ row products in increasing giant order followed by at most $O(\sqrt{B_2})$ individual atoms of the first non-trivial row in the nondecreasing key order of Proposition 32.4 (babies increasing within the row) – $O(\sqrt{B_2})$ gcds, not a scan of B_2 atoms – stopping at the first vanishing two-sided atom. The tested values are exactly $[1, B_2]$, so the curve factors N iff the labels of Proposition 32.4 differ. Per curve the ladder work is $O(B_1)$ group operations. The evaluation of F at the giants is done in fixed blocks of 4096 points, each block evaluated independently, to bound memory; this is not the single quasi-linear multipoint evaluation of the complexity statement after Conjecture 33.1, and no asymptotic claim is made for the blocked form.

Measurements. By *counted multiplications* we mean the number of modular multiplications performed in the Montgomery ladders and differential additions, the products of projective Z-coordinates, the batch inversions and the row products; the coefficient arithmetic inside the product tree and the multipoint evaluation is not counted. With the explicit parameters $(64, 4096)$, on 500 random prime pairs from $[2^{18}, 2^{19})$ (unordered, distinct, from a seeded generator) plus the 116 distinct pairs of the one-large-prime separation tail of that interval, all 616 were factored (maximum 3 curves, mean 1.08, counts 565/50/1 at one/two/three curves; 145 by stage 1, 471 by stage 2; the algorithm succeeded no later than the one-large-prime separation index – the first curve whose one-large-prime bits differ at p and q – on every pair, and strictly earlier on 54.7%). On 20 balanced semiprimes per size ($q/p < \sqrt{2}$, from a seeded generator) at $32, 40, 48, 56, 64$ bits with $u = 3$, $C = 2$ and N-derived parameters, the mean curve count is 1.0–1.2, the mean counted multiplications per modulus are $1104, 2715, 8920, 22\,725, 48\,358$, and the fitted exponent of the counted multiplications against the bit size is 0.175 against the prediction $1/(2u) = 0.167$ up to logarithmic factors, with a curve-count exponent of 0.003. Scaling, on single instances with wall-clock times and peak resident memory: blocking the multipoint evaluation cut a 128-bit $u = 3$ run from 179 s and 6.46 GiB to 96 s and 3.01 GiB; over the 120-, 124-, 128- and 136-bit instances the peak memory fits the empirical line $1.46\,\mathrm{KiB} \cdot B_1 + 13$ MiB; fixed $u = 3$ reached 136 bits in 203 s with 7.78 GiB, while larger u trades memory for curves ($u = 4$: 144 bits in 129 s and 405 MiB with 3 curves, 160 bits in 81 s and 1.44 GiB with 1 curve; $u = 5$: 180 bits in 671 s and 422 MiB with 10 curves). A greedy multi-u schedule – choosing at each step, among the curves $\sigma = 6, \ldots, 17$ at each of $u \in \{2.5, 3, 3.5, 4, 4.5\}$, the one resolving the most pairs per unit of B_1 – was slightly worse than the best separating fixed-u prefix from the same candidates on the two populations tested (152 against 150; 324 against 264 in the work proxy $\sum B_1$): cheap high-u coordinates leave an expensive cleanup tail.

35. What the route can and cannot give

Write $L_N[s, c] := \exp\big((c + o(1))(\log N)^s (\log\log N)^{1-s}\big)$. Each curve's exposure probability is modelled heuristically by the semismoothness [3] of a group or point order of size $N^{1/2+o(1)}$, so,

under the random-signature and semismoothness heuristics extrapolated uniformly as u grows, the optimised fixed-list model is $L_N[1/2,1]$ – the classical ECM heuristic, here as the cost of a fixed list – against the number field sieve's $L_N[1/3,(64/9)^{1/3}]$ [64]: in leading L-formulas, $2^{65.9}$ against $2^{63.9}$ operations at 512 bits, $2^{98.5}$ against $2^{86.8}$ at 1024, $2^{146.4}$ against $2^{116.9}$ at 2048 – formal values of the leading formulas, not operation counts of implementations. Exact labels, torsion families, better chains and batching change constants; the route's content is determinism and the finite certificates, not speed.

Proposition 35.1 (non-uniform derandomisation)**.** *Let $U_n = L_{2^n}[1/2,1]$ uniformly bound the expected running time – measured in bit operations, each random bit read costing one operation – of the Lenstra–Pomerance algorithm [54] on n-bit inputs. For every n there is an advice string w_n of length $O(nU_n)$ such that a deterministic algorithm with access to w_n factors every n-bit integer in time $O(nU_n)$.*

Proof. Run the algorithm with a fixed random string truncated after $4U_n$ steps; by Markov it succeeds with probability $\geq 3/4$ on every input. With $n+1$ independent strings the failure probability on a fixed input is $\leq 4^{-(n+1)}$, and the union bound over the $< 2^n$ inputs shows that some choice of strings succeeds on all of them. Each truncated run consumes $O(U_n)$ random bits, so the concatenated advice has length $O(nU_n)$. □

This is Adleman's argument [1] with subexponential advice. A uniform deterministic subexponential algorithm would follow from an explicit construction of such advice; a smoothness separating family is one candidate shape for it, and Conjecture 33.1 is the fixed-u version.

Proposition 35.2 (Miller's descent)**.** *Let p,q be distinct odd primes, $N=pq$, $\gcd(\alpha,N)=1$, $\operatorname{ord}_p(\alpha)=2^{j_p}u_p$ and $\operatorname{ord}_q(\alpha)=2^{j_q}u_q$ with u_p,u_q odd, and let $E=2^s m$ (m odd) be a common multiple of the two orders, as delivered by a two-sided $p-1$ collision $\alpha^E\equiv 1 \pmod N$. Testing $\gcd(x_j-1,N)$ along the binary chain $x_j=\alpha^{2^j m}$, $0\leq j\leq s$, yields a proper factor iff $j_p\neq j_q$.*

Proof. Since both orders divide E, $u_p,u_q \mid m$ and $j_p,j_q\leq s$, so $x_j\equiv 1 \pmod p$ first at $j=j_p$ and $\pmod q$ first at $j=j_q$. If $j_p<j_q$ then $\gcd(x_{j_p}-1,N)=p$, and symmetrically; if $j_p=j_q>0$ then x_{j_p-1} has order 2 modulo both odd primes, hence is -1 modulo both, and no gcd in the chain is proper; if $j_p=j_q=0$ there is no 2-adic descent. □

This is the criterion for Miller's binary chain [59] only: descent by an odd prime dividing one order and not the other can recover further cases when E can be factored.

A positive-density success theorem. The asymptotic running-time conclusions above for the explicit fixed-list smooth-order route are heuristic or conditional. We give one unconditional positive-density statement about the smooth-order route. It is an elementary corollary of two classical theorems on shifted primes, and we record it as the rigorous reference point of the route: it concerns the success set of a deterministic *partial* factoring algorithm – one allowed to fail on a complementary positive proportion – and is not an average running-time bound for an algorithm that factors every input. Let $P^+(n)$ denote the largest prime factor of n. Friedlander [26] proved that for every $\theta > 1/(2\sqrt{e}) = 0.30326\ldots$ there is $c_1(\theta)>0$ with $\#\{p\leq x : P^+(p-1)\leq x^\theta\}\geq c_1 x/\log x$ for all large x – a positive proportion of the primes; the record exponents 0.2961 [4] and 0.2844 [51] are stated as infinitude results (infinitely many primes p with $P^+(p-1)\leq p^\theta$), not as positive proportions. Fouvry [25] proved that there is a fixed exponent $\alpha>2/3$ – one may take $\alpha=0.6687$ in the published numerical form – and $c_2>0$ with $\#\{q\leq x : P^+(q-1)\geq q^\alpha\}\geq c_2 x/\log x$. The algorithm $\mathcal{A}_\theta$ is Pollard's $p-1$ method [68] with base 2: on input N put $B:=\lceil N^{\theta/2}\rceil$, $E:=\prod_{\ell\leq B,\ \ell \text{ prime}} \ell^{\lfloor \log N/\log \ell\rfloor}$ and $g:=\gcd(2^E-1 \bmod N,\ N)$, and output g if $1<g<N$. It is deterministic. Since $\log_2 E\leq \pi(B)\log_2 N$, binary exponentiation (prime power

by prime power, without materialising E) costs at most $2\pi(B)\log_2 N + O(1)$ multiplications modulo N; with $\pi(B) \ll B/\log B$ and $\log B = (\theta/2)\log N + O(1)$ this is $O_\theta(B) = O_\theta(N^{\theta/2})$. Prime generation and the final gcd add only factors polynomial in $\log N$ to the bit complexity.

Proposition 35.3 (positive-density success of the $p-1$ method)**.** *Fix $\theta \in (1/(2\sqrt{e}), 1/2)$ and $\varepsilon \in (0, 1/4)$ with $\theta(1-\varepsilon) > 1/(2\sqrt{e})$. There are $c = c(\theta,\varepsilon) > 0$ and x_0 such that for every $x \ge x_0$, at least $c\,\pi(x)^2$ of the ordered pairs (p,q) of distinct primes in $(x^{1-\varepsilon}, x]$ satisfy $\mathcal{A}_\theta(pq) = p$. Hence a positive relative proportion of the semiprimes $N = pq$ with $x^{1-\varepsilon} < p, q \le x$ are factored deterministically in $O_\theta(N^{\theta/2})$ multiplications modulo N. For fixed ε the admissible exponent $\theta/2$ can approach $1/(4\sqrt{e}(1-\varepsilon))$; letting $\varepsilon \downarrow 0$ and $\theta \downarrow 1/(2\sqrt{e})$ jointly, subject to $\theta(1-\varepsilon) > 1/(2\sqrt{e})$, brings it arbitrarily close to $1/(4\sqrt{e}) = 0.151632\ldots$.*

Proof. Put $\theta' := \theta(1-\varepsilon)$ and $S := \{p \in (x^{1-\varepsilon}, x] : P^+(p-1) \le x^{\theta'}\}$; by Friedlander's theorem at θ', $|S| \ge c_1(\theta')x/\log x - \pi(x^{1-\varepsilon}) \ge \frac{1}{2}c_1 x/\log x$ for large x. Let $\Omega := \{q : \mathrm{ord}_q(2) \le x^{1-\alpha}\}$; a prime with $\mathrm{ord}_q(2) = d$ divides $2^d - 1$, which has fewer than d prime factors, so $|\Omega| \le \sum_{d \le x^{1-\alpha}} d \le x^{2-2\alpha} = x^{0.6626}$. Let $G := \{q \in (x^{1-\varepsilon}, x] : P^+(q-1) \ge q^\alpha\} \setminus \Omega$; by Fouvry's theorem $|G| \ge \frac{1}{2}c_2 x/\log x$ for large x. Take $p \in S$, $q \in G$, $p \ne q$, and $N := pq \in (x^{2-2\varepsilon}, x^2]$. (i) $B \ge N^{\theta/2} > x^{\theta'} \ge P^+(p-1)$, and every prime power $\ell^v \,\|\, p-1$ has $\ell^v < N$, so $v \le \lfloor \log N/\log \ell \rfloor$; hence $p-1 \mid E$ and $p \mid 2^E - 1$. (ii) If $q \mid 2^E - 1$ then $\mathrm{ord}_q(2) \mid E$ is B-smooth. Since $B \le x^\theta + 1 < x^{\alpha(1-\varepsilon)} < q^\alpha \le P^+(q-1)$ for large x (as $\alpha(1-\varepsilon) > 3\alpha/4 > 1/2 > \theta$), the prime $P^+(q-1)$ does not divide $\mathrm{ord}_q(2)$, and $\mathrm{ord}_q(2) \mid q-1$ gives $\mathrm{ord}_q(2) \le (q-1)/P^+(q-1) < q^{1-\alpha} \le x^{1-\alpha}$, i.e. $q \in \Omega$, a contradiction. So $g = \gcd(2^E - 1, N) = p$. (iii) The pairs number at least $|S||G| - |S \cap G| \ge \frac{1}{4}c_1 c_2 (x/\log x)^2 - x \ge c\,\pi(x)^2$ for large x, with $c := c_1 c_2/7$, by $\pi(x) \le 1.26\,x/\log x$ [74]. Each semiprime pq arises from at most two ordered pairs, so the semiprimes factored number at least $c\,\pi(x)^2/2$, a positive proportion of the fewer than $\pi(x)^2/2$ semiprimes of the family. □

Remarks. (a) On $S \times G$ neither the two-sided outcome $g = N$ nor the failure $g = 1$ occurs; on the whole population both do. (b) For independent primes near x the smoothness heuristic predicts that $p-1$ is x^θ-smooth with probability $\varrho = \rho(1/\theta)$ (Dickman's ρ); if in addition the exceptional event $\mathrm{ord}_p(2) \mid E$ with $p-1$ not B-smooth has density $o(1)$, the idealised one-sided success rate is $2\varrho(1-\varrho)$. At $\theta = 0.31$, $\rho(3.2258\ldots) = 0.02997$ and $2\varrho(1-\varrho) = 0.0582$ – a Dickman proxy, not an established asymptotic; the lower-bound set S of the proof has the smaller parameter $\theta' = \theta(1-\varepsilon)$ and Dickman density $\rho(1/\theta')$. We ran $\mathcal{A}_{0.31}$ on 2000 balanced semiprimes ($q/p < \sqrt{2}$, from a seeded generator) at each of $40, 48, 56, 64$ bits, with $B = \lceil N^{0.155} \rceil$ between 67 and 967: it factored $204, 178, 165, 152$ of them ($0.102, 0.089, 0.083, 0.076$, binomial standard errors below 0.007), returned $g = N$ for $8, 8, 6, 7$, and failed on the rest – decreasing over these four sizes and above the proxy, consistent with a finite-size excess of smooth shifted primes; the data do not establish convergence. The order anomaly of step (ii) – a factor whose shifted prime is not B-smooth but whose order of 2 divides E – occurred for $2, 2, 0, 0$ of the 2000 moduli. (c) With the stage 2 of Section 34 ($B_2 = B^2$, cost still $B^{1+o(1)}$), a (B, B^2)-semismooth $p-1$ is a sufficient exposure condition on the p-side, not a necessary one ($\mathrm{ord}_p(2)$ may be a proper divisor of $p-1$); a one-sided theorem again needs the other side excluded, which the Fouvry set provides by the same order argument when $B^2 < P^+(q-1)$, i.e. for $2\theta < \alpha(1-\varepsilon)$. In the same samples exactly one of $p-1, q-1$ is (B, B^2)-semismooth for $0.496, 0.482, 0.504, 0.483$ of the moduli (a measurement, not an asymptotic statement); a positive-proportion theorem for $p-1 = (x^\theta\text{-smooth}) \times (\text{prime} \le x^{2\theta})$ at some $\theta < 1/(2\sqrt{e})$ would lower the exponent of Proposition 35.3 accordingly, and we have not found one: the shifted-prime results available [26, 25, 4, 51] concern smoothness or a large prime factor alone. (d) Harvey's $N^{1/5}$ [31] is a worst-case bound valid for every N; given any slack $\delta > 0$, choosing first the balance parameter $\varepsilon > 0$

small and then θ with $1/(2\sqrt{e}(1-\varepsilon)) < \theta < 1/(2\sqrt{e})+2\delta$, Proposition 35.3 says that on a positive relative proportion of the corresponding nearly balanced semiprimes a plain algorithm from 1974 provably succeeds within $O(N^{1/(4\sqrt{e})+\delta})$ modular multiplications. It is not an improvement of the worst-case guarantee, and a fallback to Harvey's algorithm on the failures leaves the expected cost at $N^{1/5+o(1)}$. The record smoothness thresholds would formally lead to exponents $0.2961/2 = 0.14805$ and $0.2844/2 = 0.1422$, each with a positive slack; the theorems behind them do not give a positive proportion. If the available smooth-prime count is $\gg x/(\log x)^C$, the same pairing gives $\gg x^2/(\log x)^{C+1}$ products, a relative proportion $\gg (\log x)^{1-C}$ of the family; an infinitude theorem gives infinitely many such products. (e) The fixed-θ run does not factor almost all semiprimes of the family: after removing the $o(\pi(x))$ small-order exceptions Ω, the Fouvry set still contains a positive proportion of primes q with $q \nmid 2^E - 1$, and pairs of two such primes form a positive proportion on which $g = 1$. More generally, fix $\eta \in (0, \varepsilon)$ with $0.66 < \alpha(1-\eta)$ and let $G_\eta := \{q \in (x^{1-\eta}, x] : P^+(q-1) \geq q^\alpha\} \setminus \Omega$, still a positive proportion of the primes of the family. For $q, r \in G_\eta$ and $N = qr \leq x^2$, every $B, B_2 \leq N^{0.33} \leq x^{0.66} < x^{\alpha(1-\eta)} < q^\alpha \leq P^+(q-1)$; if the stage-1 exponent has prime support at most B then $\mathrm{ord}_q(2) \nmid Ev$ for every $1 \leq v \leq B_2$ (otherwise $P^+(q-1) \nmid \mathrm{ord}_q(2)$, forcing $\mathrm{ord}_q(2) < q^{1-\alpha} \leq x^{1-\alpha}$, contrary to $q \notin \Omega$), and likewise for r. So pairs from G_η form a positive relative proportion on which the single-exponent base-2 $p-1$ method, with any stage-1 and stage-2 bounds up to $N^{0.33}$, returns no factor: a limitation of that method, not a lower bound for all algorithms of that cost.

36. Open problems

Prove or break Conjecture 33.1 for an explicit curve family, including uniform control as u grows; find a separating family with per-member success rate near $1/2$ and an alphabet richer than the residual order; determine whether the separation tail structured by $p \bmod 3$ can be removed by mixing families; and prove a positive-proportion theorem for shifted primes $p - 1 = (x^\theta\text{-smooth}) \times (\text{prime} \leq x^{2\theta})$ below $\theta = 1/(2\sqrt{e})$, which would lower the exponent of Proposition 35.3. Any of the first three could change constants or coverage; heuristically, an exponent change requires leaving the smooth-order paradigm.

Reproducibility. The label statistics of Table 12 are exact computations on the full prime populations of the stated intervals, by adaptive refinement of the label partition as described in Section 33.3; the tests of Section 34 and the $p-1$ samples of remark (b) use pseudorandom populations from a seeded generator (seeds 3, 9 and 11 respectively) and are reproducible from those seeds; the counted multiplications and curve counts are deterministic; the wall-clock times and peak memory figures are single measurements on one machine and are reported as such. The code and the archived outputs are in the repository [23].

Part 6. Coefficient floors of number-field-sieve polynomial selection

Summary. Polynomial selection for the number field sieve chooses a polynomial f of degree d and a root m of f modulo N; in the unskewed model the sieve's figure of merit is the norm product $P = \|f\|_\infty m$, and this part asks how small it can be made. For a fixed root the coefficient vectors form a lattice of determinant N, so a polynomial with $\|f\|_\infty \leq N^{1/(d+1)}$ exists, and for a random root of a balanced semiprime $N^{1/(d+1)-o(1)}$ is a proven floor. When the root is free, a counting heuristic puts the minimum of P at $N^{(2d-1)/(d^2+d-1)}$, below the scale $N^{2/(d+1)}$ of the base-m and Kleinjung constructions, and complete enumerations at small sizes fit that exponent at $d = 2$ and $d = 3$. For two polynomials with a common root the resultant gives the rigorous floor $\|f\|_\infty^e \|g\|_\infty^d \gg N$, and a count of the pairs with a common root modulo a generic N gives, within the standard heuristic optimisation, the existence bound $\alpha\delta \geq 2 - o(1)$ for the two-polynomial dial: the GNFS constant 1.923 is then its existence optimum, and

the SNFS constant is out of reach for generic N at every pair of degrees. The room below the base-m scale is lower order in L-notation, and reaching it by search is estimated to cost a fractional power of N in trials: the floor asks for an algebraic small-coefficient method, not search. The rigorous content is the two floors; the rest is heuristic, with its finite verification stated as such.

37. Introduction

The general number field sieve [53, 9] factors N in heuristic time $L_N[1/3, (64/9)^{1/3}]$ with $L_N[s,c] = \exp((c+o(1))(\ln N)^s(\ln\ln N)^{1-s})$. Its polynomial selection step chooses $f \in \mathbb{Z}[x]$ of degree d with a root m modulo N; the sizes of $\|f\|_\infty$ and m enter the exponent through the norm product, and in practice through Murphy's E-value, skew, rotation and root properties [62, 45, 6]. This part isolates the size question – how small can $\|f\|_\infty$ and m be made jointly, and what is the consequence for the constant c – and answers it with two proven floors, one heuristic floor checked by exact enumeration, one classical identity read as a search problem, and one heuristic optimisation. Throughout, $\|f\|_\infty := \max_i |f_i|$ is the largest absolute value of a coefficient of $f = \sum_i f_i x^i$, and N is called *generic* when the counting heuristic of Proposition 39.1 – roots of bounded polynomials equidistributed modulo N – is assumed for it. The asymptotic conclusion, within the two-polynomial heuristic model developed here, is negative and quantified: coefficient size does not change the constant $(64/9)^{1/3}$ for generic N, and any remaining gain from polynomial selection is lower order. In numbers: the random-root floor is $N^{1/(d+1)-o(1)}$ (Proposition 38.1); the joint heuristic floor $N^{(2d-1)/(d^2+d-1)}$ is matched by exact enumerations with fitted exponents 0.607 ± 0.009 at $d=2$ (predicted 0.6) and 0.460 ± 0.014 at $d=3$ (predicted 5/11); the resultant floor $\|f\|_\infty^e\|g\|_\infty^d \gg N$ together with the heuristic count of pairs with a common root gives $\alpha\delta \geq 2 - O(1/(d+e))$ for generic N, hence the GNFS constant 1.923 as the existence optimum and not the SNFS constant 1.526; for degree two, selection is a bounded near-square search whose event counts are consistent with a Poisson model; and the gap between the joint floor and the base-m scale is 9–24 bits in the norm product at 512–1024 bits for $d = 5, 6$, which direct search would need about $N^{0.09}$–$N^{0.10}$ trials to cross and $N^{0.12}$–$N^{0.14}$ to reach the floor.

38. The fixed-root lattice

For a residue m let $L_{N,m} = \{(f_0,\dots,f_d) \in \mathbb{Z}^{d+1} : \sum_i f_i m^i \equiv 0 \pmod N\}$, the kernel of $(f_i) \mapsto \sum f_i m^i \bmod N$, a lattice of determinant N.

Proposition 38.1. *(i) $L_{N,m}$ contains a nonzero vector of sup-norm $\leq N^{1/(d+1)}$. (ii) Let $N = pq$ with distinct primes p, q, and let $H < \min(p,q)$. For m uniform on $\mathbb{Z}/N$,*

$$\Pr[\exists f \neq 0,\ \|f\|_\infty \leq H,\ f(m) \equiv 0 \pmod N] \leq d^2(2H+1)^{d+1}/N,$$

so, when $\min(p,q) \geq N^{1/(d+1)}$ – in particular for balanced semiprimes – for every fixed $\varepsilon > 0$ the shortest vector at random m has sup-norm $\geq N^{1/(d+1)-\varepsilon}$ with probability at least $1 - d^2 3^{d+1} N^{-\varepsilon(d+1)}$.

Proof. (i) is Minkowski on the cube. (ii): a nonzero f with $\|f\|_\infty \leq H$ is nonzero modulo p and modulo q, hence has at most d roots modulo each and at most d^2 modulo N by the Chinese remainder theorem ($x^2 - 1$ attains four); there are fewer than $(2H+1)^{d+1}$ such f. □

The Minkowski vector represents a nonzero polynomial of degree at most d; exact degree, irreducibility and usability for the sieve are not asserted. Exact Euclidean shortest-vector lengths (LLL-seeded Fincke–Pohst enumeration in dimension $d+1 \leq 7$), normalised by $N^{1/(d+1)}$ and averaged over 20 balanced semiprimes ($q/p < \sqrt{2}$) per size at 64, 96 and 128 bits, are 0.62–0.80 for random m, matching the Gaussian-heuristic constants $V_n^{-1/n}$, V_n the volume of the

Euclidean unit ball in $\mathbb{R}^n$ (0.62–0.80 for $n = 3, \dots, 7$). For the base-m choice $m = \lfloor N^{1/d} \rfloor$, simultaneous Dirichlet approximation of $r/m, \dots, r/m^{d-1}$ with $r = N - m^d$ gives vectors of sup-norm $O_d(N^{(d-1)/d^2})$, a ratio $N^{-1/(d^2(d+1))}$ to the generic fixed-root scale $N^{1/(d+1)}$ that predicts the measured change of the ratio from 64 to 128 bits ($0.291, 0.574, 0.839$ predicted against $0.30, 0.56, 0.84$ measured for $d = 3, 4, 6$); for special $N = 2^k - c$ the lattice contains $2^s x^d - c$ and the ratio collapses.

39. The joint floor

When the root is free, the figure of merit is $P(f, m) = \|f\|_\infty m$, where for a nonzero residue root r we put $m := \min(r, N - r) \in [1, N/2]$ and, replacing $f(x)$ by $f(-x)$ when $r > N/2$ (which preserves $\|f\|_\infty$), assume $f(m) \equiv 0 \pmod N$ (the root 0 is excluded, since $x^d + N$ has $P = 0$). Throughout, the admissible polynomials are those of exact degree d, irreducible over $\mathbb{Q}$ (the content is not restricted; every minimiser found below is primitive), with a nonzero root modulo N; $f(m)$ may be any nonzero multiple of N, and $\|f\|_\infty$ is the largest absolute value among all $d + 1$ coefficients.

Proposition 39.1 (heuristic)**.** *Let $N = pq$ be generic and $d \geq 2$. Among irreducible f of degree d with a nonzero root modulo N,*

$$\min P(f, m) = N^{(2d-1)/(d^2+d-1)+o(1)}, \qquad \|f\|_\infty = N^{(d-1)/(d^2+d-1)+o(1)}, \quad m = N^{d/(d^2+d-1)+o(1)}$$

at the minimiser.

Derivation. There are two constraints. (i) Counting: the number of polynomials with $\|f\|_\infty \leq H$ is $\asymp H^{d+1}$; a polynomial has on average one root modulo N, and the roots are assumed equidistributed, so the number of pairs with $m \leq M$ is $\asymp H^{d+1}M/N$; existence requires $(d + 1) \log_N H + \log_N M \geq 1$. (ii) Nontriviality: an irreducible f has $f(m) \neq 0$, so $|f(m)| \geq N$ and $(d + 1)\|f\|_\infty m^d \geq N$. Minimising $\log_N H + \log_N M$ on this polyhedron gives the vertex $((d - 1)/(d^2 + d - 1), d/(d^2 + d - 1))$. Constraint (ii) is a theorem and gives the rigorous $P \geq (N/(d + 1))^{1/d}$; constraint (i) is the heuristic. □

The three construction regimes in this measure are: Dirichlet vectors at $m = N^{1/d}$, $P = N^{(2d-1)/d^2}$; base-m and Kleinjung's regime $\|f\|_\infty \approx m \approx N^{1/(d+1)}$, $P = N^{2/(d+1)}$; and the joint floor $N^{(2d-1)/(d^2+d-1)}$ – $N^{0.75}, N^{0.667}, N^{0.6}$ at $d = 2$ and $N^{0.36}, N^{0.333}, N^{0.310}$ at $d = 5$ (the balanced base-m expansion attains the scale $N^{2/(d+1)}$ without the irreducibility requirement, and the constructions of [45] attain it in practice). The gap between the last two is $N^{(d-1)/((d+1)(d^2+d-1))}$, lower order in L-notation: about 9–24 bits of norm product at 512–1024 bits for $d = 5, 6$.

Refined count. After imposing the leading-term cutoff $m \gtrsim (N/f_d)^{1/d}$ and quotienting by mirror orbits, the expected number $\Lambda_d(x)$ of pairs with $P \leq x$ is a finite sum over coefficient shells whose leading term is $\frac{2}{3}x^5/N^3$ ($d = 2$) and $\frac{16}{15}x^{11/2}/N^{5/2}$ ($d = 3$); under a Poisson model with these intensities $\mathbb{E}P_{\min} = 0.996\, N^{3/5}$ and $0.912\, N^{5/11}$.

Complete enumeration. Using the known factorisation, the exact minimum of P over the admissible polynomials was computed by complete enumeration (square-root tables modulo p, q for $d = 2$, value tables for $d = 3$); the pruning bounds are exact consequences of $|f(\pm m)| \geq N$, the enumeration checks at run time that its finite search windows were not exceeded, and the routine is checked against brute force on small moduli in the accompanying test suite. For 51 balanced semiprimes at 24–44 bits ($d = 2$) the ordinary least-squares slope of $\log_2 P_{\min}$ on $\log_2 N$ is 0.6067 ± 0.0088 (predicted 0.6); for 42 distinct moduli among 44 generated instances at 16–32 bits ($d = 3$) it is 0.4603 ± 0.0137 (predicted $5/11 = 0.4545$; 0.4595 ± 0.0131 with the two duplicated moduli counted twice). The base-m control – one balanced expansion per N

at $m = \lceil N^{1/(d+1)} \rceil$ with $f(m) = N$, not restricted to irreducible f – fits 0.6669 against 2/3 and 0.5037 against 1/2. The row means of the coefficient and root exponents are 0.18–0.20 and 0.40–0.41 ($d = 2$) and 0.13–0.17 and 0.28–0.29 ($d = 3$); individual instances spread more widely. The ratio $P_{\min}/\mathbb{E}P_{\min}$ ranges over 0.70–1.09 rowwise; under the Poisson model $\Lambda_d(P_{\min})$ should be exponential with mean one, and it has mean 0.67 and 0.49 (Kolmogorov–Smirnov $p = 0.048$ and 0.001 on the generated instances): the model is borderline for $d = 2$ and rejected for $d = 3$, the minima lying below its prediction. A search through $8N^{(d-1)/(d^2+d-1)}$ leading coefficients (two adjacent candidate bases per leading coefficient, balanced lower digits, irreducible results only) found 50/51 and 37/44 exact minima; in the cubic case the mean excess was 0.070 bits – finite evidence for the predicted search scale, not an asymptotic search theorem.

40. Two polynomials: the resultant floor

Proposition 40.1. *Let $f, g \in \mathbb{Z}[x]$ be coprime over $\mathbb{Q}$, of degrees d, e, with a common root modulo each prime dividing squarefree N. Then $N \mid \operatorname{Res}(f, g)$ and*

$$N \le |\operatorname{Res}(f,g)| \le \|f\|_2^e \, \|g\|_2^d \le (d+1)^{e/2}(e+1)^{d/2}\|f\|_\infty^e \|g\|_\infty^d.$$

In particular two quadratics with a common root modulo N satisfy $\|f\|_\infty \|g\|_\infty \ge \sqrt{N}/3$.

Proof. $\operatorname{Res}(f, g) = Af + Bg$ with $A, B \in \mathbb{Z}[x]$ (adjugate of the Sylvester matrix), so a common root modulo p gives $p \mid \operatorname{Res}$; coprimality makes the resultant nonzero; Hadamard's inequality on the Sylvester matrix gives the norm bounds. □

Montgomery's two-quadratics construction [61], as described in [62], produces pairs at this floor up to a constant factor for the N to which it applies. For the usual linear side $g = x - m$ the inequality is $\|f\|_\infty m^d \gg_d N$, the nontriviality constraint of Proposition 39.1. The inequality is standard; its reading as a coefficient floor for nonlinear pairs is the point here. Exact (d, d) floors – complete enumeration over primitive, coprime, irreducible pairs of exact degree d with a nonzero common root, on 18 moduli at 16–28 bits ($d = 2$) and 21 moduli at 16–24 bits ($d = 3$) – give $P_2 = \|f\|_\infty \|g\|_\infty$ at 1.4–1.9 times the resultant lower bound $N^{1/d}/(d + 1)$, with fitted slopes 0.482 ± 0.005 ($d = 2$) and 0.313 ± 0.015 ($d = 3$) approaching $1/d$ from below; 16 of the 21 cubic minimisers have a monic side of sup-norm 1, whose partner is a short vector of the product of the two prime ideals of $\mathbb{Z}[\alpha]$ determined by the chosen modular roots, α a root of the monic side, and that "tiny-f" route has fitted exponents $0.5004, 0.3332, 0.2483$ for $d = 2, 3, 4$ on 33, 34 and 41 moduli at 24–96 bits, agreeing with $1/d$ to three decimal places. Under the balanced-minima heuristic for the orthogonal lattice, the constructible geometric-progression pairs [61, 47, 70, 18] (root-lift of a d-th root of N modulo a prime $p' \approx N^{1/d}$, orthogonal lattice of determinant $\approx N^{1-1/d}$) are predicted to reach only $P_2 \approx N^{2(d-1)/d^2}$ (measured slopes 0.4991 ± 0.0018, 0.4493 ± 0.0030, 0.3744 ± 0.0034 on ten moduli at each of $40, 64, 96, 128$ bits, against $0.5, 0.444, 0.375$), which equals the floor only at $d = 2$.

41. Degree two: a bounded near-square search

Lemma 41.1 (bridge)**.** *Let $f = ax^2 + bx + c \in \mathbb{Z}[x]$, $a > 0$, and $f(m) = kN$ with $k \ne 0$; put $y = 2am + b$. Then $y^2 - 4akN = b^2 - 4ac = \operatorname{disc} f$. Conversely, given (a, k, y) and $b \equiv y \pmod{2a}$, the integers $m = (y - b)/(2a)$ and $c = (b^2 - (y^2 - 4akN))/(4a)$ satisfy $f(m) = kN$. If $\|f\|_\infty \le H$ then $|\operatorname{disc} f| \le 5H^2$, sharp at $(H, H, -H)$.*

Proof. $(2am + b)^2 = 4a(am^2 + bm + c) + b^2 - 4ac = 4akN + \operatorname{disc} f$. Conversely $b \equiv y \pmod{2a}$ gives $b^2 \equiv y^2 \pmod{4a}$, so $c \in \mathbb{Z}$ and substitution gives $f(m) = kN$; the discriminant bound is $|b^2 - 4ac| \le H^2 + 4H^2$. □

The pairs (f, m) with $a > 0$ therefore correspond bijectively to quadruples (a, k, y, b) with $b \equiv y \pmod{2a}$; the choice of b within its class selects a translate $f(x + t)$, which preserves y. A bounded known-root quadratic is a near-square $4akN = y^2 - \Delta$ with $|\Delta| \le 5H^2$, and at the joint floor $H = N^{1/5}$ degree-two selection is the search for $a \le N^{1/5}$ and a bounded multiplier k with $4akN$ within $O(N^{2/5})$ of a square, plus a congruent representative passing the coefficient filters (the near-square condition is necessary, not sufficient). The identity is classical (Lehman [49], MPQS's $(ax + b)^2 - N = aQ(x)$ [77], Montgomery [61]); Harvey's deterministic algorithm [31] shares the multiplier exponent $N^{1/5}$ and the identity $U^2 - 4abN = v^2$ but requires a square residual, where selection needs only a small one. The expected number of bounded $k = 1$ pairs is $\lambda(H) \approx 4H^{5/2}/\sqrt{N}$ under equidistributed phases, the finite sum behind this leading term being what the computation predicts. For 200 balanced semiprimes at each of $40, 56, 72, 80$ bits, exhaustive enumeration through the bridge at $H = CN^{1/5}$, $C = 0.75, 1, 1.5$, gives pair means within 10% of the finite-sum prediction in 8 of the 12 (size, C) cells, within 12% in three more, and 25% below it at $C = 0.75$ and 56 bits; at $C = 1$ the observed means are 3.5–4.3 against the predicted 3.8–4.0. The pair variance is about ten times the mean, because translates and scalings $(at^2, yt, \Delta t^2)$ are exact dependencies; the scaling-primitive classes have variance approximately equal to the mean, and their means reproduce the zero fractions through $e^{-\mu}$ to within 0.024 – consistent with a Poisson model once the two dependencies are removed, a diagnostic rather than a goodness-of-fit test. The chirp phases $(y^2 - 4aN)/(2y)$ are consistent with a uniform uncorrelated null (variance 0.0830 against $1/12 = 0.0833$; pooled Kolmogorov–Smirnov $p = 0.21$ on 50 phases per modulus treated as independent, again a diagnostic; lag-one correlation 0.000 ± 0.004). Among the moduli with at least one $k = 1$ candidate at $C = 1.5$ (784 of the 800), the best $k = 1$ product has median 0.93–0.98 $N^{3/5}$ from 40 to 80 bits, an upper bound at the scale of the exact minima for those moduli.

42. The constant dial

Write the norm product of a two-polynomial relation scheme, up to lower-order factors, as $X = N^{\alpha/d} A^{\delta d}$: α is the N-content the pair carries per unit $1/d$ and δ the growth of the norm product per unit degree. Minimising $\ln X$ over d gives $\ln X = 2\sqrt{\alpha\delta \ln N \ln A}$; with A^2 candidate pairs (a, b), $|a|, |b| \le A$, Dickman smoothness with independent sides, a smoothness bound B and B^2 linear algebra, the usual balance gives

$$c(\alpha\delta) = \Big(\frac{32\,\alpha\delta}{9}\Big)^{1/3}$$

for the constant of $L_N[1/3, c]$ – 1.923 at $\alpha\delta = 2$ (GNFS: $\|f\|_\infty \approx m \approx N^{1/(d+1)}$, $\alpha = 2$, $\delta = 1$) and 1.526 at $\alpha\delta = 1$ (SNFS: $\|f\|_\infty = O(1)$, $m = N^{1/d}$). With A^t candidates the balance gives $c_t = (64\alpha\delta/(9t))^{1/3}$; if a t-dimensional sieve carries N-content $\alpha = 2(t - 1)$ – the value obtained by extrapolating the two-polynomial count, an assumption we do not derive – then $c_t = (128(t - 1)/(9t))^{1/3}$ is minimal at $t = 2$ within this model. The model is heuristic and calibrated on the two known constants.

Proposition 42.1 (the floor of the dial)**.** *Let f, g be coprime of degrees $d \ge e \ge 1$ with a common root modulo N, $F = \ln\|f\|_\infty$, $G = \ln\|g\|_\infty$, $L = \ln N$, so that $\alpha = d(F + G)/L$ and $\delta = (d+e)/d$. (i) Proposition 40.1 gives $eF + dG \ge L - O_{d,e}(1)$, hence $F + G \ge L/d - O(1)$ and*

$$\alpha\delta \;\ge\; 1 + \frac{e}{d} - O_{d,e}\Big(\frac{1}{L}\Big).$$

(ii) [heuristic] For generic N, the number of pairs with $\|f\|_\infty \le e^F$, $\|g\|_\infty \le e^G$ and a common root is about $e^{(d+1)F+(e+1)G}/N$, so existence needs $(d + 1)F + (e + 1)G \ge L$; minimising $F + G$

under (i) and (ii) gives $F+G = 2L/(d+e+1)$ for $e < d$ (at the vertex $F = L(d-e-1)/(d(d+1)-e(e+1))$, $G = L(d-e+1)/(d(d+1)-e(e+1))$) and L/d for $e = d$, hence

$$\alpha\delta \;\geq\; \frac{2(d+e)}{d+e+1} - o(1) \quad (e<d), \qquad \alpha\delta \geq 2 - o(1) \quad (e=d).$$

Consequences, within this heuristic model. For generic N the existence optimum of the two-polynomial dial is $\alpha\delta = 2 - o(1)$, i.e. $c = (64/9)^{1/3}$: GNFS's constant is optimal among two-polynomial relation schemes, and the room below it – the joint floor of Proposition 39.1 ($\alpha\delta = (2d-1)(d+1)/(d^2+d-1) \to 2$), non-monic linear sides ($2(d+1)/(d+2)$), skew – is lower order. The SNFS constant requires $\alpha\delta = 1 + o(1)$, which generic N does not admit at any pair of degrees; special N realise it through $f(m) = kN$ with $\|f\|_\infty = O(1)$. A modular root of a bounded polynomial buys only a (d,d) pair at the resultant floor through the ideal lattice, with $\alpha\delta = 2$ – a tie with GNFS (the (d,d) pair at its floor has the norm product $N^{1/d}A^{2d} = N^{2/D}A^D$ with $D = 2d$, that of a linear pair at base-m scale with degree parameter D), and the geometric-progression pairs have $\alpha\delta \to 4$, $c = 2.423$, worse. Constant-level gains for generic N must come from outside the two-polynomial dial – side multiplicity (Coppersmith's multiple-polynomial NFS [16], 1.902) and amortisation (Coppersmith's factory [16], 1.639 per target after a shared 2.007 precomputation) – or from outside the relation paradigm.

The free-root frontier (heuristic). Searching T leading coefficients $a_d = 1, \ldots, T$ with $m \approx (N/T)^{1/d}$ and balanced lower digits, the best of T trials has $P(T) \approx \max(T, (m/2)T^{-1/(d-1)})\, m$, which crosses the base-m scale $N^{2/(d+1)}$ at $T_\times = N^{2(d-1)/((d+1)(3d-2))}$ and reaches the joint floor at $T = N^{(d-1)/(d^2+d-1)}$. On three balanced semiprimes per size ($d = 3$ at 48–96 bits, $d = 4$ at 64–128 bits, budgets T checked at the powers of two up to $T_{\max} = 2^{18}$), six of the seven measured crossovers lie within one bit of the prediction and the seventh ($d = 4$, 128 bits) came 1.3 bits earlier; the theory curve is within 0.5 bits of the measured minimum at every crossover, and in the six rows with $T_{\max} \geq N^{(d-1)/(d^2+d-1)}$ the measured minimum ends 0.1–0.5 bits below the predicted floor, the formula continuing to fall past the floor while the measurement stops there. Heuristic extrapolation to NFS sizes puts the crossover budget at 2^{46}–2^{105} trials and the floor 2^{62}–2^{141} for $d = 5, 6$ at 512–1024 bits: realising the joint floor appears to require an algebraic method that makes $d-1$ base-m digits small at a free root, rather than direct search; within the dial, the resulting gain is lower order and does not change c.

43. What is open

The open questions are whether the joint floor can be reached in polynomial time (an algebraic small-digit selection); whether balanced (d,d) pairs at the resultant floor exist without a tiny-polynomial root (the counting heuristic predicts many); and, outside this part, whether any relation source has norms that do not carry $N^{\Theta(1/d)}$ per side while growing like $A^{\Theta(d)}$ – within the heuristic framework used here, the only place where the exponent $1/3$, rather than its constant, could move.

Reproducibility. Every number above is the output of a deterministic computation from the stated parameters: exact shortest-vector lengths, complete enumerations of the exact minima with pruning bounds derived from $|f(\pm m)| \geq N$ (checked against brute force on small moduli in the accompanying test suite), exhaustive enumeration through the bridge, and the coefficient search, on populations of balanced semiprimes ($q/p < \sqrt{2}$) drawn from a seeded generator (seeds 97, 111, 5, 131 and 23 for the fixed-root, single-polynomial, two-polynomial, bridge and frontier computations). Methods: all enumerations and root tests are in exact integer arithmetic; the reported slopes are ordinary least-squares fits of $\log_2 P_{\min}$ against $\log_2 N$ with the slope's standard error from the residuals; the Poisson checks compare the empirical distribution of

$\Lambda_d(P_{\min})$ with the exponential law of mean one by the Kolmogorov–Smirnov statistic and its asymptotic p-value; all statistics are in double precision. The code and the archived outputs accompany the paper.

References

[1] L. M. Adleman, Two theorems on random polynomial time, *Proc. 19th IEEE Symp. Foundations of Computer Science* (1978), 75–83.

[2] D. Aggarwal, U. Maurer, Breaking RSA generically is equivalent to factoring, *Advances in Cryptology – EUROCRYPT 2009*, LNCS 5479, Springer, 2009, 36–53.

[3] E. Bach, R. Peralta, Asymptotic semismoothness probabilities, *Math. Comp.* 65 (1996), 1701–1715.

[4] R. C. Baker, G. Harman, Shifted primes without large prime factors, *Acta Arith.* 83 (1998), 331–361.

[5] P. T. Bateman, R. A. Horn, A heuristic asymptotic formula concerning the distribution of prime numbers, *Math. Comp.* 16 (1962), 363–367.

[6] S. Bai, C. Bouvier, A. Kruppa, P. Zimmermann, Better polynomials for GNFS, *Math. Comp.* 85 (2016), 861–873.

[7] D. Boneh, G. Durfee, Y. Frankel, An attack on RSA given a small fraction of the private key bits, *Advances in Cryptology – ASIACRYPT '98*, LNCS 1514, Springer, 1998, 25–34.

[8] D. J. Bernstein, T. Lange, Computing small discrete logarithms faster, *Progress in Cryptology – INDOCRYPT 2012*, LNCS 7668, Springer, 2012, 317–338.

[9] J. P. Buhler, H. W. Lenstra, Jr., C. Pomerance, Factoring integers with the number field sieve, in *The development of the number field sieve*, LNM 1554, Springer, 1993, 50–94.

[10] R. P. Brent, Some integer factorization algorithms using elliptic curves, *Austral. Comput. Sci. Comm.* 8 (1986), 149–163 (reporting Suyama's parametrisation).

[11] D. A. Burgess, On character sums and primitive roots, *Proc. London Math. Soc.* (3) 12 (1962), 179–192.

[12] P. Bürgisser, On defining integers and proving arithmetic circuit lower bounds, *Comput. Complexity* 18 (2009), 81–103.

[13] N. A. Carella, Note: integer factoring methods III, arXiv:0707.4468 (2007); unrefereed.

[14] N. A. Carella, Deterministic integer factorization algorithms, arXiv:1308.2891 (2013); unrefereed.

[15] Q. Cheng, On the ultimate complexity of factorials, *Theoret. Comput. Sci.* 326 (2004), 419–429.

[16] D. Coppersmith, Modifications to the number field sieve, *J. Cryptology* 6 (1993), 169–180.

[17] D. Coppersmith, Small solutions to polynomial equations, and low exponent RSA vulnerabilities, *J. Cryptology* 10 (1997), 233–260.

[18] N. Coxon, On nonlinear polynomial selection for the number field sieve, preprint, arXiv:1109.6398 (2011).

[19] I. Damgård, M. Koprowski, Generic lower bounds for root extraction and signature schemes in general groups, *Advances in Cryptology – EUROCRYPT 2002*, LNCS 2332, Springer, 2002, 256–271.

[20] N. D. Elkies, C. T. McMullen, Gaps in $\sqrt{n}$ mod 1 and ergodic theory, *Duke Math. J.* 123 (2004), 95–139.

[21] D. El-Baz, J. Marklof, I. Vinogradov, The two-point correlation function of the fractional parts of $\sqrt{n}$ is Poisson, *Proc. Amer. Math. Soc.* 143 (2015), 2815–2828.

[22] T. Estermann, Einige Sätze über quadratfreie Zahlen, *Math. Ann.* 105 (1931), 653–662.

[23] S. Ward, Maestro, *factorlab and latticelab*: code, tests and archived results behind the reports, version 2026.09, https://github.com/iGentAI/factorlab.

[24] M. Filaseta, O. Trifonov, On gaps between squarefree numbers II, *J. London Math. Soc.* 45 (1992), 215–221.

[25] É. Fouvry, Théorème de Brun–Titchmarsh; application au théorème de Fermat, *Invent. Math.* 79 (1985), 383–407.

[26] J. B. Friedlander, Shifted primes without large prime factors, in: *Number Theory and Applications* (Banff, 1988), NATO Adv. Sci. Inst. Ser. C 265, Kluwer, Dordrecht, 1989, 393–401.

[27] Y. Gao, Y. Feng, H. Hu, Y. Pan, On factoring and power divisor problems via rank-3 lattices and the second vector, Cryptology ePrint Archive 2025/1004; to appear in *Math. Comp.*

[28] R. L. Graham, D. E. Knuth, O. Patashnik, *Concrete Mathematics*, 2nd ed., Addison-Wesley, 1994.

[29] S. D. Galbraith, J. M. Pollard, R. S. Ruprai, Computing discrete logarithms in an interval, *Math. Comp.* 82 (2013), 1181–1195.

[30] W. B. Hart, A one line factoring algorithm, *J. Aust. Math. Soc.* 92 (2012), 61–69.

[31] D. Harvey, An exponent one-fifth algorithm for deterministic integer factorisation, *Math. Comp.* 90 (2021), 2937–2950; arXiv:2010.05450.

[32] D. Harvey, M. Hittmeir, A log-log speedup for exponent one-fifth deterministic integer factorisation, *Math. Comp.* 91 (2022), 1367–1379; arXiv:2105.11105.

[33] D. Harvey, M. Hittmeir, Deterministic methods for finding elements of large multiplicative order, arXiv:2601.11131 (2026).
[34] D. R. Heath-Brown, Square-free values of n^2+1, *Acta Arith.* 155 (2012), 1–13.
[35] N. Heninger, H. Shacham, Reconstructing RSA private keys from random key bits, *Advances in Cryptology – CRYPTO 2009*, LNCS 5677, Springer, 2009, 1–17.
[36] M. Herrmann, A. May, Solving linear equations modulo divisors: on factoring given any bits, *Advances in Cryptology – ASIACRYPT 2008*, LNCS 5350, Springer, 2008, 406–424.
[37] M. J. Hinek, Another look at small RSA exponents, *Topics in Cryptology – CT-RSA 2006*, LNCS 3860, Springer, 2006, 82–98.
[38] M. J. Hinek, *On the security of some variants of RSA*, PhD thesis, University of Waterloo, 2007.
[39] M. Hittmeir, A babystep-giantstep method for faster deterministic integer factorization, *Math. Comp.* 87 (2018), 2915–2935.
[40] M. Hittmeir, A time-space tradeoff for Lehman's deterministic integer factorization method, *Math. Comp.* 90 (2021), 1999–2010.
[41] C. Hooley, On the power free values of polynomials, *Mathematika* 14 (1967), 21–26.
[42] N. Howgrave-Graham, Finding small roots of univariate modular equations revisited, *Cryptography and Coding*, LNCS 1355, Springer, 1997, 131–142.
[43] M. N. Huxley, *Area, Lattice Points, and Exponential Sums*, London Math. Soc. Monographs (N.S.) 13, Oxford University Press, 1996.
[44] A. Ya. Khinchin, *Continued Fractions*, University of Chicago Press, 1964.
[45] T. Kleinjung, On polynomial selection for the general number field sieve, *Math. Comp.* 75 (2006), 2037–2047.
[46] P. Koiran, Valiant's model and the cost of computing integers, *Comput. Complexity* 13 (2004), 131–146.
[47] N. Koo, G. H. Jo, S. Kwon, On nonlinear polynomial selection and geometric progression (mod N) for number field sieve, IACR Cryptology ePrint Archive, Report 2011/292 (2011).
[48] J. Kärkkäinen, P. Sanders, S. Burkhardt, Linear work suffix array construction, *J. ACM* 53 (2006), 918–936.
[49] R. S. Lehman, Factoring large integers, *Math. Comp.* 28 (1974), 637–646.
[50] H. W. Lenstra, Jr., Factoring integers with elliptic curves, *Ann. of Math.* 126 (1987), 649–673.
[51] J. D. Lichtman, Primes in arithmetic progressions to large moduli, and shifted primes without large prime factors, arXiv:2211.09641 (2022).
[52] R. J. Lipton, Straight-line complexity and integer factorization, *Algorithmic Number Theory (ANTS-I)*, LNCS 877, Springer, 1994, 71–79.
[53] A. K. Lenstra, H. W. Lenstra, Jr., M. S. Manasse, J. M. Pollard, The number field sieve, in *The development of the number field sieve*, LNM 1554, Springer, 1993, 11–42.
[54] H. W. Lenstra, Jr., C. Pomerance, A rigorous time bound for factoring integers, *J. Amer. Math. Soc.* 5 (1992), 483–516.
[55] C. Lutsko, A. Sourmelidis, N. Technau, Pair correlation of the fractional parts of αn^θ, arXiv:2106.09800 (2021).
[56] C. Lutsko, N. Technau, m-point correlations of the fractional parts of αn^θ, *Amer. J. Math.*, to appear.
[57] U. M. Maurer, Factoring with an oracle, *Advances in Cryptology – EUROCRYPT '92*, LNCS 658, Springer, 1993, 429–436.
[58] J. McKee, R. G. E. Pinch, Further attacks on server-aided RSA cryptosystems, manuscript, 1998. Its abstract (on the second author's publication page) states a method of cost $O(N^{1/4}/\beta)$ for moduli whose $p-1$ and $q-1$ share the factor β; the text itself we could not retrieve, and the statements used here are those restated in [38, §6.3].
[59] G. L. Miller, Riemann's hypothesis and tests for primality, *J. Comput. System Sci.* 13 (1976), 300–317.
[60] P. L. Montgomery, Speeding the Pollard and elliptic curve methods of factorization, *Math. Comp.* 48 (1987), 243–264.
[61] P. L. Montgomery, Small geometric progressions modulo n, unpublished note, 1993; see also the account in the thesis of Murphy above.
[62] B. A. Murphy, *Polynomial selection for the number field sieve integer factorisation algorithm*, PhD thesis, Australian National University, 1999.
[63] M. Nair, Power free values of polynomials, *Mathematika* 23 (1976), 159–183.
[64] A. K. Lenstra, H. W. Lenstra, Jr. (eds.), *The Development of the Number Field Sieve*, Lecture Notes in Math. 1554, Springer, 1993.
[65] I. Nir, Deterministically finding an element of large order in $\mathbb{Z}_n^*$, arXiv:2605.09592 (2026).
[66] K. K. Norton, Numbers with small prime factors, and the least kth power non-residue, *Mem. Amer. Math. Soc.* 106 (1971).

[67] Z. Oznovich, B. L. Volk, On deterministically finding an element of high order modulo a composite, *Proc. ACM-SIAM Symposium on Discrete Algorithms (SODA 2026)*, 6379–6391.
[68] J. M. Pollard, Theorems on factorization and primality testing, *Proc. Cambridge Philos. Soc.* 76 (1974), 521–528.
[69] J. M. Pollard, Theorems on factorization and primality testing, *Proc. Cambridge Philos. Soc.* 76 (1974), 521–528.
[70] T. Prest, P. Zimmermann, Non-linear polynomial selection for the number field sieve, *J. Symbolic Comput.* 47 (2012), 401–409.
[71] M. Radziwiłł, K. Shubin, Poissonian pair correlation for αn^θ mod 1, *Int. Math. Res. Not. IMRN* 2024, no. 9, 7654–7688.
[72] R. L. Rivest, A. Shamir, Efficient factoring based on partial information, *Advances in Cryptology – EUROCRYPT '85*, LNCS 219, Springer, 1986, 31–34.
[73] O. Robert, P. Sargos, Three-dimensional exponential sums with monomials, *J. reine angew. Math.* 591 (2006), 1–20.
[74] J. B. Rosser, L. Schoenfeld, Approximate formulas for some functions of prime numbers, *Illinois J. Math.* 6 (1962), 64–94.
[75] V. Shoup, Lower bounds for discrete logarithms and related problems, *Advances in Cryptology – EUROCRYPT '97*, LNCS 1233, Springer, 1997, 256–266.
[76] M. Shub, S. Smale, On the intractability of Hilbert's Nullstellensatz and an algebraic version of "NP $\neq$ P?", *Duke Math. J.* 81 (1995), 47–54.
[77] R. D. Silverman, The multiple polynomial quadratic sieve, *Math. Comp.* 48 (1987), 329–339.
[78] V. Strassen, Einige Resultate über Berechnungskomplexität, *Jber. Deutsch. Math.-Verein.* 78 (1976/77), 1–8.
[79] N. Technau, N. Yesha, On the correlations of n^α mod 1, arXiv:2006.16629 (2020).
[80] G. Tenenbaum, *Introduction to Analytic and Probabilistic Number Theory*, 3rd ed., Grad. Stud. Math. 163, Amer. Math. Soc., 2015.
[81] M. J. Wiener, Cryptanalysis of short RSA secret exponents, *IEEE Trans. Inform. Theory* 36 (1990), 553–558.

iGent AI
Email address: `sward@igent.ai`

iGent AI