A PRIMITIVE-RATIO CUT AND A 7/12 BOUND FOR THE *abc* EXCEPTIONAL SET

RUIYI ZHANG

ABSTRACT. We prove $N_1(X) \ll_\varepsilon X^{7/12+\varepsilon}$, where $N_\lambda(X)$ counts the *abc* triples of exponent λ in $[1, X]^3$. The previous bound proved in full was $X^{13/20+\varepsilon}$, and Bernert (arXiv:2506.13364, Section 4) sketched $X^{8/13+\varepsilon}$. We keep the anatomic reduction of Bernert, Browning, Lichtman and Teräväinen. It gives 22 inequalities for the exponents of a box. We add two exact rows per side for the root classes of index at least seven, and a cut. The cut fixes all but two coprime products on two sides. The remaining pairs are primitive points of a lattice. Its determinant is a product of roots on the third side. This gives 1 560 195 further inequalities. An exact Farkas certificate over 8192 cells and 10 820 leaves gives exponent 7/12 for the box count. On each leaf the multipliers sum to at most 18/7. This pays for the losses at $\lambda = 1$.

1. INTRODUCTION

For this paper $\mathbb{N} = \{1, 2, \ldots\}$. For $n \neq 0$ put $\mathrm{rad}(n) = \prod_{p|n} p$, and write $\omega(n) = \#\{p : p \mid n\}$, with $\omega(1) = 0$. Following [2], call $(a, b, c) \in \mathbb{N}^3$ with $\gcd(a, b, c) = 1$ an *abc triple of exponent* λ if

$$a + b = c, \qquad \mathrm{rad}(abc) < c^\lambda, \tag{1.1}$$

Let $N_\lambda(X)$ count these triples in $[1, X]^3$. The *abc* conjecture of Masser and Oesterlé says that $N_\lambda(X) = O_\lambda(1)$ for $\lambda < 1$, so $\lambda = 1$ is the endpoint of interest.

The bound $N_\lambda(X) \ll_\varepsilon X^{2\lambda/3+\varepsilon}$ is classical [2, Proposition 1.1]; it comes from de Bruijn's count [4] of integers with restricted radical. For $\lambda \geq 2$, Kane's result [6] gives $N_\lambda(X) \ll_\varepsilon X^{\lambda-1+\varepsilon}$ [2, Section 1]. On $2 \leq \lambda \leq 3$ it matches $N_\lambda(X) \gg X^{\lambda-1-\varepsilon}$. [2, Section 1] counts it in Kane's result and states it for every $\lambda > 1$. It is originally due to unpublished work of Granville; the question is Mazur's. It can hold only for $\lambda \leq 3$, since $N_\lambda(X) \leq \#\{(a, b, c) \in [1, X]^3 : a + b = c\} = O(X^2)$. Browning, Lichtman and Teräväinen [3, Theorem 1.2] gave the first power saving below $2\lambda/3$, for $0.99 < \lambda < 1.001$. Bernert [1, Theorem 1] proved

$$N_\lambda(X) \ll_\varepsilon X^{(23\lambda+3)/40+\varepsilon} \qquad (0 < \lambda < 2), \tag{1.2}$$

so $N_1(X) \ll_\varepsilon X^{13/20+\varepsilon}$. In the merged paper, [2, Theorem 1.2] states (1.2) for $0 < \lambda \leq 2$. Its Theorem 1.3 gives 0.6 for $0 < \lambda < 1$, with a constant depending on λ. At $\lambda = 1$: a remark at the end of [1, Section 4] sketches $8/13 = 0.6153\ldots$; Li [7] obtains $56/85 = 0.6588\ldots$ from the method of [3], its limit without further optimisation by his account.

Theorem 1. *$N_1(X) \ll_\varepsilon X^{7/12+\varepsilon}$, where $7/12 = 0.5833\ldots$*

2020 *Mathematics Subject Classification.* Primary 11D75; Secondary 11N25, 11H06, 11Y16.

Key words and phrases. abc conjecture, exceptional set, radical, anatomy of integers, Farey fractions, linear programming certificate.

Since $N_\mu \le N_1$ for $\mu \le 1$, Theorem 1 also improves [2, Theorem 1.3] on $0 < \mu < 1$, with a constant independent of μ. It is the best known exponent for $61/69 < \mu \le 1$; for $\mu \le 61/69$, (1.2) is at least as good.

Section 2 recalls the anatomic reduction of [3, 1, 2]. Fix $\varepsilon_0 > 0$ and $d = d(\varepsilon_0)$. Every triple lies in a box. In a box, each side is a fixed coefficient times $\prod_{i \le d} \mathfrak{s}_i^{i}$. The roots $\mathfrak{s}_i$ have fixed dyadic sizes. Distinct side values are coprime. Let $\mathcal{B}_d$ be the number of triples in a box and $D = \log_X \mathcal{B}_d$. The proof is the chain

$$N_1(X) \ll (\log X)^2 \max_{Y \le X,\ \mu \le 1} N_\mu^*(Y), \qquad N_\mu^*(Y) \ll Y^{\varepsilon_0} \max \mathcal{B}_d,$$

$$D \le \frac{7}{12} + \frac{18}{7}\rho, \qquad \rho \le 19\varepsilon_0,$$

from (2.1), (2.3), Theorem 2 with Lemma 7.1, and Lemma 7.2. Since $\frac{18}{7} \cdot 19 + 1 = \frac{349}{7}$, choosing ε_0 with $\frac{349}{7}\varepsilon_0 < \varepsilon/2$ proves Theorem 1. The inequalities for D are linear in the 25 exponents

$$(a, b, c;\ a_1, \dots, a_6, b_1, \dots, b_6, c_1, \dots, c_6;\ v_a, v_b, v_c;\ D), \tag{1.3}$$

with $L = a + b + c$. The 22 inequalities of [2, Proposition 6.1] give $D \le 3/5$ there.

We add two kinds of rows. First, the tail (Section 3). The root classes of index at least seven cost X^{τ_s} to enumerate, $\tau_s = s - \sum_{i \le 6} s_i$. Their product is at least X^{v_s}, with $v_s \ge 7\tau_s$. So the tail gives two exact rows per side and can serve as a modulus. Second, the cut (Section 5). Choose a side σ, and let y, z be the other two. Choose a nonempty $T \subseteq \{1, \dots, 6, \mathrm{tail}\}$ of root classes of σ, and $R_y, R_z \subseteq [6]$, not both empty. The tail is never free. Enumerate all other roots. The selected roots form a modulus M. Put $q = \gcd(\gcd R_y, \gcd R_z)$, $\gcd \emptyset = 0$, and

$$u = \prod_{r \in R_y} \mathfrak{y}_r^{r/q}, \qquad v = \prod_{s \in R_z} \mathfrak{z}_s^{s/q},$$

with $\mathfrak{y}_r, \mathfrak{z}_s$ the free roots of y, z; an empty product is 1. Then $u \equiv \xi v \pmod M$ for one of $X^{o(1)}$ units ξ, and

$$\gcd(u, v) = 1, \qquad u \equiv \xi v \pmod M, \qquad 0 < u \le U, \quad 0 < v \le V,$$

a primitive point of a lattice of determinant M; Lemma 4.1 counts these by $1 + 2UV/M$. This is the cut (5.3). There is one per (σ, T, R_y, R_z):

$$3\,(64^2 - 1)\,127 = 1\,560\,195 \tag{1.4}$$

cuts.

Section 6 proves $D \le 7/12$ from these rows. The thirteen min-disjunctions of [2, Proposition 6.1] give 8192 cells. Each cell closes along a binary tree of sign covers. Its leaves carry nonnegative rational Farkas multipliers. Only 424 cells need a cut; 732 distinct cuts occur. On a leaf the multipliers sum to at most 18/7, whence the term $\frac{18}{7}\rho$. Section 7 uses it at $\lambda = 1$. There the row $L \le 1$ of [2] fails. Its proof in [2, Theorem 1.3] needs $\varepsilon_0 \le (1 - \lambda)/3$. So for $\lambda < 1$ the constants, d among them, depend on λ; at $\lambda = 1$ no $\varepsilon_0 > 0$ serves. We use $L \le 1 + 3\varepsilon_0/4$ instead. The L-native Fourier and geometry rows cost $18\varepsilon_0$ and $6\varepsilon_0$. Hence $\rho = 18\varepsilon_0 + o_{\varepsilon_0, d}(1)$.

The proof in [2] does not use the determinant method; [3] used it together with Thue equations. Lemma 4.1 is its simplest case. The cut applies it to two sides at once, with modulus M built from roots on the third side.

Imported statements are cited where used. Citations of [2] refer to arXiv: 2410.12234v2 (to appear in C. R. Math.); each imported statement was checked

against it. The dyadic reduction (2.1) is only stated in [2, Section 2]; Section 2 proves it. Besides the divisor bound and its consequence for $x^j - y^j = n$ [2, Propositions 2.3 and 3.1], only [2, Proposition 4.1] rests on an outside result, [5, Lemma 3]. Nothing else in [2, Sections 2–6] is cited from elsewhere. No conjecture is assumed, and no estimate for rational points on curves is used. The finite step, Theorem 2, is a certificate checked by four independent programs (Appendix B). Every other step is proved in the text.

2. The anatomic reduction

Throughout this section $0 < \lambda \le 2$, $\varepsilon > 0$ is arbitrary, $\varepsilon_0 > 0$ is small and fixed before X, and $d = d(\varepsilon_0)$. Implied constants may depend on $\lambda, \varepsilon, \varepsilon_0, d$, never on the box, the coefficients or X.

Write $N^*_\lambda(X)$ for the number of triples counted by $N_\lambda(X)$ that satisfy in addition $c \ge X/2$ and $\operatorname{rad}(abc) \ge c^\lambda/2$. A dyadic decomposition and the pigeonhole principle give

$$N_\lambda(X) \ll_\lambda (\log X)^2 \max_{1 \le Y \le X} \max_{\mu \le \lambda} N^*_\mu(Y), \tag{2.1}$$

as in [2, Section 2]. The bounds below for $N^*_\mu(Y)$ are uniform in $0 < \mu \le \lambda$ (Lemma 7.2).

Proof of (2.1). Let (a, b, c) be counted by $N_\lambda(X)$. Then $c \ge 2$ and $a, b < c \le X$. Choose $l \ge 1$ with $c \in (X2^{-l}, X2^{1-l}]$ and put $Y = X2^{1-l} \le X$. Thus $Y/2 < c \le Y$. Put

$$\mu_i = \lambda - i/\log_2 X \qquad (0 \le i \le \lceil \lambda \log_2 X \rceil).$$

Here $\log_2 X = (\log X)/(\log 2)$. There are $O(\log X)$ such exponents, all at most λ, and the last is at most 0. With $r = \operatorname{rad}(abc)$, we have $r < c^{\mu_0}$ and $r \ge 1 \ge c^{\mu_m}$ for the last index m. Choose the largest i with $r < c^{\mu_i}$. Then $r \ge c^{\mu_{i+1}} \ge c^{\mu_i}/2$, since $c \le X$. Hence the triple is counted by $N^*_{\mu_i}(Y)$. There are $O(\log X)$ choices of Y and $O(\log X)$ choices of μ_i, which proves the claim. □

For $n \ge 1$ write $[n] = \{1, \dots, n\}$.

Let $d \in \mathbb{N}$ and $(a_0, b_0, c_0) \in \mathbb{N}^3$. Let $\mathbf{A}, \mathbf{B}, \mathbf{C} \in \mathbb{N}^d$ have entries powers of two, 1 included. Put $\mathfrak{a}_0 = a_0$, $\mathfrak{b}_0 = b_0$, $\mathfrak{c}_0 = c_0$ and

$$\mathcal{B}_d = \#\Big\{ (\mathfrak{a}_i, \mathfrak{b}_i, \mathfrak{c}_i)_{i=1}^d \in \mathbb{N}^{3d} : \begin{array}{c} \mathfrak{a}_i \sim A_i,\ \mathfrak{b}_i \sim B_i,\ \mathfrak{c}_i \sim C_i \quad (1 \le i \le d), \\ a_0 \textstyle\prod_{i \le d} \mathfrak{a}_i^i + b_0 \prod_{i \le d} \mathfrak{b}_i^i = c_0 \prod_{i \le d} \mathfrak{c}_i^i, \\ \gcd\big(\prod_{i=0}^d \mathfrak{a}_i, \prod_{i=0}^d \mathfrak{b}_i, \prod_{i=0}^d \mathfrak{c}_i\big) = 1 \end{array} \Big\}. \tag{2.2}$$

For a side write $S_j \in \{A_j, B_j, C_j\}$ for its dyadic root size. Equation (2.2) is the B_d of [2], under our dyadic convention, with no condition within a side. We write $\mathcal{B}_d$ to avoid a clash with the size B_d.

Here $x \sim A$ means $A \le x < 2A$, as in [1]; [2] uses $x \in (A, 2A]$. Both give $x/2 \le A \le x$ and $\#\{x : x \sim A\} = A$, so no inequality of [2] changes. The difference is at $x = 1$, which no interval $(A, 2A]$, $A \in \mathbb{N}$, contains. Roots equal to 1 occur: in [2, Lemma 2.2], $m_j = \prod_{p^j \| n} p$ and $n_j = m_j$ for $j \ne k$, the folded index. So $n_j = 1$ unless some prime has multiplicity exactly j. As n has at most $\omega(n)$ distinct multiplicities, all but at most $\omega(n)$ indices $j \le d$, $j \ne k$, give $n_j = 1$. So in the proof of [2, Proposition 2.1] we read "the unique powers of two such that

$a_i \sim A_i$" with $A_i = 1$ when $a_i = 1$. By [2, Proposition 2.1], each $\varepsilon_0 > 0$ has a $d = d(\varepsilon_0)$ with

$$N^*_\lambda(X) \ll_{\varepsilon_0} X^{\varepsilon_0} \max \mathcal{B}_d, \tag{2.3}$$

the maximum over pairwise coprime $1 \le a_0, b_0, c_0 \le X^{\varepsilon_0}$ and dyadic $A_i, B_i, C_i \in [1, X]$ with

$$\begin{gathered} 2^{-O_d(1)} X^\lambda \le \prod_{j \le d} A_j B_j C_j \le X^{\lambda + 3\varepsilon_0/4}, \\ \prod_{j \le d} A_j^j \le X, \qquad \prod_{j \le d} B_j^j \le X, \qquad X^{1-\varepsilon_0/2} \le \prod_{j \le d} C_j^j \le X, \end{gathered} \tag{2.4}$$

where the power of two depends only on d: by $\prod_j n_j \ge \mathrm{rad}(n)$ below and $c \ge X/2$, $\prod_j A_j B_j C_j \ge 2^{-3d} \mathrm{rad}(abc) \ge 2^{-3d-\lambda-1} X^\lambda$. This suffices, since d is fixed before X. Fix such a box. For a side $s \in \{a, b, c\}$ set

$$s_i = \log_X S_i \ (1 \le i \le d), \qquad s = \log_X \prod_{i \le d} S_i = \sum_{i \le d} \log_X S_i. \tag{2.5}$$

Index 0 is not a coordinate: a_0, b_0, c_0 are the integer coefficients. From here on a, b, c denote coordinates, not the triple. Boxes with $\mathcal{B}_d = 0$ do not attain the maximum in (2.3) unless $N^*_\lambda(X) = 0$. So $\mathcal{B}_d \ge 1$, and we set

$$D = \log_X \mathcal{B}_d \ge 0.$$

For the tail put

$$\tau_s = s - \sum_{i \le 6} s_i = \sum_{7 \le j \le d} \log_X S_j, \qquad v_s = \sum_{7 \le j \le d} j \log_X S_j. \tag{2.6}$$

Put $L = a + b + c$ and $p_i = a_i + b_i + c_i$; a bare p is a prime. Moving a dyadic endpoint at scale X changes the root coordinates by

$$\Delta s_i, \ \Delta s = O_d\left(\frac{1}{\log X}\right). \tag{2.7}$$

This does not cover D. Its resolution loss is in Lemma 7.2 and Table 7.1.

The left inequality of (2.4) uses $\prod_{j \le d} n_j \ge \mathrm{rad}(n)$ with no loss, while [2, Lemma 2.2] states only $\mathrm{rad}(n) \le n^{\varepsilon_0} \prod_j n_j$. In the proof of that lemma $m_j = \prod_{p^j \| n} p$, and, for the folded index k, $n_j = m_j$ for $j \ne k$ and $n_k = m_k \prod_{j > d} m_j^{\lfloor j/k \rfloor}$. Since $d \ge k$, $\lfloor j/k \rfloor \ge 1$ for $j > d$, so

$$\prod_{j \le d} n_j \ \ge \ \prod_{j \le d} m_j \prod_{j > d} m_j \ = \ \mathrm{rad}(n).$$

Lemma 2.1. *The three side values of a box are pairwise coprime. Hence every factor of one side is a unit modulo every factor of another side.*

Proof. A prime dividing two side values divides the third by the equation in (2.2). If it divides $a_0 \prod_{i=1}^d \mathfrak{a}_i^i$, it divides $\prod_{i=0}^d \mathfrak{a}_i$; likewise on the other sides. This contradicts the gcd condition in (2.2). □

The index $i = 0$ in this gcd condition is needed. With $d = 1$, the choice

$$a_0 = b_0 = 1, \quad \mathfrak{a}_1 = \mathfrak{b}_1 = 2, \quad c_0 = 4, \quad \mathfrak{c}_1 = 1$$

satisfies the condition with $i \geq 1$: $\gcd(2,2,1) = 1$. Yet its side values 2, 2, 4 share the prime 2. The full condition excludes it: $\gcd(2,2,4) = 2$.

Nothing is assumed within a side. Roots of one side may share primes with a_0. Unenumerated roots are recovered by a divisor bound, not by squarefreeness.

2.1. **The inherited rows.** In the coordinates (2.5), the propositions of [2] give the 22 inequalities of [2, Proposition 6.1]. We display them in exact limiting form, as Section 6 uses them. At finite resolution each right-hand side rises by the amount in Lemma 7.2 and Table 7.1; (2.9) and (2.10) hold exactly.

The radical bound in (2.4) gives [2, (6.1)],

$$L \leq 1, \tag{2.8}$$

derived there for $\lambda < 1$. At $\lambda = 1$ only $L \leq 1 + 3\varepsilon_0/4$ holds (Lemma 7.2). Roots sum into their side, [2, (6.2)]:

$$\sum_{i \leq 6} s_i \leq s \qquad (s = a, b, c). \tag{2.9}$$

The side-value bound in (2.4) gives [2, (6.3)],

$$7s - 1 \leq 6s_1 + 5s_2 + 4s_3 + 3s_4 + 2s_5 + s_6 \qquad (s = a, b, c). \tag{2.10}$$

Two sides determine the third up to a divisor factor [2, Proposition 2.3]; this is [2, (6.4)],

$$D \leq \min(a + b,\ a + c,\ b + c). \tag{2.11}$$

Geometry of numbers [2, Proposition 4.1] gives the thirteen min-disjunctions [2, (6.5)–(6.9)], labelled (G1)–(G13) in Appendix A. The Fourier estimate [2, Proposition 3.1] gives [2, (6.10)],

$$\sum_{s=a,b,c} \max(s_2 + s_4 + s_6,\ s_3 + s_6,\ s_5) \leq 4L - 6D. \tag{2.12}$$

One argument per maximum gives 27 linear rows; Section 6 uses these.

The `Python` listing of [2, Appendix A] transcribes the first line of [2, (6.9)] as `p1+a2+b3+a3`, where the displayed formula reads $p_1 + a_2 + b_3 + c_3$. We use the displayed formula. It is what [2, Proposition 4.1] gives with the printed second argument $1 - a_2 - 2b_3 - 2c_3$. For (2.10) we use

$$7s - 1 \leq 7s - \sum_{i \leq d} i\, s_i = \sum_{i \leq 6} (7 - i) s_i + \sum_{7 \leq i \leq d} (7 - i) s_i \leq \sum_{i \leq 6} (7 - i) s_i, \tag{2.13}$$

since $7 - i \leq 0$ for $i \geq 7$. In our system (2.10) is the sum of the two tail rows of Section 3, so it is redundant.

3. The multiplicity-at-least-seven tail

The tail rows we use are

$$\tau_s \geq 0, \qquad v_s \geq 7\tau_s, \qquad \sum_{i \leq 6} i\, s_i + v_s \leq 1 \qquad (s = a, b, c). \tag{3.1}$$

The first is (2.9). The second: $\log_X S_j \geq 0$, with weight $j \geq 7$ in v_s and weight 1 in τ_s. The third is the side-value bound in (2.4). Adding the last two,

$$(v_s - 7\tau_s) + (1 - \sum_{i \leq 6} i s_i - v_s) \geq 0,$$

that is, $\sum_{i\le 6} is_i + 7\tau_s \le 1$; with $\tau_s = s - \sum_{i\le 6} s_i$ this is (2.10). With the upper row $v_s \le d\tau_s$, these rows describe the real pairs (τ_s, v_s) exactly. We omit that row, so the system $\mathcal{S}$ of Section 6 is free of d.

Lemma 3.1. *Fix a side $s \in \{a, b, c\}$ and let $M_{\mathrm{tail}} = \prod_{7\le j\le d} \mathfrak{s}_j^j$ with $\mathfrak{s}_j \sim S_j$. Then*

$$\begin{aligned} \#\{(\mathfrak{s}_7, \dots, \mathfrak{s}_d) \in \mathbb{N}^{d-6} : \mathfrak{s}_j \sim S_j\} &\le \prod_{7\le j\le d} S_j = X^{\tau_s}, \\ M_{\mathrm{tail}} &\ge \prod_{7\le j\le d} S_j^j = X^{v_s}, \end{aligned} \tag{3.2}$$

and every prime p dividing M_{tail} does so with $\mathrm{ord}_p(M_{\mathrm{tail}}) \ge 7$, where ord_p denotes the p-adic valuation.

Proof. Each interval $[S_j, 2S_j)$ contains exactly S_j integers, which gives the count, and $\mathfrak{s}_j \ge S_j$ gives the modulus bound. If $p \mid \mathfrak{s}_j$ with $j \ge 7$ then $\mathrm{ord}_p(M_{\mathrm{tail}}) \ge j \ge 7$. □

The count is exact for $[N, 2N)$ and for $(N, 2N]$. A closed $[N, 2N]$ would overcount by a factor of up to 2^{d-6}, attained when all $S_j = 1$.

So the tail costs X^{τ_s} to enumerate. The bound is on roots, not on $\mathrm{rad}(M_{\mathrm{tail}})$, and uses no squarefreeness. On the modulus side σ of Section 5, M_{tail} divides the side value and is at least X^{v_σ}. Hence T may contain the tail.

3.1. **Folded roots.** In [2, Lemma 2.2] the very high multiplicities are folded into one index k. The folded root need not be squarefree; the remainder goes into n_0. The proof of [2, Proposition 2.1] uses that lemma with parameter $\varepsilon_0/4$. So

$$k = \lceil 4/\varepsilon_0 \rceil, \qquad d = 6\lceil (4/\varepsilon_0)^2 \rceil.$$

We may assume $\varepsilon_0 \le 4/7$. Then $k \ge 7$, and $d \ge 6(4/\varepsilon_0)^2 > 4/\varepsilon_0 + 1 > k$. So the folded root lies in the tail, M_{tail} still divides the side value, and $\mathrm{ord}_p(n_k^k) = k\,\mathrm{ord}_p(n_k) \ge 7$ for $p \mid n_k$. The overlap of n_0 and n_k is within one side, so Lemma 2.1 is untouched. Unenumerated roots are recovered by Lemma 4.3. The folding costs two things, both already charged: the factor $X^{3\varepsilon_0/4}$ in the upper radical bound of (2.4), and the choice of a_0, b_0, c_0 in the X^{ε_0} of (2.3). It changes neither Lemma 3.1 nor (3.1).

4. Three lattice and divisor lemmas

Lemma 4.1. *Let $M \ge 1$, let ξ be a unit modulo M and let $U, V \ge 1$. Then*

$$\#\{(u, v) \in \mathbb{N}^2 : u \le U,\ v \le V,\ \gcd(u, v) = 1,\ u \equiv \xi v\ (M)\} \le 1 + \frac{2UV}{M}. \tag{4.1}$$

Proof. Put $\mathcal{L}_\xi = \{(u, v) \in \mathbb{Z}^2 : u \equiv \xi v\ (M)\}$, of index M in $\mathbb{Z}^2$. Any two of its points have determinant in $M\mathbb{Z}$. Positive collinear primitive points coincide. So the primitive points of $\mathcal{L}_\xi$ in $[1, U] \times [1, V]$ have distinct slopes; order them by slope. Consecutive points span with 0 a triangle of area $\ge M/2$. These triangles have disjoint interiors and lie in $[0, U] \times [0, V]$. With P points, $(P - 1)M/2 \le UV$. □

The proof does not use that ξ is a unit. The constant 2 is sharp. For odd $n \ge 3$ take $U = n$, $V = n + 4$, $M = n^2 + 2n - 4$, $\xi = n - 1$. The primitive points are $(n - 1, 1)$, $(n, n + 4)$, $(1, n + 3)$, with consecutive determinants M, and $(P - 1)M/(UV) \to 2$.

Lemma 4.2. *Let $1 \le q \le 6$, let $M \ge 1$ and let K be a unit modulo M. The congruence $\eta^q \equiv K \pmod M$ has at most $8^{\omega(M)}$ solutions in units η, hence $X^{o(1)}$ solutions uniformly for $M \le X^{O(1)}$.*

Proof. Odd prime power: the unit group is cyclic, so the kernel of $\eta \mapsto \eta^q$ has order at most $q \le 6$. Modulus 2^e, $e \le 2$: the unit group has order at most 2. Modulus 2^e, $e \ge 3$: $(\mathbb{Z}/2^e\mathbb{Z})^\times \cong C_2 \times C_{2^{e-2}}$, so the kernel has order $\gcd(q,2)\gcd(q,2^{e-2}) \le 8$; the 2-part of $q \le 6$ is at most 4. Multiply by the Chinese remainder theorem. As $\omega(M) \ll \log M/\log\log M$ for $M \ge 3$,

$$8^{\omega(M)} = \exp\big(O(\log M/\log\log M)\big) = X^{O(1/\log\log X)} \qquad (M \le X^{O(1)}).$$

The form $C^{\omega(M)} \ll_{C,\delta} M^\delta$, for fixed $C \ge 1$ and $\delta > 0$, alone gives no rate. For $M \le 2$ there is one class. No bound $B^{\omega(M)}$ with $B < 8$ holds. For $q = 4$, $e \ge 4$, the group $C_2 \times C_{2^{e-2}}$ has $\gcd(4,2)\gcd(4,2^{e-2}) = 8$ fourth roots of unity. And $q = 4$ occurs in (5.2). Local kernel size 8 occurs only for 2^e, $e \ge 4$, $q = 4$. So for $\omega(M) \ge 2$ the product of local kernel sizes is at most $8 \cdot 6^{\omega(M)-1} < 8^{\omega(M)}$. □

Lemma 4.3. *Write $\operatorname{div}(N) = \#\{m \in \mathbb{N} : m \mid N\}$ for $N \in \mathbb{N}$. For $r \ge 1$ and positive integers $e_1, \dots, e_r$,*

$$\#\{(x_1,\dots,x_r) \in \mathbb{N}^r : x_1^{e_1}\cdots x_r^{e_r} = N\} \le \operatorname{div}(N)^{r-1} \ll_{r,\delta} N^\delta \qquad (\delta > 0),$$

uniformly in $e_1, \dots, e_r$. For $N \ge 3$ the maximal-order bound gives the rate

$$\operatorname{div}(N)^{r-1} = \exp\big(O_r(\log N/\log\log N)\big),$$

so for $N \le X^{O_d(1)}$ and $r \le d$ the count is $X^{O_d(1/\log\log X)}$.

Proof. Each x_i divides N. The equation determines x_r from $x_1, \dots, x_{r-1}$, for any e_i. Apply the divisor bound and, for the rate, the maximal order of $\operatorname{div}(N)$. □

Thus an unfactored product of roots determines those roots up to $X^{o(1)}$.

5. The cut

Fix a box and (σ, T, R_y, R_z). Let (y, z) be the other two sides in alphabetical order: (b,c), (a,c), (a,b) for $\sigma = a, b, c$. Write $[6] = \{1,2,3,4,5,6\}$. Choose a nonempty

$$T \subseteq \{1,2,3,4,5,6,\mathrm{tail}\},$$

$$C_T = \sum_{t \in T \cap [6]} \sigma_t + \mathbf{1}_{\mathrm{tail} \in T}\tau_\sigma, \qquad W_T = \sum_{t \in T \cap [6]} t\,\sigma_t + \mathbf{1}_{\mathrm{tail} \in T} v_\sigma. \tag{5.1}$$

Choose free sets $R_y, R_z \subset [6]$, not both empty, and put

$$q = \gcd(\gcd R_y, \gcd R_z), \qquad F = \sum_{r \in R_y} r\,y_r + \sum_{s \in R_z} s\,z_s, \tag{5.2}$$

with $\gcd \emptyset = 0$. As $R_y \cup R_z \ne \emptyset$, $1 \le q \le 6$ and q divides every element of $R_y \cup R_z$. By Section 3.1, $d = 6\lceil (4/\varepsilon_0)^2 \rceil \ge 294$, so the tail notation applies. Write $\max\{t, 0\} = \max\{t, 0\}$.

Proposition 5.1. *Let a box satisfy the constraints after (2.3), with $\mathcal{B}_d \ge 1$. Then*

$$D \le \Big(y - \sum_{r \in R_y} y_r\Big) + \Big(z - \sum_{s \in R_z} z_s\Big) + C_T + \max\{\frac{F}{q} - W_T, 0\} + o_{\varepsilon_0,d}(1), \tag{5.3}$$

uniformly over the finitely many choices of (σ, T, R_y, R_z).

For later use write

$$\mathsf{A} = \Big(y - \sum_{r\in R_y} y_r\Big) + \Big(z - \sum_{s\in R_z} z_s\Big) + C_T, \qquad \mathsf{B} = F/q - W_T.$$

The limiting form of (5.3) is $D \le \mathsf{A} + \max(0, \mathsf{B})$.

Proof. Sizes. Let σ_0, y_0, z_0 be the coefficients and Σ_i, Y_i, Z_i the dyadic root sizes of the σ, y, z sides. Besides (2.2), the proof uses the dyadic ranges and $\prod_i A_i^i$, $\prod_i B_i^i$, $\prod_i C_i^i \le X$. So all products and cofactors are at most $X^{O_d(1)}$, and Lemmas 4.2 and 4.3 cost $X^{o(1)}$. The coefficients enter no size estimate.

Enumeration. Write m_i, n_i, n_i' for the roots of the σ-, y- and z-sides, so $m_i \sim \Sigma_i$, $n_i \sim Y_i$, $n_i' \sim Z_i$. The coefficients a_0, b_0, c_0 are fixed by the box. Enumerate: the roots of y outside R_y; of z outside R_z; the selected roots of σ. Tails go through Lemma 3.1. The y- and z-tails are never in $R_y, R_z \subseteq [6]$, so they are enumerated too. The cost is

$$X^{(y-\sum_{r\in R_y} y_r)+(z-\sum_{s\in R_z} z_s)+C_T+o_{\varepsilon_0,d}(1)}. \tag{5.4}$$

Modulus. The tail roots of σ are m_j, $j \ge 7$. Put

$$M = \prod_{t\in T\cap[6]} m_t^t \, M_{\mathrm{tail}}^{\mathbf{1}_{\mathrm{tail}\in T}}, \qquad M_{\mathrm{tail}} = \prod_{7\le j\le d} m_j^j. \tag{5.5}$$

Then M divides the σ-side value. Also $m_t \ge \Sigma_t = X^{\sigma_t}$ (left-closed intervals), and Lemma 3.1 bounds the tail. So

$$M \ge X^{W_T}. \tag{5.6}$$

Put

$$u = \prod_{r\in R_y} n_r^{r/q}, \qquad v = \prod_{s\in R_z} (n_s')^{s/q}, \tag{5.7}$$

an empty product being 1; the exponents are integers by (5.2).

Congruence. By Lemma 2.1, modulo M every factor of the y- and z-sides is a unit: coefficients, enumerated roots, u, v. The equation of (2.2) modulo M gives

$$(uv^{-1})^q \equiv K \pmod M \tag{5.8}$$

with

$$K \equiv \epsilon E_z E_y^{-1} \pmod M, \qquad \epsilon = \begin{cases} -1, & \sigma = c, \\ +1, & \sigma \ne c, \end{cases} \tag{5.9}$$

where

$$E_y = y_0 \prod_{i\in[6]\setminus R_y} n_i^i \prod_{7\le j\le d} n_j^j,$$
$$E_z = z_0 \prod_{i\in[6]\setminus R_z} (n_i')^i \prod_{7\le j\le d} (n_j')^j.$$

So K is a unit fixed by the enumerated data.

Count. By Lemma 4.2, $\xi = uv^{-1}$ lies in one of $X^{o(1)}$ classes. By Lemma 2.1, $\gcd(u, v) = 1$. Also

$$u \le 2^{\sum_{r\in R_y} r/q} \prod_{r\in R_y} Y_r^{r/q} =: U, \qquad v \le 2^{\sum_{s\in R_z} s/q} \prod_{s\in R_z} Z_s^{s/q} =: V,$$

with $U, V \geq 1$; the powers of 2 cost $O_d(1/\log X)$, as $\sum_{r\leq 6} r = 21$. Hence

$$U = X^{\sum_{r\in R_y} r\, y_r/q + O_d(1/\log X)}, \qquad V = X^{\sum_{s\in R_z} s\, z_s/q + O_d(1/\log X)},$$

so $UV = X^{F/q+O_d(1/\log X)}$, and Lemma 4.1 with (5.6) bounds the pairs (u, v) in one class by

$$1 + \frac{2UV}{M} \ll_d X^{\max\{F/q - W_T, 0\} + O_d(1/\log X)}. \tag{5.10}$$

Recovery. Given u, v, Lemma 4.3 recovers the free roots up to $X^{o(1)}$. Then the y- and z-side values are known, and the equation fixes the σ-side value. Divide by σ_0 and the selected roots:

$$N = \prod_{i\in[6]\setminus T} m_i^i \left(\prod_{7\leq j\leq d} m_j^j \right)^{1-\mathbf{1}_{\mathrm{tail}\in T}}.$$

Lemma 4.3 recovers these unselected roots up to $X^{o(1)}$. Multiplying (5.4) by $X^{o(1)}$ and (5.10) gives (5.3). □

If one free set is empty, $U = 1$ or $V = 1$, and Lemma 4.1 still applies. The exponent q depends only on the index sets R_y, R_z, not on which roots equal 1. Counting coefficients is charged to (2.3): inside a box a_0, b_0, c_0 are fixed.

Count: 64 subsets of [6] per counted side, 127 sets T, three sides σ. Without the pair of empty free sets this is $3(64^2 - 1)127 = 1\,560\,195$, as in (1.4).

6. The finite system

Let $\mathcal{S}$ be the following system in the 25 nonnegative variables (1.3):

- (2.8), (2.9), (2.10), (2.11), the 27 rows of (2.12), and the last two rows of (3.1) per side: 43 linear rows (the first row of (3.1) is (2.9));
- the thirteen min-disjunctions (G1)–(G13) of Appendix A;
- the 1 560 195 cuts (5.3), without the $o_{\varepsilon_0,d}(1)$.

Here L, p_i and τ_s are linear forms, not coordinates. The coefficients a_0, b_0, c_0 are not coordinates either. Coefficients, dyadic product bounds, tail enumeration and the $X^{o(1)}$ recovery costs are all inside these rows. So the theorem below is a finite statement about (1.3).

Theorem 2. *Every nonnegative solution of $\mathcal{S}$ satisfies $D \leq 7/12$.*

Proof. Write x for the 25-vector in (1.3), and write e_D for the unit vector in its D coordinate. Each min-disjunction holds in one of its two branches; branch 1 of (Gj) is $D \leq L - P_j$ and branch 2 is $D \leq L - Q_j$. For $\eta \in \{1, 2\}^{13}$ let the cell $\mathcal{C}_\eta$ be the intersection of the selected branches. A feasible point lies in at least one of these 8192 cells; cells may overlap. In a cell, a cut $D \leq \mathsf{A} + \max(0, \mathsf{B})$ is covered exactly by

$$\{D \leq \mathsf{A} + \mathsf{B}\} \;\cup\; \{\mathsf{B} \leq 0,\ D \leq \mathsf{A}\}. \tag{6.1}$$

For each cell the certificate gives a full binary tree. Completeness is recursive: either the sole path is empty, or all paths begin with one cut descriptor, both + and − children occur, and the same holds for both suffix sets. The checker enforces: distinct leaf paths; no leaf a prefix of another; no descriptor repeated on a path. Let $x \in \mathcal{S}$ lie in the cell. By (6.1), x descends at each node into a child. So it

reaches a leaf whose rows all hold at x. A + edge gives $D \le \mathsf{A} + \mathsf{B}$; a − edge gives $\mathsf{B} \le 0$ and $D \le \mathsf{A}$. At $\mathsf{B} = 0$ either edge is allowed.

At a leaf ℓ of $\mathcal{C}$, the set $I(\mathcal{C}, \ell)$ consists of: the 43 linear rows; the thirteen branches of $\mathcal{C}$; the rows its path authorizes. No other row is admissible. All rows of $I(\mathcal{C}, \ell)$ hold at x. The certificate gives $\lambda_i \ge 0$ in $\mathbb{Q}$ on $I(\mathcal{C}, \ell)$, and $\lambda_i = 0$ elsewhere. Write each row as $\langle g_i, x\rangle \le h_i$, and put $V_\ell = \sum_{i \in I(\mathcal{C},\ell)} \lambda_i h_i$. The checker verifies at every leaf

$$\sum_i \lambda_i g_i \ge e_D \text{ componentwise}, \qquad \sum_i \lambda_i h_i \le \frac{7}{12}. \tag{6.2}$$

All variables being nonnegative,

$$D = \langle e_D, x\rangle \le \Big\langle \sum_i \lambda_i g_i, x \Big\rangle \le \sum_i \lambda_i h_i \le \frac{7}{12}. \tag{6.3}$$ □

The 43 linear rows alone give $D \le \frac{2}{3}$ (bound each maximum in (2.12) below by 0). Adding the disjunctions gives the exact optimum 3/5, the 0.6 of [2, Theorem 1.3]. Adding the cuts gives 7/12. Gains: 1/15 from the disjunctions, 1/60 from the cut, 1/12 in total.

A leaf by hand. Take branch 2 of (G2) and (G3), branch 1 of (G5), any branch of the others. Multipliers: 1 on (2.8), 1/4 on the Q arms of (G2) and (G3), 1/2 on the P arm of (G5). Then

$$\sum_i \lambda_i g_i = e_D + \tfrac{1}{2}(e_{a_1} + e_{b_1} + e_{c_1}) \ge e_D, \qquad \sum_i \lambda_i h_i = \tfrac{1}{2} \le \tfrac{7}{12}, \qquad \sum_i \lambda_i = 2 \le \tfrac{18}{7}.$$

The coefficients of a, b, c, p_2, a_3, b_3 cancel; D has 1 and p_1 has 1/2. So this leaf gives $D \le 1/2$. Three bits are fixed, so the vector is valid on $2^{10} = 1024$ cells; the certificate uses it on 200.

The certificate is `cert_leaves_tail.jsonl` (Appendix B). Rebuild every row from (2.8)–(2.12), (3.1), (5.3) and (G1)–(G13); check (6.2) exactly. This gives Table 6.1. The $o_{\varepsilon_0,d}(1)$ dropped from (5.3) is paid in Lemma 7.2. Put

$$\Lambda := \max_{\text{certificate leaves}} \sum_i \lambda_i,$$

the sum over all multipliers at the leaf, sign rows $\mathsf{B} \le 0$ included.

TABLE 6.1. The certificate of Theorem 2. A no-cut cell has a sole empty path; otherwise it has a nontrivial complete cut trie. The certificate value is V_ℓ from the proof.

branch cells $\eta \in \{1,2\}^{13}$	8192
cells with a sole empty path (no cut)	7768
cells with a nontrivial complete cut trie	424
leaves	10 820
internal binary nodes	2628
Λ	18/7
distinct cuts (5.3) used	732
worst certificate value	7/12
leaves with certificate value exactly 7/12	89
minimum over leaves and coordinates of $\big(\sum_i \lambda_i g_i - e_D\big)_k$	0

6.1. **The optimum.** The bound 7/12 is the exact optimum of $\mathcal{S}$. For $0 \le t \le 1/24$, let $x(t)$ be the point

$$\begin{aligned} a = \tfrac{1}{3},\ a_1 = \tfrac{1}{8},\ a_4 = \tfrac{5}{24},\ v_a = t; \qquad b = \tfrac{1}{4},\ b_1 = \tfrac{1}{8},\ v_b = \tfrac{7}{8}; \\ c = \tfrac{5}{12},\ c_1 = \tfrac{1}{8},\ c_3 = \tfrac{7}{24}; \qquad D = \tfrac{7}{12}, \end{aligned} \tag{6.4}$$

with every unlisted coordinate zero. The restriction $t \le 1/24$ is $a_1 + 4a_4 + v_a = 23/24 + t \le 1$. For each such t: $L = 1$; the 43 linear rows hold; one arm of each (Gj) holds. All 1 560 195 cuts hold, with minimum bound exactly 7/12. Here $t = v_a$ enters (5.3) only via W_T ($\sigma = a$, tail $\in T$), where it lowers B. So each cut bound is nonincreasing in t, and $t = 0$, $t = 1/24$ suffice. At $t = 0$ the minimum is attained 170 335 times. The cut family is $b \leftrightarrow c$ symmetric, and $b \leftrightarrow c$ maps $x(1/24)$ to

$$\begin{aligned} a = \tfrac{1}{3},\ a_1 = \tfrac{1}{8},\ a_4 = \tfrac{5}{24},\ v_a = \tfrac{1}{24}; \qquad b = \tfrac{5}{12},\ b_1 = \tfrac{1}{8},\ b_3 = \tfrac{7}{24}; \\ c = \tfrac{1}{4},\ c_1 = \tfrac{1}{8},\ v_c = \tfrac{7}{8}; \qquad D = \tfrac{7}{12}, \end{aligned} \tag{6.5}$$

with every unlisted coordinate zero. There it is attained 227 551 times. The points $x(0)$ and (6.5) lie in different cells. Their arms differ in (G3), (G4), (G5), (G7), (G9), (G10), (G12), (G13). These are points of $\mathcal{S}$, not abc triples; they give no lower bound for $N_1(X)$.

7. Stability and the endpoint

Lemma 7.1. *Let $\rho \ge 0$ and $x \ge 0$. Suppose that, with right-hand sides raised by ρ,*

- *every linear row of $\mathcal{S}$ holds at x;*
- *at least one arm of each min-disjunction holds at x;*
- *every cut holds at x in the form $D \le \mathsf{A} + \max(0, \mathsf{B}) + \rho$.*

The sign rows $\mathsf{B} \le 0$ of the descent are case distinctions and are not relaxed. Then

$$D \le \frac{7}{12} + \frac{18}{7}\rho. \tag{7.1}$$

Proof. Each relaxed min-disjunction has a valid arm; choose a relaxed cell. The relaxed sign cover descends to a leaf as in Theorem 2. Then (6.3) with relaxed right-hand sides adds

$$\rho \sum_{i\ \text{non-sign}} \lambda_i \le \rho \sum_i \lambda_i \le \Lambda\rho.$$

Table 6.1 gives $\Lambda = 18/7$. For a relaxed cut the cover (6.1) becomes $\{D \le \mathsf{A} + \mathsf{B} + \rho\} \cup \{\mathsf{B} \le 0,\ D \le \mathsf{A} + \rho\}$. The sign row stays exact, so nothing is relaxed twice. □

Lemma 7.2. *Let $0 < \mu \le 1$. Choose $Y_0 = Y_0(\varepsilon_0, d)$ such that, for $Y \ge Y_0$, each displayed loss $O_d(1/\log Y)$, $O_d(1/\log\log Y)$, $o_{\varepsilon_0,d}(1)$ is at most ε_0. Let $Y \ge Y_0$, and take a nonempty box admissible in* (2.3) *at scale Y. Then every row of $\mathcal{S}$ holds in a μ-free form. Its right-hand side rises by at most the amount in Table 7.1. In particular* (2.8) *becomes $L \le 1 + 3\varepsilon_0/4$. Consequently*

$$\rho = 18\varepsilon_0 + o_{\varepsilon_0,d}(1) \tag{7.2}$$

is admissible in Lemma 7.1. Increasing Y_0 if necessary, $\rho \le 19\varepsilon_0$ for $\mu \le 1$ and $Y \in [Y_0, X]$.

TABLE 7.1. The relaxation carried by each family of $\mathcal{S}$.

family	source	raised by
(2.8)	[2, (6.1)]	$3\varepsilon_0/4$
(2.9), (2.10), (3.1)	definitions and (2.4)	0
(2.11)	Lemma 4.3	$O_d(1/\log\log Y)$
Appendix A	[2, Proposition 4.1]	$6\varepsilon_0 + o_{\varepsilon_0,d}(1)$
(2.12)	[2, Proposition 3.1], D-coefficient 6	$18\varepsilon_0 + o_{\varepsilon_0,d}(1)$
(5.3)	Proposition 5.1	$o_{\varepsilon_0,d}(1)$

Proof. Run Section 2 at scale Y with λ replaced by μ and $\log_Y$ in place of $\log_X$ in (2.5). In (2.4), μ occurs only in the radical constraint. Its upper half is

$$L \le \mu + 3\varepsilon_0/4 \le 1 + 3\varepsilon_0/4, \tag{7.3}$$

the one row of $\mathcal{S}$ that depends on μ. Its lower half is

$$\mu \le L + O_d\Big(\frac{1}{\log Y}\Big). \tag{7.4}$$

The proof of [2, Theorem 1.3] uses the radical bound in the weaker form

$$\prod_{j\le d} A_j B_j C_j \le Y^{\mu+3\varepsilon_0} \le Y.$$

This needs $\varepsilon_0 \le (1-\mu)/3$, which fails at $\mu = 1$ for every $\varepsilon_0 > 0$. It is the only use of $\mu < 1$ in the proof of [2, Theorem 1.3], and the source of its μ-dependent constant. We keep (7.3) as a relaxed row instead.

By (7.4), μ leaves the other rows. The other constraints of (2.4) are μ-free. So (2.9), (2.10), (3.1) and Lemma 3.1 are μ-free and exact. The constants in [2, Propositions 2.1, 3.1, 4.1] depend only on ε_0 and d.

Pair row. [2, Proposition 2.3] has no rate in Y; we prove (2.11) directly. There are exactly $\prod_i A_i B_i = Y^{a+b}$ choices of the a- and b-roots. With c_0 fixed, the equation in (2.2) then gives $N = \prod_i \mathfrak{c}_i^i \le 2^{O_d(1)}Y$, and Lemma 4.3 leaves $\mathrm{div}(N)^{d-1} = Y^{O_d(1/\log\log Y)}$ choices. Likewise for the other two pairs. So (2.11) costs $O_d(1/\log\log Y)$, as in Table 7.1.

Cut. Proposition 5.1 uses only the dyadic ranges, (2.2) and the fixed-d lemmas. So it holds verbatim at scale Y for every $\mu \le 1$. Its divisor and root-class factors cost $O_d(1/\log\log Y)$.

Fourier and geometry. The ε_0-exponents are $3\varepsilon_0$ in [2, Proposition 3.1], and $4\varepsilon_0$, $6\varepsilon_0$ in [2, Proposition 4.1]. There μ appears as $Y^{2\mu/3}$, resp. Y^{μ} and $Y^{\mu-1}$. By (7.4), $\mu \to L$ costs $4 \cdot O_d(1/\log Y)$, resp. $O_d(1/\log Y)$. Our Fourier row is θ times the source inequality, with

$$\theta = \frac{[D]_{\mathrm{row}}}{[D]_{\mathrm{source}}} = \frac{6}{1} = \frac{[L]}{[\mu]} = \frac{4}{2/3} = \frac{1}{1/6} = 6$$

on the D, L and root terms. So $3\varepsilon_0$ becomes $18\varepsilon_0$. The rows (G1)–(G13) have coefficient 1 on D and L, so they rise by at most $6\varepsilon_0$.

These are the entries of Table 7.1; the largest is (7.2). All constants depend only on ε_0 and d, not on μ. The Y-dependent errors are the displayed $O_d(1/\log Y)$, $O_d(1/\log\log Y)$ and $o_{\varepsilon_0,d}(1)$ terms. They include the constants of [2, Propositions 3.1, 4.1], as in Appendix A. $Y \ge Y_0$ absorbs them.

□

Lemma 7.2 removes the restriction $\lambda < 1$ of [2, Theorem 1.3]. At $\mu = 1$ the row (2.8) is lost, and Lemma 7.1 prices the relaxation ρ at $\frac{18}{7}\rho$ in the exponent. The constant does not depend on μ. $\mathcal{S}$, its certificate and $18/7$ are free of ε_0 and d. So choosing ε_0 from $18/7$ is not circular.

8. Proof of Theorem 1

Fix $\varepsilon > 0$. Choose $\varepsilon_0 > 0$ with $\frac{349}{7}\varepsilon_0 < \varepsilon/2$ and with folded index at least seven (Section 3.1). This fixes $d = d(\varepsilon_0)$. The factors $8^{\omega(M)}$ and $\mathrm{div}(N)^{O_d(1)}$ are $Y^{o(1)}$ for fixed d (Lemmas 4.2, 4.3); they need no budget. Dyadic constants and the count of boxes are inside (2.3).

Let $\mu \le 1$ and $Y \le X$. If $\mu \le 0$, then $N^*_\mu(Y) = 0$, as $\mathrm{rad}(abc) \ge 1 \ge c^\mu$. Let $0 < \mu \le 1$. If $Y < Y_0(\varepsilon_0, d)$, then $N^*_\mu(Y) \le Y^3 \ll_{\varepsilon_0,d} 1$. If $Y \ge Y_0$, then for a nonempty box Lemma 7.2 raises every row of $\mathcal{S}$ by at most $\rho \le 19\varepsilon_0$. So the stability and reduction losses are at most

$$\frac{18}{7}\,19\varepsilon_0 + \varepsilon_0 = \frac{349}{7}\varepsilon_0 < \frac{\varepsilon}{2}.$$

Thus Theorem 2 and Lemma 7.1 give

$$\mathcal{B}_d \le Y^{7/12+(342/7)\varepsilon_0}. \tag{8.1}$$

By (2.3) at scale Y, at a further cost Y^{ε_0}, $N^*_\mu(Y) \ll_\varepsilon Y^{7/12+\varepsilon/2}$. The constant depends only on ε. Finally (2.1) with $\lambda = 1$ absorbs $(\log X)^2$ and the maxima into the remaining $\varepsilon/2$. □

Appendix A. The thirteen min-disjunctions

We work at scale X with exponent λ; Lemma 7.2 uses the same derivation at scale Y with μ. Put $p_i = a_i + b_i + c_i$, $L = a + b + c$. Each row is [2, Proposition 4.1] for one triple $I, I', I'' \subseteq [6]$. A sum of two terms is at most twice the larger, so

$$D \le \max(\lambda + 4\varepsilon_0 - P, \lambda + 6\varepsilon_0 - Q) + O_{\varepsilon_0,d}(1/\log X),$$

where $P = \sum_{i\in I} a_i + \sum_{i\in I'} b_i + \sum_{i\in I''} c_i$ and $Q = 1 - \sum_{i\in I}(i-1)a_i - \sum_{i\in I'}(i-1)b_i - \sum_{i\in I''}(i-1)c_i$. With $C_{\varepsilon_0,d}$ the constant of the source estimate, the constant and the factor 2 cost $\log(2C_{\varepsilon_0,d})/\log X = O_{\varepsilon_0,d}(1/\log X)$. By the lower bound in (2.4) and (2.5),

$$2^{-O_d(1)}X^\lambda \le \prod_{j\le d} A_j B_j C_j = X^L, \qquad \lambda \le L + O_{\varepsilon_0,d}(1/\log X).$$

Hence $\min(P, Q) \le L - D + 6\varepsilon_0 + O_{\varepsilon_0,d}(1/\log X)$. The certificate uses the limiting form below, free of ε_0; the slack $6\varepsilon_0 + O_{\varepsilon_0,d}(1/\log X)$ is charged in Lemma 7.2 and Table 7.1. The choices, from the proof of [2, Theorem 1.3]:

row 1 :	$I = I' = I'' = \{1,2,4\}$;
rows 2–4 :	$I = I' = \{1,2\},\ I'' = \{1,2,3\}$ and permutations;
rows 5–7 :	$I = \{1,2\},\ I' = I'' = \{1,2,3\}$ and permutations;
rows 8–10 :	$I = I' = \{1,2\},\ I'' = \{1,3\}$ and permutations;
rows 11–13 :	$I = \{1,2\},\ I' = I'' = \{1,3\}$ and permutations.

Written out:

(G1)	$P_1 = p_1 + p_2 + p_4,$	$Q_1 = 1 - p_2 - 3p_4$
(G2)	$P_2 = p_1 + p_2 + a_3,$	$Q_2 = 1 - p_2 - 2a_3$
(G3)	$P_3 = p_1 + p_2 + b_3,$	$Q_3 = 1 - p_2 - 2b_3$
(G4)	$P_4 = p_1 + p_2 + c_3,$	$Q_4 = 1 - p_2 - 2c_3$
(G5)	$P_5 = p_1 + p_2 + a_3 + b_3,$	$Q_5 = 1 - p_2 - 2a_3 - 2b_3$
(G6)	$P_6 = p_1 + p_2 + b_3 + c_3,$	$Q_6 = 1 - p_2 - 2b_3 - 2c_3$
(G7)	$P_7 = p_1 + p_2 + c_3 + a_3,$	$Q_7 = 1 - p_2 - 2c_3 - 2a_3$
(G8)	$P_8 = p_1 + a_3 + b_2 + c_2,$	$Q_8 = 1 - 2a_3 - b_2 - c_2$
(G9)	$P_9 = p_1 + b_3 + a_2 + c_2,$	$Q_9 = 1 - 2b_3 - a_2 - c_2$
(G10)	$P_{10} = p_1 + c_3 + a_2 + b_2,$	$Q_{10} = 1 - 2c_3 - a_2 - b_2$
(G11)	$P_{11} = p_1 + a_2 + b_3 + c_3,$	$Q_{11} = 1 - a_2 - 2b_3 - 2c_3$
(G12)	$P_{12} = p_1 + b_2 + a_3 + c_3,$	$Q_{12} = 1 - b_2 - 2a_3 - 2c_3$
(G13)	$P_{13} = p_1 + c_2 + a_3 + b_3,$	$Q_{13} = 1 - c_2 - 2a_3 - 2b_3$

and each row is used as the disjunction $D \le L - P_j$ or $D \le L - Q_j$.

Appendix B. Computer checks

External inputs, supplied with the paper: the certificate tree, its leaf scan, the stability scan, the two witness scans. Everything else is in the text, including the witness coordinates of Section 6.1.

The displayed system alone gives the 8192 cells, the cut-free optimum 3/5, and the root-cell scan: 7768 cells need no cut, 424 do. The file `cert_leaves_tail.jsonl` gives the rest. It holds the cut descriptors and their order on paths, 10 820 leaves, 2628 internal nodes, 732 distinct cuts among the 1 560 195, and the multipliers. All arithmetic is exact. No program imports a generator, a numerical package or an LP solver.

Format: one metadata record, then one record per leaf with exactly the fields `cell`, `path`, `kind`, `lam`.

- `cell`: a list in $\{1,2\}^{13}$; entry j (zero-based) selects `("dis",j,cell[j])`.
- `path`: a list of `[cut-descriptor,sign]` pairs.
- `lam`: a sparse list of `[row-key,"p/q"]`, multipliers positive.
- `kind`: always `"bound"`; only the fourth program reads it.

A cut descriptor is `[sigma,R_y,R_z,T]`: R_y, R_z increasing subsets of [6], not both empty; T a nonempty increasing subset of $\{1,\dots,6,\text{tail}\}$; $(y,z) = (b,c),(a,c),(a,b)$ for $\sigma = a,b,c$. The quantities $q, C_T, W_T, F, \mathsf{A}, \mathsf{B}, \tau_s$ are those of (5.1), (5.2) and after (5.3). A path step is a descriptor and a sign, e.g. `[["b",[3],[3],[4,7]],"+"]`.

Admissible rows at a leaf in cell $\mathcal{C}$: the 43 linear rows; the thirteen rows selected by $\mathcal{C}$; the rows its path authorizes, namely `cutP` for a + step and `cutMA`, `cutMB` for a − step. Nothing else is admissible. The text defines every row and (6.2). It fixes neither the tree nor a numbering of cuts; the file names each cut by its descriptor. The four programs below are the certificate checks; the statistics scripts are not.

program	what it checks
`verify_712_independent.py`	Tables 6.1 and B.2; each leaf's rows are valid on its cell and path; (6.2) at every leaf; that the declared target is 7/12 and that no leaf's certificate value exceeds it; the points $x(0)$ and (6.5) of Section 6.1 against all 1 560 195 instances of (5.3). It also runs four smoke tests, corruptions it must reject: a malformed cell, a missing branch (a deleted sibling leaf), a negative multiplier and a sign-mismatched row. Soundness rests on (6.2) at every leaf with the admissibility and complete-trie checks of Section 6.
`verify_cert_tail_independent.py`	the same cell and leaf structure and the same Farkas conditions, from a second implementation sharing no code with the first
`check_stability_constant.py`	$\Lambda = 18/7$ of Lemma 7.1, and the largest number of nonzero multipliers on a leaf; it reads the multipliers only and is not a validity check
`blind_fourth/verify_blind_third.py`	a code-independent reconstruction from the displayed formulas, checking row authorization, complete binary path tries, the Farkas conditions and the stability scan; its self-test rejects malformed certificates

Table B.1 is generated from the verifier's row labels. Here $s \in \{a, b, c\}$, $(x, y) \in \{(a, b), (a, c), (b, c)\}$, $(i_a, i_b, i_c) \in \{0, 1, 2\}^3$. The branch tag is $b \in \{1, 2\}$, not the side b. The program indexes the tail by 7. In the disjunction row, the exact index

TABLE B.1. Verifier row labels and their paper sources.

verifier label	definition and source
`("L",)`	(2.8), from [2, (6.1)]
`("62",s)`	(2.9), from [2, (6.2)]
`("63",s)`	(2.10), from [2, (6.3)]
`("T7",s)`	$v_s \geq 7\tau_s$, the second row of (3.1)
`("TV",s)`	$\sum_{i \leq 6} i s_i + v_s \leq 1$, the third row of (3.1)
`("64",x,y)`	(2.11) for the pair (x, y), from [2, (6.4)]
`("F",i_a,i_b,i_c)`	(2.12), where on side s the index $i_s = 0, 1, 2$ selects respectively $s_2 + s_4 + s_6$, $s_3 + s_6$, and s_5, from [2, (6.10)]
`("dis",j,b)`	$0 \leq j \leq 12$: the P arm of (G_{j+1}) if $b = 1$, and the Q arm if $b = 2$. The order is (6.5), (6.6)a/b/c, (6.7)ab/bc/ca, (6.8)a/b/c, (6.9)a/b/c.
`("cutP",sigma,R_y,R_z,T)`	$D \leq \mathsf{A} + \mathsf{B}$, the plus child of (6.1).
`("cutMB",sigma,R_y,R_z,T)`	$\mathsf{B} \leq 0$, the sign row of the minus child.
`("cutMA",sigma,R_y,R_z,T)`	$D \leq \mathsf{A}$, the bound row of the minus child.

correspondence is `("dis",j,b)` $\leftrightarrow (G_{j+1})$ for $j = 0, \ldots, 12$. The verifier stores $e_2 = 1 - Q_{j+1}$, since its second branch is $1 - e_2 \leq L - D$.

All $2^{13} = 8192$ arm combinations are represented, with no pruning among the thirteen min-disjunctions. Every cell is non-empty, so none is infeasible; some cells contain others and are checked all the same. Within each cell (6.1) is applied to selected cut rows.

In each cell the paths must form a complete binary trie, as in the proof of Theorem 2. Duplicate paths and prefix leaves are rejected. Hence $10\,820 = 8192 + 2628$, as for any full binary forest with 8192 trees. Table B.2 gives further statistics; the proof uses none of them.

TABLE B.2. Further statistics of the certificate.

largest number of distinct cuts in one cell	35
largest number of leaves in one cell	42
largest path depth	13
distinct multiplier vectors	4545
largest number of nonzero multipliers on one leaf	21
distinct multiplier denominators	4009
largest multiplier denominator	29 322 517
nonzero multipliers in the whole certificate	126 499
each of the thirteen disjunctions used in both branches	yes
second-worst certificate value	9154/15693

TABLE B.3. SHA-256 of the files this paper depends on.

file	SHA-256
`cert_leaves_tail.jsonl`	`0a92f82f24e3dc4bca87f4592633ba ceebaa6a91423048561deb23b7031fd5e7`
`verify_712_independent.py`	`2e40f9527368c3d85b5553615f2445f 81409bfe4c0a68267646d4e41ac4a2b75`
`verify_cert_tail_independent.py`	`226e2ed54b620983df5dd311fae529 3d40e33055ce1270cb2637c8910d5059fe`
`check_stability_constant.py`	`0e491fb351fafebb4445268a29703bd9 ec00fced4a56db64a4c4d4d04c582451`
`blind_fourth/verify_blind_third.py`	`e733c750ce60e64236008d185e1725e0 83cb7961f58b7868117428c8a989f316`

STATEMENTS AND DECLARATIONS

Competing interests. The author has no competing interests.

Funding. No funding was received for this work.

Data availability. The certificate, the four checking programs of Appendix B and their outputs accompany this paper; their SHA-256 values are in Table B.3.

REFERENCES

[1] C. Bernert, *The exceptional set in the abc conjecture*, preprint (2025), arXiv:2506.13364v1.

[2] C. Bernert, T. Browning, J. D. Lichtman, and J. Teräväinen, *Bounds on the exceptional set in the abc conjecture*, to appear in C. R. Math.; arXiv:2410.12234v2, revised 9 May 2026.

[3] T. Browning, J. D. Lichtman, and J. Teräväinen, *Bounds on the exceptional set in the abc conjecture*, preprint (2024), arXiv:2410.12234v1; earlier version of [2].

[4] N. G. de Bruijn, *On the number of integers $\leq x$ whose prime factors divide n*, Illinois J. Math. **6** (1962), 137–141, `https://doi.org/10.1215/ijm/1255631814`.

[5] D. R. Heath-Brown, *Diophantine approximation with square-free numbers*, Math. Z. **187** (1984), 335–344, `https://doi.org/10.1007/BF01161951`.

[6] D. M. Kane, *On the number of ABC solutions with restricted radical sizes*, J. Number Theory **154** (2015), 32–43, `https://doi.org/10.1016/j.jnt.2014.09.021`.

[7] R. Li, *On the exceptional set in the abc conjecture*, preprint (2025), arXiv:2507.02885v1.

INDEPENDENT RESEARCHER, COLUMBIA, MARYLAND, USA
Email address: `ryzfsu@gmail.com`