Series · Markdown source

II. The cyclotomic host and the ordinary prime ramp

1. The field and its character group
2. The auxiliary quadratic primes
   - Progression counts and the exceptional kernel
   - The second moment
   - The maximal estimate and the two discards
   - Centering at 1 and selecting both primes
   - Selecting a pair with prescribed quadratic class numbers
3. From prime selection to real cutoffs
   - A conditional proposition for prescribed starts at conductor one
4. Fitting the geometry at the same cutoff
   - Arithmetic data for the host
   - Thick digits and translated orbits
   - Tops and weighted Frobenius-orbit estimates
   - Carrying the sets into the geometry
5. The finite prime-power identity and the thick mean
   - The comparison error
   - The geometric upper bound needed at the interface

Write $\Lambda(n)$ for the von Mangoldt function and γ for Euler's constant. The ordinary prime ramp at a real $x \geq 1$ is

$$G_z(x) = \sum_{1 \leq n \leq x} \Lambda(n)(x/n - 1) - x \log x + (1 + \gamma)x.$$

For a primitive nonprincipal Dirichlet character χ, write $\ell_\chi = L'(1, \chi)/L(1, \chi)$ and

$$G_\chi(x) = \sum_{1 \leq n \leq x} \Lambda(n)\chi(n)(x/n - 1) + x\ell_\chi.$$

We derive a bound for the difference between $G_z(x)$ and the normalized translate average of the finite prime sum minus its main term. The average uses a cyclotomic CM field depending on x. Its geometric upper bound is an input from the companion papers. All release references below are pinned to commit `e7c8972c701d3396947808ab20cc16dc14623d33`. Ordinary ramp · Primitive ramp

span_of_dilates gives the general CM balance identity, the conditional per-host readout with its size bracket, and a statement-level account of the two-stage prime selection in rStation paper 192 v1, Theorems 4.7, 4.9 and 4.17. Its Lemma 4.20 outlines the centering step, for which we use the finite-tail estimate of Paper I. We specialize the same released argument to conductor one. Section 2 proves the progression and moment estimates in this family, including the growing-modulus step raised in the public review, and shows that the

good pair can also satisfy $v_2 h(-4qd) = v_2 h(-8qd) = 3$. Section 4 constructs the thick and top sets and proves their orbit estimates. Section 5 writes the full translate readout as four signed ramps and retains the geometric bound as a hypothesis. Paper 192 v1 · Public review

1. The field and its character group

Proposition 1 (proved here). Let $e \geq 2$ be an integer. Let q, d be primes with $q \equiv 3 \bmod 8$, $d \equiv 7 \bmod 16$, $q \mid d+1$, and $q \neq 3$. Put $B = 2^{e+2}$, $J_r(a) = (a/r)$ for $r = q, d$, and let

$$H = (\mathbb{Z}/B\mathbb{Z})^\times/\{\pm 1\}, \qquad E = \mathbb{Q}(\zeta_B)^+(\sqrt{-q}, \sqrt{-d}), \qquad K = \mathbb{Q}(\sqrt{-d}), \qquad F = E^+.$$

Then H is cyclic of order $n = 2^e$, E is CM with $[E:\mathbb{Q}] = 4n$, and E has conductor Bqd. In the source's arithmetic frame the dyadic level is $QS = B$ and the level index o is one. Field input and output · Kernel characterization · Frame and Artin map

Proof. The three moduli B, q, d are coprime. Classical cyclotomic Galois theory, the quadratic Gauss sum, and the Chinese remainder theorem give a surjection

$$(\mathbb{Z}/Bqd\mathbb{Z})^\times \longrightarrow (\mathbb{Z}/2)^2 \times H, \qquad a \longmapsto \big(\epsilon_q(a), \epsilon_d(a), [a]\big), \quad (-1)^{\epsilon_r(a)} = J_r(a).$$

Its fixed field is E. Equivalently its kernel consists of the classes with $a \equiv \pm 1 \bmod B$ and $J_q(a) = J_d(a) = 1$. This is exactly the fixed-field condition `IsChiField` for the character of conductor one. An ambient cyclotomic field of any larger admissible modulus gives the same fixed field by reduction of units. The group H is generated by the class of 5: the identity $v_2(5^{2^j} - 1) = j + 2$ follows by repeated squaring, and the 2^e distinct powers of 5 give all units congruent to $1 \bmod 4$. A faithful character of H detects $1 + B/2$ and hence has conductor B. The characters J_q, J_d have conductors q, d. Their least common multiple with B is the conductor of E. In the release $B \mid QS \mid \operatorname{lcm}(B, 1)$, which also gives $QS = B$. Complex conjugation negates both square roots and fixes the real cyclotomic field. Its fixed field $F = \mathbb{Q}(\zeta_B)^+(\sqrt{qd})$ is totally real, and E/F is totally imaginary quadratic. Artin map · Dyadic divisibility · CM type

For the averaging argument write the two bits in the order (ϵ_d, ϵ_q). The source's Artin map stores the q-sign first and the d-sign second, and corrects its last coordinate by the class of the q-sign. This correction is trivial in H at conductor one. Thus in our coordinates the group is $G = (\mathbb{Z}/2)^2 \times H$, complex conjugation is $c = ((1,1),0)$, and $\Phi = \{((0,b),h) : b \in \mathbb{Z}/2,\ h \in H\}$ is the group over K. It indexes the CM embeddings after a choice sending $\sqrt{-d}$ to $i\sqrt{d}$. Its cardinality is $f = 2n$. In the release $H_0 \subset H$, the subgroup used for translates, has index o. Since the powers of the conductor-one character form a singleton, $o = 1$, the kept set is $\{\mathbf{1}\}$, and $H_0 = H$. Every translate has full top share, because $t - u \in H_0$ for all $t, u \in H$ and the prescribed share is $4/(3o+1) = 1$. Original coordinate order · CM type · Colour and top coordinates · Kept set and shares · Index

Ramification and splitting (proved here). For $p \nmid Bqd$, the three components of Frobenius are the two Legendre signs and $[p]$. The inertia at d is the order-two subgroup generated by $((1,0),0)$, since the other two factors are unramified there. Both inertia and a lift of Frobenius at 2 and q belong to Φ: K is unramified and split there by $(-d/2) = 1$ and $(-d/q) = 1$, the latter from $q \mid d+1$. For every odd prime $p \mid d+1$ the same congruence gives $(-d/p) = 1$. For $r = q, d$, quadratic reciprocity identifies J_r with the primitive character of discriminant $-r$, including the value at 2 by the supplementary law. The q-row

therefore has the eigenvector $\sqrt{-q}$ with the specified sign, $QS \le n^2$ for $e \ge 2$, and the condition on odd primes dividing QS is vacuous. These are the ordinary specialization of the field and splitting rows. Pair to splitting rows · Field rows · Host inertia conventions

If ϑ is a character of G, its primitive Dirichlet character is the product of the primitive representative of its H-factor and the indicated powers of J_q, J_d. Indeed the product agrees on the units by the Artin map, and primitive characters inducing the same common-level character agree. Equivalently at p^k its value is the average of ϑ over $k\,\mathrm{Frob}_p + I_p$, where I_p is inertia and Frob_p is a lift. If the restriction to I_p is trivial, this is its unramified value. Otherwise the average vanishes by orthogonality and the primitive local Euler factor is absent. In particular the principal representative takes value 1 also at $2, q, d$. These primitive values give the same L-functions and centers at 1 as the corresponding host characters. Primitive value and balance definitions · Primitive twist comparison

2. The auxiliary quadratic primes

Fix $0 < \varepsilon < 1/52$ and a real $A \ge 1$. At a real scale $Y > 1$ put $z = Y^{\varepsilon/5}$, $z' = z(\log Y)^2$, $U_1 = e^z$, $U_2 = e^{z'}$. Choose q in $(\sqrt{U_1}, U_1]$ with $q \equiv 3 \bmod 8$. For each such q take d in $(\sqrt{U_2}, U_2]$ in the reduced class $a \bmod M$ specified by

$$M = 16q \prod_{3 \le p \le z,\ p \text{ prime}} p, \qquad a \equiv 7 \bmod 16, \qquad a \equiv -1 \bmod M/16.$$

Eventually $q > z$, $d > q$, and $\log M \ll z$ uniformly by the elementary Chebyshev bound for $\sum_{p \le z} \log p$. The boxes use heights $Q_0 = 8U_1$ and $Q_1 = 8U_1U_2$. For every sufficiently large Y the released selection gives the following inequalities for one pair, uniformly at every real $y \in [Y, Y^A]$:

$$|G_\chi(y)| \le y^{1/2+\varepsilon/2}, \qquad |\ell_\chi| \le 4 \log y \quad (\chi = J_q, J_d, J_qJ_d). \tag{1}$$

It also gives $L(s, \rho J_q^i J_d^j) \ne 0$ for $\Re s > 399/400$, $|\Im s| \le Q_j$, for every primitive real ρ of conductor dividing 8 and every $(i, j) \in \{0,1\}^2 \setminus \{(0,0)\}$. All these products are primitive at their product conductors. The extra real bases are needed for estimates involving the full host dual. Progression · Window and box predicates · Export of estimates

Progression counts and the exceptional kernel

Published inputs. For a modulus $D \ge 8$ divisible by 8, apply Thorner and Zaman's relative prime-counting theorem to $\mathbb{Q}(\zeta_D)/\mathbb{Q}$, with the full Galois group as its abelian subgroup. The fixed field is $\mathbb{Q}$, and the largest character conductor is at most D. A reduced class b is a singleton Frobenius class. For $\log U \gg \log D$, the theorem gives

$$\pi(U; D, b) = \frac{\mathrm{Li}(U) - \chi_D(b)\,\mathrm{Li}(U^{\beta_D})}{\varphi(D)}(1 + O(\mathcal{R}_D(U))), \qquad \mathcal{R}_D(U) = e^{-c\log U/\log D} + e^{-c\sqrt{\log U}}.$$

Here $\mathrm{Li}(U) = \int_2^U dt/\log t$, and the term with the real nonprincipal character χ_D and its real zero β_D is omitted if the theorem specifies no such pair. The character at level D is induced by a primitive real character. The release proves a relative progression estimate of the same form under $C \log D \log(\log D + 2) \le \log U$. The argument below fits this latter range as well. We also use Page's theorem in its conductor-uniform form: among primitive real characters of conductor $\le V$, at most one can

have a zero in $(1 - c_P/\log V, 1)$ for an absolute $c_P > 0$. Finally Siegel's theorem gives $1 - \beta \gg f^{-1}$ for a primitive real character of conductor f, with an ineffective constant. At a larger modulus the additional Euler factors have no zeros at a positive real β, so both zero estimates apply to the inducing character. Thorner–Zaman, Theorem 1.4 · Page's theorem · Page's theorem in the density paper, Corollary 2.5 · Released progression bound · Release assembly · Released Page argument · Siegel input

Write $L = \log Y$, $Z_* = (Y^A)^{400}$, $N_* = \lfloor Z_* \rfloor$ and $V = m_0 Z_*^2$. In stage one let $(m_0, a_0, U, \varepsilon_*) = (8, 3, U_1, \varepsilon)$. In stage two, with q fixed in its candidate set, use $(M, a, U_2, 6\varepsilon/5)$. In each case $m_0 > 0$, $8 \mid m_0$ and a_0 is reduced. The stage parameters are $0 < \varepsilon_* \le 1/4$, $K = A \ge 1$, $\tau = \varepsilon/2 > 0$ and any fixed discard share $0 < \nu < 1$. The height constraints are

$$\log U \ge \tfrac{1}{4} L^2 \log m_0, \qquad Y^{\varepsilon_*/6} \le \log U \le Y^{\varepsilon_*/5} L^2.$$

For the fixed modulus 8, these constraints follow from $\log U_1 = z$ and $z/L^3 \to \infty$. For the growing modulus M, $\log M \le \log 16 + \log q + z \log 4 \le 4z$ eventually. Hence $\frac{1}{4} L^2 \log M \le zL^2 = \log U_2$. Also $Y^{\varepsilon_*/6} = z$ and $Y^{\varepsilon_*/5} = zY^{\varepsilon/25}$ at the second stage, giving the other two bounds. In addition, the statement of `stage` requires every prime divisor p of m_0 to satisfy $p \le w = Y^{\varepsilon_*/5}$ or $p > Z_*$. The only prime divisor at stage one is 2, and at stage two 2 and the primes in the product are at most $z \le w$. The remaining prime q satisfies $\log q > z/2 > 400AL = \log Z_*$ eventually. Thus the divisor primes that can support a term of the short sum are at most w, as used in the second moment below. Stage statement · Scale and modulus conditions · Modulus bound · Second-stage verification

At stage one $D = 8s \le 8Y^{800A} = 8(\log U_1)^{4000A/\varepsilon}$, so ordinary Siegel–Walfisz is available at this stage. Its log-power saving can be chosen to absorb both the $\varphi(s)$ classes in the fiber and the desired factor $Y^{-(A+1)}$. At stage two even the base modulus exceeds every fixed power of $\log U_2$, since $\log M \ge \log q > z/2$ and $\log\log U_2 = (\varepsilon/5)L + 2\log L$. We use the relative estimate above at the full modulus $D = m_0 s$ in both stages, as in the release's `twistK`. Indeed $\log D \le \log V = \log m_0 + 800AL$. Eventually $\log V \le z/L^2$ at stage one and $\log V \le 5z$ at stage two, so $\log U/\log V \ge L^2/5$ in either case. Since $\log\log U = O(L)$ and $\log U \ge z \gg L^4$, this gives $C \log D \log(\log D + 2) \le \log U$ for any fixed C and $\mathcal{R}_D(U) = O(e^{-c' L^2})$. The same bounds give $\sqrt{U} > N_*^2$ eventually, so every r in $P = \{r : \sqrt{U} < r \le U,\ r \equiv a_0 \bmod m_0\}$ exceeds N_*^2. These estimates are uniform in q with ε, A fixed. Released relative progression estimate · Range comparison · Public review of paper 192

Lower bound for the candidate sets (proved here). At $D = m_0$, if the exceptional term is subtracted, then

$$\mathrm{Li}(U) - \mathrm{Li}(U^\beta) \ge \frac{U(1 - U^{\beta-1})}{\log U} \gg \frac{U}{\log U} \min(1, (1-\beta)\log U).$$

The first inequality follows by integrating on $[U^\beta, U]$, where the possible exceptional β is near 1 and $U^\beta \ge 2$ eventually. With $f \le m_0$, Siegel's bound gives $\min(1, (1-\beta)\log U) \gg m_0^{-1}$. Dividing by $\varphi(m_0) \le m_0$ yields a lower bound $\gg U/(m_0^2 \log U)$ for the main term in all three cases. The relative error is $o(1)$ even in the subtractive case. Because $\log m_0/\log U = O(L^{-2})$, subtracting at most $\sqrt{U}$ small primes leaves $|P| \ge U^{9/10}$ eventually. The estimate holds uniformly in the allowed first primes. Released candidate count

Progression correlation (proved here). For odd squarefree $1 < s \le N_*^2$ coprime to m_0, let $\xi_s(b) = (b/s)$ on $(\mathbb{Z}/s\mathbb{Z})^\times$. Its conductor is exactly s, since its local character at each prime dividing s is

nontrivial. On the reduced classes modulo $D = m_0 s$ lying over a_0, sum the progression theorem weighted by $\xi_s(b)$. The ordinary main term sums to zero. The Chinese remainder theorem writes $\chi_D = \chi_{m_0}\chi_s$, with the first factor constant on this fiber. The exceptional main term sums to zero unless $\chi_s = \xi_s$, by character orthogonality on the second factor. In the remaining case the base conductor of χ_D has prime-to-m_0 part exactly s. Thus at most one $s = s_*$ can have both this matching condition and $1 - \beta_D < c_P/\log V$, by Page's theorem for the inducing characters of conductor $\le V$. The Page threshold is common to all s, and the exceptional kernel is independent of the short-sum coefficients and cutoff y. Release exceptional-kernel argument · Orthogonality on the progression

For each $s \neq s_*$, the exceptional term either cancels or has $\mathrm{Li}(U^{\beta_D})/\mathrm{Li}(U) \ll (\log U)e^{-c_P \log U/\log V}$, by $\mathrm{Li}(U) \gg U/\log U$ and $\mathrm{Li}(U^{\beta_D}) \ll U^{\beta_D}$. This ratio and $\mathcal{R}_D(U)$ are both $O(e^{-c''L^2})$ after reducing $c'' > 0$. To control the error relative to the actual number of primes, sum the *nonnegative* main terms over the fiber. When the exceptional term matches ξ_s, its unweighted sum is zero, so the total main term is $\mathrm{Li}(U)/\varphi(m_0)$. If the exceptional term does not match, the weighted main sum is already zero and the sum of the absolute errors is at most $O(\mathcal{R}_D(U))$ times the total main term, even when the latter is small. In both cases the total main term is within a factor $1 + O(\mathcal{R}_D(U))$ of the unweighted prime count in that fiber up to U. The count can omit primes dividing s, all of which are below $\sqrt{U}$. Removing the other primes through $\sqrt{U}$ changes either sum by at most $\sqrt{U}$, with $\sqrt{U}/|P| \le U^{-2/5}$. Thus for all sufficiently large Y, for every such $s \neq s_*$,

$$\left|\sum_{r\in P}(r/s)\right| \le |P|\,Y^{-(A+1)}. \tag{1a}$$

This is the bound needed from `twistK`. Correlation statement · Normalization by the actual progression count · Final subtraction

The second moment

Proposition (proved here). There is an absolute C such that, at either stage for every sufficiently large Y, every real $Y \le y \le Y^A$, and every complex sequence $|\kappa(n)| \le 1$, one has

$$\sum_{r\in P}|S_r(y)|^2 \le Cy\log y\,|P|, \qquad S_r(y) = \sum_{1\le n\le N_*} c_y(n)\kappa(n)J_r(n), \qquad c_y(n) = \Lambda(n)\min(1, y/n). \tag{1b}$$

The onset can depend on ε, A and is uniform in the allowed q at the second stage. This is the ordinary specialization of the released moment, including its uniformity over arbitrary bounded coefficients. Moment statement · Source moment expansion

For n, n' in the sum, let $s(n,n')$ be the product of the primes outside m_0 occurring to odd exponent in nn'. Since $r > N_*^2$, no prime dividing nn' equals r. Quadratic reciprocity and the supplement at 2 give

$$J_r(n)J_r(n') = \omega(n,n')\,(r/s(n,n')),$$

The sign $\omega(n,n')$ is constant on P, since the residues of r modulo 8 and modulo each $p \mid m_0$ are fixed. Each $s(n,n')$ is odd, squarefree, coprime to m_0, and at most N_*^2. On expanding the square and applying $|\kappa| \le 1$, the upper bound divided by $|P|$ is

$$\sum_{s(n,n')=1} c_y(n)c_y(n') + \sum_{s(n,n')=s_*} c_y(n)c_y(n') + Y^{-(A+1)}\left(\sum_n c_y(n)\right)^2.$$

The first two terms use the trivial bound $|P|$ for the correlation, with the second term zero when Page's exception is absent. The final term uses (1a). By the Chebyshev estimate $\sum_{n\le t}\Lambda(n) \ll t$ and partial summation, $\sum_{n\le N_*} c_y(n) \ll y(1+\log(Z_*/y))$. Thus the last term is $O_A(y(\log Y)^2/Y)$, which is $\le y\log y$ times an absolute constant eventually. Squarefree kernel · Source application of the correlation bound

Only prime powers have positive mass. Put

$$W_p = \sum_{p^j\le N_*,\, j\ge 1} c_y(p^j), \qquad R = \sum_{\substack{p^j\le N_*,\, j\ge 1\\ p\mid m_0 \text{ or } 2\mid j}} c_y(p^j).$$

For $p\le y$, the powers $p^j\le y$ cost at most $\log y$, and the remaining geometric series at most $2\log p$. For $p>y$, $W_p\le 2y\log p/p$. So $\sup_p W_p \ll \log y$. The same-prime contribution is bounded by $\sum_p W_p^2 \ll y\log y$. Indeed for $p\le\sqrt{y}$ use the bound $O(\sqrt{y}\log^2 y)$, for $\sqrt{y}<p\le y$ use $\log^2 y\le 2\log y\log p$ and Chebyshev, and for $2^j y<p\le 2^{j+1}y$, $j\ge 0$, use

$$\sum_{2^j y<p\le 2^{j+1}y}\frac{\log^2 p}{p^2} \ll \frac{\log(2^{j+1}y)}{2^j y}.$$

The latter estimate follows by taking one factor $\log p/p^2$ at its upper bound and using Chebyshev on the remaining $\log p$. Its sum gives the desired contribution for $p>y$. Diagonal estimate

The mass of all even powers is at most $\sum_{h\ge 1}\Lambda(h)\min(1,y/h^2)\ll\sqrt{y}$, since $\Lambda(h^2)=\Lambda(h)$ and Chebyshev applies first to $h\le\sqrt{y}$ and then on the dyadic intervals above it. The powers on primes dividing m_0 cost $O(w\log y)$ by the individual prime bound and the bound on primes that divide m_0 and occur in the sum. Thus $R\ll\sqrt{y}+w\log y\ll\sqrt{y}$ because $w\le y^{1/5}$. Even powers and modulus primes

To classify the pairs, each prime power contributes a kernel label. It is 1 for the mass in R, or its prime base p for an odd power with $p\nmid m_0$. The pair kernel multiplies labels modulo squares. It equals 1 only for two label-1 terms or two terms with the same prime label, giving a total weight at most $R^2+\sum_p W_p^2$. If s_* is prime, its only pairs join that label to label 1 and have weight at most $2W_{s_*}R$. If s_* is a product of two distinct primes, its pairs join those two labels and have weight at most $2(\sup_p W_p)^2$. No other s_* occurs among the pairs. In particular the exceptional weight is always $O(\log^2 y+\sqrt{y}\log y)$. Combining these estimates proves (1b). Kernel pair cases · Released moment assembly

The maximal estimate and the two discards

For $[T,W]\subset[Y,Y^A]$ with $W\le 2T$, the finite ramp $S_r(y)$ is continuous and absolutely continuous on the block, with right derivative

$$S_r'(v) = v^{-1}\sum_{n\le N_*} c_v(n)\kappa(n)\mathbf{1}_{n>v}J_r(n).$$

The bounded coefficients in (1b) may depend on v. Thus almost everywhere $\sum_{r\in P}|S_r'(v)|^2\le C(\log v)|P|/v$. Absolute continuity and Cauchy–Schwarz bound each square

increment from T, uniformly at every point in the block, by $(W-T)\int_T^W |S_r'(v)|^2\,dv$. Hence

$$\sum_{r\in P}\max_{T\le y\le W}|S_r(y)|^2 \le 2\sum_{r\in P}|S_r(T)|^2 + 2(W-T)\int_T^W \sum_{r\in P}|S_r'(v)|^2\,dv \le 4CT\log(2T)|P|. \tag{1c}$$

Put $\tau=\varepsilon/2$. The fraction violating $|S_r(y)|\le y^{1/2+\tau}/2$ anywhere on one block is $O(T^{-2\tau}\log(2T))$. Summing over the blocks beginning at $T=2^jY$ and stopped at Y^A gives $O_\varepsilon(Y^{-\tau})$, using $\log(2T)\ll_\tau T^\tau$ and the geometric series. This is the source's finite-block argument and works at both endpoints of the full interval. Derivative and maximal argument · Dyadic sum

Thorner and Zaman count zeros with multiplicity over the primitive characters of conductor at most Q, including a possible exceptional real zero. Their $N(\sigma,Q)$ counts $\Re\rho>\sigma$ and $|\Im\rho|\le Q$. For $Q\ge 3$ and $\sigma\ge 39/40$, Theorem 1.2 gives $N(\sigma,Q)\le 10^{88}(10^{421}Q^{99})^{1-\sigma}$. A character failing the box contributes at least one zero. Even if the box includes its left edge $\sigma_0=399/400$, every such zero is counted at each $39/40<\sigma<\sigma_0$. Taking $\sigma\uparrow\sigma_0$ on the right bounds the number of discarded characters by $10^{88}(10^{421}Q^{99})^{1/400}=O(Q^{99/400})$. Since $99/400<1/4$, this gives the required $O(Q^{1/4})$ discard. The release obtains a bound of this order by zero detection with a mollifier. Published definition (1.2) and Theorem 1.2 · Released box count

At stage one the coefficient families are the primitive representatives ρ of the characters modulo 8. At stage two they are the primitive products ρJ_q^i with $i=0,1$ for the already selected q. Their conductors f_0 are respectively $\le F_1=8$ and $\le F_2=8U_1$. There are at most four and eight coefficients, respectively. In the fresh variable r, the product with J_r is primitive of conductor f_0r by coprimality and the Chinese remainder theorem. With $Q=FU$, $F=F_i$, one has $F\le\sqrt{U}$ eventually. Choose one bad character of conductor f_0r for each r that fails the box for any coefficient. A conductor cannot witness two different primes $r>\sqrt{U}$, since both would then divide the conductor and each cofactor would exceed F. Hence the density discard has relative size $O((FU)^{1/4}/U^{9/10})=O(U^{-21/40})$. By (1c) and the finite union over coefficients, the total fraction removed in either stage tends to zero. Coefficient families · Primitive twist and union · Conductor injection

Centering at 1 and selecting both primes

The release uses the contour bound `tail_contour` for the centering step. Paper I, Proposition 3 proves the finite-tail bound needed at the two stage heights directly from the classical local partial-fraction estimate. The release's contour inequality below assumes a nonprincipal χ, conductor at most Q, the box above $399/400$ through height $Q\ge 2$, $Z\ge 1$, $0<h\le Z$, and a fixed $0<\delta\le 1/800$. Put $\mathcal{S}_\chi(t)=\sum_{n\le t}\Lambda(n)\chi(n)(t/n-1)$. In this range the bound is

$$|\mathcal{S}_\chi(Z+h)-\mathcal{S}_\chi(Z)+h\ell_\chi| \ll_\delta \log Q\left((\log Q)hZ^{-1/400+\delta}+Z^2/Q\right).$$

Its proof in the release shifts a smoothed Perron difference to $\Re s=399/400+\delta$, using a bound for L'/L inside the box and a truncation at height $Q/2$. The following calculation records the use of this stronger contour inequality in the release. The proof via Proposition 3 of Paper I gives the required prime selection independently of this inequality. The elementary difference quotient of $\mathcal{S}_\chi$ differs from $\sum_{n\le Z}\Lambda(n)\chi(n)/n$ by at most $(h+1)\log(Z+h)/Z$, by summing only the new integers in $(Z,Z+h]$. Consequently the finite tail $\mathcal{T}_\chi(Z)=-\ell_\chi-\sum_{n\le Z}\Lambda(n)\chi(n)/n$ satisfies

$$|\mathcal{T}_\chi(Z)| \ll_\delta (\log Q)^2 Z^{-1/400+\delta} + Z^2 \log Q/(hQ) + (h+1)\log(Z+h)/Z. \tag{1d}$$

Tail contour estimate · Elementary difference

At $X_* = Y^A$ take $Z = X_*^{400}$, $h = Z/X_*^2$, $\delta = \varepsilon_*/(1600A)$. The stage heights obey $Y^{\varepsilon_*/6} \le \log Q \le Y^{\varepsilon_*/4}$ eventually. Multiplying (1d) by X_*, the first term is $O_\delta(Y^{3\varepsilon_*/4})$, the second is $O_\delta(X_*^{403}\log Q/Q)$, and the third is $O(\log X_*/X_*)$. Their sum is at most Y^{ε_*} eventually. For the primitive product $\chi(n) = \kappa(n)J_r(n)$, finite splitting at $y \le X_*$ gives

$$G_\chi(y) = -S_r(y) - y\,\mathcal{T}_\chi(Z).$$

Because $\varepsilon_* \le 6\varepsilon/5 < 1/2$, the tail costs at most $y^{1/2+\tau}/2$ at large Y, leaving the short sum controlled by the maximal estimate. Release tail parameters · Box from short sum

For the logarithmic derivative bound use (1d) again with $Z = (\log Q)^{1000}$, $h = Z^{1-1/200}$, and $\delta = 1/4000$. The first term is $O((\log Q)^{-1/4})$, the second is $O((\log Q)^{1006}/Q)$, and the third is $o(1)$. The classical bound $\sum_{n\le Z}\Lambda(n)/n \le \log Z + O(1)$ then gives $|\ell_\chi| \le 1000\log\log Q + O(1)$. At both stage heights $\log\log Q \le (\varepsilon/5)\log Y + 2\log\log Y + O(1)$. Since $200\varepsilon < 4$, this gives $|\ell_\chi| \le 4\log Y \le 4\log y$ eventually. This use of the original ε follows from the actual choices of U_1, U_2. The enlarged parameter $\varepsilon_* = 6\varepsilon/5$ is used for admissibility of the second stage. Released derivative estimate · Center and primitive handoff

Alternatively, after the ramp bound at the left endpoint Y has been obtained with Proposition 3 of Paper I, the identity defining the ramp and Chebyshev's bound give $|\ell_\chi| \le 4\log Y$ directly. Finite-tail proof and center estimate

We have proved the two relative share assertions used below. For every fixed $0 < \nu < 1$, every sufficiently large real Y has a first-stage subset of share at least $1-\nu$. For each q in that subset, its own second-stage good set has share at least $1-\nu$. The thresholds are uniform in q. The selection control comes from separate box and maximal discards at each stage. Choosing q first and then d gives (1), using the trivial ρ, and gives all the real-base boxes using the full primitive family from modulus 8. These are the conductor-one contents of `GoodQ`, `GoodD` and `PairGood`. Quantified shares · Stage union · Final output

Selecting a pair with prescribed quadratic class numbers

Proposition 2 (proved here using classical and published inputs). For every ε, A as above, for every sufficiently large Y the pair in (1) can be chosen with $v_2 h(-4qd) = v_2 h(-8qd) = 3$, where $h(D)$ is the class number of the imaginary quadratic field of fundamental discriminant D. This refinement uses Stevenhagen's Rédei reciprocity and the Thorner–Zaman effective Chebotarev theorem. It strengthens the released pair choice for this specialization. Good fractions · Rédei reciprocity, §§3–4 and 7 · Effective Chebotarev, Theorem 1.4

Proof. Write $q = u^2 + 2v^2$ with $u, v > 0$. The norm-Euclidean ring $\mathbb{Z}[\sqrt{-2}]$ gives this representation because $(-2/q) = 1$. For a root $w^2 = q \bmod d$, define $\sigma(q,d) = (\frac{q+uw}{d})$. The root exists because $(d/q) = (-1/q) = -1$ and quadratic reciprocity gives $(q/d) = 1$. The two possible arguments have product $2qv^2 \bmod d$, a nonzero square since $d > q$ and $(2/d) = 1$. The sign is independent of the root. Pair congruences

Use the three ramified primes $2, q, d$ as columns of the Rédei matrices, and order the rows by the discriminants $-4, -q, -d$ or $-8, -q, -d$. Encoding $+1$ as 0, the two matrices are

$$\begin{pmatrix} 1 & 1 & 1 \\ 1 & 1 & 1 \\ 0 & 0 & 0 \end{pmatrix}, \qquad \begin{pmatrix} 1 & 0 & 1 \\ 1 & 0 & 1 \\ 0 & 0 & 0 \end{pmatrix}.$$

For example, $(-q/2) = -1,\ (-d/2) = 1,\ (-q/d) = (d/q) = -1,\ (-d/q) = 1$. The other off-diagonal entries follow from the signs of $-4, -8$ at q, d, and each column sums to zero. Genus theory gives 2-rank two and 4-rank one, so the 2-primary class group has the form $C_2 \times C_{2^a}$, $a \geq 2$. In both fields the product of the ramified ideals over $2, d$ is the nonzero element of $\mathrm{Cl}[2] \cap 2\mathrm{Cl}$. Its vector $(1, 0, 1)$ belongs to each kernel and differs from the principal relations $(0, 1, 1)$ and $(1, 1, 1)$ from the quadratic radicals. The genus character corresponding to $-d$ spans the genus characters admitting a quartic lift. The one-entry 8-rank matrices therefore have entries $[-d, q, 2d]$ and $[-d, 2q, 2d]$ in additive Rédei notation. Genus and 8-rank formulas

Rédei reciprocity and linearity identify both entries with $[q, 2q, -d]$. In detail, the second-kind identities give $[-d, 2, 2d] = 0$ and $[-d, q, qd] = 0$, so

$$[-d, 2q, 2d] = [-d, q, 2d] = [-d, q, 2q] = [q, 2q, -d].$$

Here the required Hilbert pairings are $(-d, q), (-d, 2), (q, 2d), (q, 2q)$. At the odd primes their values are $+1$ by $(-d/q) = (q/d) = (2/d) = 1$, $(2d/q) = (2/q)(d/q) = 1$, and $(q, 2q) = (q, -2)$ with $(-2/q) = 1$. The real symbols are trivial, so the product formula gives the dyadic values. Also $(-d, d) = 1$ and $(2, 2d) = (2, -d) = 1$. These identities verify the pairings for both applications of linearity. The three discriminants in each triple have common gcd one. The second-kind decompositions $-8d = (-d) \cdot 8$ and $-4qd = (-d) \cdot 4q$ give the two vanishing symbols by Proposition 7.10. For the two class-group entries, the associated cyclic quartic extension of the quadratic base is unramified. The Artin action of the ideal of norm $2d$ in Definition 4.4 therefore gives the symmetric symbol of Definition 7.8 with second squareclass q or $2q$. Rédei reciprocity and the vanishing symbols, Theorem 7.9 and Proposition 7.10

Let $k = \mathbb{Q}(\sqrt{2})$, $D = k(\sqrt{q})$, and $L = D(\sqrt{q + u\sqrt{q}})$. The norm $2qv^2$ of $q + u\sqrt{q}$ shows that neither that element nor its quotient by 2 is square in $\mathbb{Q}(\sqrt{q})$, hence $[L : D] = 2$. Over k, the map $\sqrt{q} \mapsto -\sqrt{q}$, $\sqrt{q + u\sqrt{q}} \mapsto v\sqrt{2q}/\sqrt{q + u\sqrt{q}}$ has square equal to the nontrivial automorphism ι over D. The involution negating $\sqrt{2}$ while fixing $\sqrt{q}$ and the chosen square root conjugates that order-four map to its inverse. This gives the dihedral extension over $\mathbb{Q}$ for the pair $(q, 2q)$. Away from $2q$ the conjugate radicands have disjoint prime supports and their product has even valuation, giving the required unramified condition. The discriminants $4q, 8q$ are respectively $4, 0 \bmod 8$, so Definition 7.6 imposes no further condition at 2 for this pair. The field is totally real since both radicands are positive. Thus the infinite-place factor in the Rédei symbol for the negative third argument $-d$ is trivial. The prime d splits in D, and its relative Artin symbol in L/D is nontrivial precisely when $\sigma(q, d) = -1$. An entry 1 gives 8-rank zero, hence $a = 2$ in each of the two 2-class groups. Dihedral interpretation, Definitions 7.6 and 7.8

To count this sign let $C = \mathbb{Q}(\zeta_M)$, with M from the second stage. It contains D, and $C \cap L = D$ because D is the maximal abelian subfield of L. The automorphism σ_a of C fixes D, so its two lifts $g_+ = (1, \sigma_a)$ and $g_- = (\iota, \sigma_a)$ to LC are central. The group $H' = \mathrm{Gal}(LC/k)$ is abelian. In H' their quotient is a square. Indeed choose the order-two automorphism b of C acting by -1 on the q-factor and trivially on the complementary factor. It fixes i, $\sqrt{2}$ and negates $\sqrt{-q}$, hence negates $\sqrt{q}$. Pair b with a generator r of

$\mathrm{Gal}(L/k)$, which has the same restriction to D. Then $(r,b)^2=(\iota,1)$. In particular every real character of H' takes equal values on g_+ and g_-. Progression automorphism · Exceptional term in effective Chebotarev

The finite Hecke conductors for LC/k divide $M\mathcal{O}_k$. For L/k, odd ramification is tame at q and absent elsewhere. At the dyadic prime $\mathfrak{p}$ of k, the 2-adic logarithm identifies the fourth powers of $1+\mathfrak{p}^3$ with $1+\mathfrak{p}^7$, since $v_{\mathfrak{p}}(4)=4$. A quartic character is trivial there and $16\mathcal{O}_k=\mathfrak{p}^8$. For C/k, the characters extend to Dirichlet characters at modulus M and pull back by the norm. A character on H' extends to the product of the two abelian groups over k, giving the same conductor bound. Thus the maximum conductor norm is at most M^2, and k has discriminant 8 and degree 2. Abelian conductors and effective Chebotarev

Apply effective Chebotarev with this common abelian subgroup and the two singleton classes, at U_2. The size threshold holds since $\log U_2=z(\log Y)^2$ and $\log M=O(z)$. The main terms, including the possible exceptional term, agree by the square calculation. Each has relative error $O(\exp(-c(\log Y)^2))$ with $c>0$ independent of q. All ramified primes divide M, so the two counts add to the progression count. Removing primes through $\sqrt{U_2}$ changes either count by at most $\sqrt{U_2}$, which is $o(|P|)$ for the second-stage candidate set P. Thus each sign has relative share $1/2+o(1)$ uniformly in the first prime. Choose $\nu=1/8$, take any first-stage good prime and intersect its second-stage good set with the negative sign. The share is at least $3/8-o(1)>0$, proving the assertion. Effective Chebotarev, Theorem 1.4 · Good fractions

3. From prime selection to real cutoffs

Cohomological inputs (specified in Paper VI). For a prime ℓ write $Ш_\ell(E,j)$ for the kernel of localization in degree one with coefficients $\mathbb{Q}_\ell/\mathbb{Z}_\ell(-j)$, for the Galois group unramified outside places over ℓ, localizing at all those places. The source uses all algebraic closures and all choices of those places. The requirement `VVle` at t, B_* asserts that each of these lengths over $\mathbb{Z}_\ell$ is finite for $1\le j<t$, and that for every finite prime set S,

$$\sum_{\ell\in S}\log\ell\sum_{1\le j<t}\operatorname{length}_{\mathbb{Z}_\ell}Ш_\ell(E,j)\ \le B_*. \tag{2}$$

For odd ℓ the two eigenspaces for complex conjugation have signs $(-1)^j$ and $(-1)^{j+1}$. Their lengths add. Below, N_E denotes the conductor of E. Localization kernel · Uniform length bound · Eigenspaces and assembly

The release uses the following three bounds, with v the cutoff before rescaling and $t(v)=\lfloor v^{401}\rfloor$. It takes the sum of their budgets.

- The dyadic input `BaAt` gives the unweighted length sum at $\ell=2$ at most $\exp(v^{4\varepsilon/5})$.
- The odd reference input `BbAt` gives the same separate bound over each fixed reference layer. At conductor one these are layers over $\mathbb{Q}(\zeta_3)$ and $\mathbb{Q}(\zeta_7)$. With these bounds, `PointHV` gives a bound $(N_Et(v))^C+C\exp(v^{4\varepsilon/5})$ on the weighted sum for the odd eigenspace of sign $(-1)^j$, whenever $e\ge2$, $t(v)\ge2$, $q,d>2t(v)+1$, and both primes exceed the fixed exclusions for the level and reference layers.
- The pointwise input `ThmBcFree` bounds the other odd weighted sum by `minusBound` for every such field and positive rank range $e,t(v)\ge2$ with the prime congruences. In the size range $\log(N_Et(v))\le2v^{\varepsilon/4}$ its bound is at most $C'\exp(8v^{\varepsilon/4})$, with C' fixed.

The bounds include finiteness of each length they cover. At conductor one, Paper VI §3.4 proves the odd minus estimate $[E:\mathbb{Q}]t(v)^2\log(t(v)N_E)$ by the equivariant Coates–Sinnott theorem for every odd coefficient prime. The odd plus estimate still needs the surviving deformation-theoretic cases of Paper VI. The release aggregates the two reference conditions and the dyadic bound on a common interval via `ThmBab`, then applies `PointHV` there. Its `VWinAt` leaves the starting point of this common interval dependent on the given pair. Three interfaces · Odd pointwise hypotheses · Proof inputs for the joint window · Budget assembly · Paper VI

Arbitrarily late windows (rescaling proved here). Fix $0<\varepsilon<1/52$, put $\alpha=2+8\varepsilon$, and take $e(v)=\lceil\alpha\log_2 v\rceil$ for $v\ge 1$. Thus $v^\alpha\le 2^{e(v)}\le 2v^\alpha$, with $2^{e(v)}\le v^4$ for $v\ge 2$. The released `VWinAt` input at exponent 401 fixes $J\ge 1$ and X_0 before X and the pair are chosen. With $Y=X^{1/J}$, it supplies, for **every** eligible pair at scale Y, a W with $Y\le W$ and $W^2\le X$, and (2) on all $v\in[W,W^2]$ with $B_*=\exp(v^\varepsilon/2)$. The proof from the three component estimates uses $\log(qd)\ll Y^{\varepsilon/5}(1+\log^2 Y)=o(Y^{\varepsilon/4})$ and $\log(2^{e(v)+2}t(v))=O(\log v)$, to fit the three budgets below $\exp(v^\varepsilon/2)$. The lower bounds for q,d also give $q,d>2X^{401}+1$ and every fixed exclusion eventually. Window quantifiers · Conductor and total budget

Choose the pair by (1) with $A=J$ before taking W. In the ordinary case the scale factor in `winPairs` is 1, since $n=2^e$. Thus for $x\in[W,W^{3/2}]$ one has $Y\le x\le X$, $x^{1/J}\le Y$, $x^\alpha\le 2^{e(x)}\le 2x^\alpha$, and (2) with $t=\lfloor x^{400}\rfloor$ by monotonicity. These are the fields of `WinDatum`. The starts $W\ge X^{1/J}$ are unbounded. Paper VI gives the general rescaling. Released rescaling · Output data · Paper VI

A conditional proposition for prescribed starts at conductor one

Fix $0<\varepsilon<1/52$, $J\ge 2$, and the admissible primes q,d of Proposition 2 at scale $Y=X^{1/J}$ with $A=J$. Write $S_a=\mathbb{Q}(\zeta_{2^a},\sqrt{qd})$ and $L_a=\mathbb{Q}(\zeta_{2^a},\sqrt{-q},\sqrt{-d})$ for $a\ge 2$. Let X_S,X_L denote the norm limits of their ordinary 2-primary ideal class groups, and let γ_2 be the cyclotomic generator of value 5. Put $E_e=\mathbb{Q}(\zeta_{2^{e+2}})^+(\sqrt{-q},\sqrt{-d})$, so $L_{e+2}=E_e(i)$.

Proposition 3 (conditional). Assume the following inputs on the pairs so selected, with the pointwise bounds uniform in X.

1. The modules X_S,X_L are free over $\mathbb{Z}_2$ of ranks two and three. If D denotes the action of γ_2 on X_S, its characteristic polynomial on $X_L\otimes\mathbb{Q}_2$ is $(T-1)\det(TI-D)$. Only this polynomial identity is needed for the determinant calculation.
2. The ordinary 2-class groups of S_2,S_3 are the coinvariants of X_S by $D-I,D^2-I$, respectively. Their orders are the 2-parts of $h(-4qd)/2$ and $h(-4qd)h(-8qd)/4$, respectively. Here h is the class number of the indicated quadratic discriminant.
3. For $j\ge 1$, restriction injects the strict localization kernel from E_e into the kernel at L_{e+2}. The dual of the latter is a quotient of $X_L/((5^j\gamma_2)^{2^e}-1)X_L$ obtained by imposing splitting at the places above 2.
4. Descent gives the reference localization bound in the 3- and 7-power towers above. The pointwise odd plus and minus bounds, including finiteness, hold for the chosen host with the budgets following (2).

Then $\operatorname{length}_{\mathbb{Z}_2} Ш_2(E_e, j) \le 3e + 4 + v_2(j)$. Every sufficiently large prescribed Z admits one pair q, d giving the arithmetic data and (2) at every real $x \in [Z, Z^{3/2}]$, with twist limit $\lfloor x^{400} \rfloor$ and budget $\exp(x^{\varepsilon}/2)$. The geometric construction at those cutoffs enters in §4.

Chems-Eddin and Müller give the minus λ-invariant and oddness of the real-layer class numbers used in the tower calculation. Paper VI §3.3 gives the additional integral rank, class-number, and restriction arguments for inputs 1–3 and treats the odd reference fields. Chems-Eddin–Müller, Theorems 4–5 · Paper VI · Strict dyadic descent · Released layer comparison

Proof. Input 2 gives $v_2 \det(D - I) = v_2 h(-4qd) - 1$. Subtracting this valuation from $v_2 \det(D^2 - I)$ gives $v_2 \det(D + I) = v_2 h(-8qd) - 1$. With the selected sign, both determinant valuations are 2. If $p = \det(D - 1)$ and $s = \operatorname{tr}(D - 1)$, then $4 + 2s + p = \det(D + 1)$ implies $s \equiv 2 \pmod 4$. Thus the characteristic polynomial of $(D - 1)/2$ has coefficients in $\mathbb{Z}_2$ and reduces to $T^2 + T + 1$. Each eigenvalue of D is $1 + 2z_0$ with $\overline{z_0} \in \mathbb{F}_4 \setminus \mathbb{F}_2$. For all positive integers j, both $5^j(1 + 2z_0) - 1$ and $5^j(1 + 2z_0) + 1$ have valuation one. The first squaring gives valuation two for the difference from 1, and each further squaring adds one. At layer $m = e$ the two eigenvalues therefore contribute $2(e + 1)$ to the determinant valuation. The trivial factor contributes $v_2(5^{j2^e} - 1) = e + 2 + v_2(j)$. The resulting determinant has finite valuation, which is the length of the cokernel on the free lattice X_L. Input 3 now gives the asserted upper bound on the host. Summing gives

$$(t - 1)(3e + 4) + v_2((t - 1)!) = O(t \log v),$$

uniformly for $t = t(v)$, since $e(v) = O(\log v)$. This fits below $\exp(v^{4\varepsilon/5})$ for every sufficiently large v. Twist and layer bounds · Dyadic budget

Paper VI proves that the odd reference towers over $\mathbb{Q}(\zeta_3)$ and $\mathbb{Q}(\zeta_7)$ have zero primary class groups, using Minkowski's bound and the ambiguous class number formula along the tower. The reference localization lengths are therefore zero by descent. The two pointwise odd host estimates still enter (2). Paper VI · Reference fields · Descent comparison

Fix $J \ge 2$ and apply Proposition 2 with $A = J$ and $Y = X^{1/J}$. For every sufficiently large X, the dyadic and reference estimates now hold throughout $[\sqrt{X}, X]$. The pointwise odd estimates and the size checks above give (2) there at exponent 401. For a prescribed sufficiently large Z take $X = Z^2$ and $W = \sqrt{X} = Z$. The resulting pair gives `WinDatum` on $[Z, Z^{3/2}]$, including $x = Z$. Pointwise height application · Cohomological window · Prior cutoff proof

4. Fitting the geometry at the same cutoff

Arithmetic data for the host

Fix the permitted constants and take $0 < \varepsilon < 1/52$ below the common existence cap. At each sufficiently large cutoff supplied by §3, use the field and frame of §1 with that cutoff's q, d, e, x. The rank input supplies m lifts, k columns, and twist limit $R.t$, with

$$m \in [x^{1/2 - \varepsilon/16},\ 2x^{1/2 - \varepsilon/16}], \qquad k \in [x^{9/2 + 10\varepsilon},\ 2x^{9/2 + 10\varepsilon}], \qquad R.t \le \lfloor x^{400} \rfloor$$

eventually. The size identity for the remaining ranks is the last equality of `RankVal`. The host extension needs $q, d > 4R.t\, R.r + 1$ and the small-prime splitting rows through $\lfloor Y^{\varepsilon/5} \rfloor$. The first follows from the

rank size bounds and the lower bounds on q, d when $Y \ge x^{1/J}$. The second follows from the progression in §2. Restrict (2) to the actual $R.t$ to obtain the cohomology input with bound $\sqrt{b}$ for $b = \exp(x^{\varepsilon})$. Paper VI §4 derives the norm primes in short intervals from the effective Chebotarev theorem over K, chooses the distinct level primes, and takes the unit ideal for the prescribed ideal of K. The character estimates for the nontrivial quadratic twists use the pair bounds of §2. With these arithmetic data fixed, the following constructions supply the two sets on this frame. Rank conditions · Twist limit · Admissibility · Host extension · Paper VI, §4

Thick digits and translated orbits

Use the coordinates $a \mapsto [5^a]$ of §1 on $H \simeq \mathbb{Z}/n$, where $n = n_0 = 2^e$ and $H_0 = H$. Write S_2 for the set of nonnegative integers stored as the host's thick digits and $T_2 \subset H$ for their image. The required length is

$$L_w = 10\lceil 405 \log x \rceil + 2\lceil 30/\varepsilon \rceil + 22.$$

Each $s \in S_2$ must satisfy $4L_w s < n$, so the digit map is injective. The condition $\mathrm{NoRel}(S_2, L_w)$ excludes

$$j(s_2 - s_1) = i(s_3 - s_1)$$

for pairwise distinct $s_1, s_2, s_3 \in S_2$ and distinct integers i, j with $1 \le |i|, |j| \le L_w$. It excludes three-term progressions and triples with unequal gaps satisfying the same coefficient bound. The source requires an absolute $C_{T_2} > 0$ with

$$\frac{n}{L_w} \exp\bigl(- C_{T_2} \sqrt{\log n \, \log L_w} \bigr) \le |T_2|, \qquad x^{-\varepsilon} n \le |T_2| \le \frac{n}{4L_w}.$$

The raw digit interval permits 0 and excludes $n/(4L_w)$. Host data · Bounded relations and orbit conclusions · Length · Required sizes

Construction. For integers $n \ge 1$ and $L \ge 2$, put $M = \lfloor n/(4L) \rfloor$ and $\lambda = \log(8L)$. If $M = 0$ or $M < \exp(4\lambda)$, take $S_2 = \{0\}$. Here $\log n \le 5\lambda \le 20 \log L$, using $M \ge n/(8L)$ for $M \ge 1$. Hence $\log(n/L) \le 15\sqrt{\log n \, \log L}$.

In the other case put $A = \log M$, $D = \lfloor \sqrt{A/\lambda} \rfloor$, $m_0 = \lfloor M^{1/D}/(4L) \rfloor$, and $Q = 4Lm_0$. Then $D \ge 2$, $M^{1/D} \ge 8L$, and $M^{1/D}/(8L) \le m_0 \le M^{1/D}$. Encode $v \in \{0, \dots, m_0 - 1\}^D$ by $\sum_{a=0}^{D-1} v_a Q^a$. Take the encodings of all vectors of a most frequent squared length. The encodings are distinct and lie in $[0, Q^D) \subseteq [0, M)$. There are at most $D(m_0 - 1)^2 + 1 \le D m_0^2$ squared lengths, so this gives at least $m_0^D/(Dm_0^2)$ digits.

For the forbidden relation with coefficient bound L, each coefficient at a power of Q has absolute value at most $(|i| + |j| + |i - j|)(m_0 - 1) < Q$. Reducing modulo Q forces the first coefficient to vanish. Dividing by Q and repeating gives the vector relation $j(v_2 - v_1) = i(v_3 - v_1)$. The three vectors would be distinct collinear points on the same sphere. Squared distance to the center along a line is strictly quadratic, so there are at most two such points.

Since $\sqrt{A/\lambda}/2 \le D \le \sqrt{A/\lambda}$, the cardinality satisfies

$$\log |S_2| \ge A - D\lambda - 2A/D - \log D \ge A - 6\sqrt{A\lambda}.$$

Using $M \ge n/(8L)$, $\lambda \le 4 \log L$, and $\log 8 \le 3\sqrt{\log n \, \log L}$ gives the required lower bound with $C_{T_2} = 15$. Both constructions have $S_2 \subseteq [0, M)$ when $M \ge 1$, so $4L|S_2| \le n$ in that range. At $L = L_w$,

we have $2 \le L_w = O_\varepsilon(\log x)$ and $n \in [x^{2+8\varepsilon}, 2x^{2+8\varepsilon}]$. Thus $\log(n/|S_2|) \le \log L_w + 15\sqrt{\log n \, \log L_w} = o(\log x)$ uniformly. This gives $\delta_2 = |T_2|/n \ge x^{-\varepsilon}$ eventually. Sphere construction · Density used in the comparison

For $t, u, h \in H$ and $\operatorname{ord}(h) > 4L_w$, every half-open integer window $[a_0, a_0 + L_w)$ contains at most two indices a with $t + ah \in u + T_2$. Three hits $a < b < c$ would give distinct digits and

$$(c-a)(s_b - s_a) - (b-a)(s_c - s_a) \equiv 0 \pmod{n}.$$

The absolute value on the left is $< n$ by the digit range, contradicting NoRel after lifting to $\mathbb{Z}$. The argument includes wraparound in H. If $d_h = \operatorname{ord}(h) \le 4L_w$, the whole orbit meets $u + T_2$ in at most one distinct point. Two hits give $d_h(s_1 - s_2) \equiv 0 \pmod{n}$, and the digit range forces $s_1 = s_2$. This point can recur in the index sequence. Also, for every $v \in H$ there are exactly $|T_2|$ translates u with $v \in u + T_2$, using $u = v - s$ for $s \in T_2$. Translated orbit bounds · Translate mean

Tops and weighted Frobenius-orbit estimates

Put $\delta = x^{-\varepsilon/8}$, $\alpha = 2 + 8\varepsilon$, and fix $C_0 \ge J$. We need $\delta n/2 \le |TT| \le 2\delta n$ and the weighted bad-orbit estimate below on this same set. The source's coset-share condition measures the discrepancy from proportions $4/(3o+1)$ on the identity coset and $3/(3o+1)$ on other cosets, with tolerance $x^{-1+\varepsilon}|TT|$. At conductor one the only coset is H, so the discrepancy is zero for every translate. The orbit tests for TT allow three-term progressions. Top selection · Coset shares

Set $X' = \exp(x^{\varepsilon/3})$ and let $\mathcal{P}$ be the primes $p < X'$ in the host's residue-zero class $\mathrm{inertClass}$ for which $d_p = \operatorname{ord}(h_p) \ge x^{\varepsilon/4}$. Here $h_p \in H$ is the cyclotomic coordinate of arithmetic Frobenius. The prime weights are $\omega_p = (\log p)/p$. By the small-prime splitting in §2, every inert-class prime exceeds $Y^{\varepsilon/5}$. Consequently the short window length satisfies

$$L_2 = 2 \sup_{p \,\mathrm{inert Class}} \left\lceil \log_p((kx^{26})^4) \right\rceil \le L_b = 2\left\lceil \frac{620 C_0}{\varepsilon} \right\rceil.$$

The supremum is over all inert-class primes. The rank bound gives $4\log k + 104 \log x \le 124 \log x$ eventually, and $Y \ge x^{1/J}$ gives $\log p > (\varepsilon/(5J)) \log x$. Prime class and weighted filter · Short window length

For each orbit $t + \langle h_p \rangle$, use its fixed representative in the source and number the positions $0, \dots, d_p - 1$ in steps of h_p. Use the parameters

$$K_1 = \left\lceil \frac{8(\alpha+2)}{\varepsilon} \right\rceil, \qquad r = K_1 + 1, \qquad U = \lceil 1/\delta \rceil, \qquad W_T = (6r+1)U,$$
$$M_0 = \left\lceil (\alpha+3) 10^{6r+1} \log x \right\rceil, \qquad G' = (M_0 + 5) W_T,$$
$$N_1 = 1 + \left\lceil \max\left(6\exp(1) G' \delta, (\alpha+3)\log_2 x\right) \right\rceil.$$

The source's good-orbit predicate is the conjunction of the following three tests on the indices selected by TT. Parameters · Short-window test · Tile and interval tests

1. Every cyclic window of L_2 indices contains at most K_1 selected indices.
2. Use the $\lfloor d_p / W_T \rfloor$ complete tiles of W_T indices starting from 0. Each tile has $6r+1$ cells of length U, numbered starting at 1. A tile is marked if every odd cell is empty and every even cell contains

exactly one selected index of even parity in the orbit numbering. For every tile index s_0, one of the tiles $s_0 + j$, $0 \le j \le M_0$, is marked, with addition modulo the number of complete tiles.

3. Every cyclic window of G' indices contains at most N_1 selected indices.

Simultaneous selection. Select each point of H independently with probability δ. For fixed $p \in \mathcal{P}$ and t, condition on selecting t. The other indices in its orbit still have independent Bernoulli laws. Uniformly for the orbits tested, G' and $(M_0+1)W_T$ are $O_\varepsilon(x^{\varepsilon/8}\log x)$ and less than d_p eventually. Also $L_2 \le L_b < d_p$. Thus individual windows have distinct indices, and the trials in test 2 use distinct tiles.

For test 1, a window containing the conditioned point requires K_1 further selections to fail, and at most L_2 windows contain that point. Each other failing window requires $K_1 + 1$ selections. The union bound is at most

$$L_2\binom{L_2}{K_1}\delta^{K_1} + d_p\binom{L_2}{K_1+1}\delta^{K_1+1} \le (L_b+2)2^{L_b}x^{-2},$$

as $d_p \le 2x^\alpha$ and $\delta^{K_1} \le x^{-(\alpha+2)}$.

For test 2, discard the trial tile containing the conditioned point if there is one. The other at least M_0 tiles are independent. An empty cell has probability $(1-\delta)^U$. A cell with exactly one selected index, of even parity, has probability at least $\lfloor U/2 \rfloor \delta(1-\delta)^{U-1}$. For $\delta \le 1/100$, the first exceeds $1/3$, by $\log(1-\delta) \ge -\delta/(1-\delta)$ and $U\delta \le 1+\delta$. The second exceeds $1/10$, since $\lfloor U/2 \rfloor \delta \ge (1-\delta)/2$. The marking probability is at least $10^{-(6r+1)}$. Summing over the at most d_p starting tiles, failure costs at most $d_p \exp(-M_0 10^{-(6r+1)}) \le 2x^{-3}$.

For test 3, the conditioned count in a window is dominated by $1 + \mathrm{Bin}(G', \delta)$. With $N_* = N_1 - 1 \ge 1$, the binomial tail at N_* is bounded by

$$\sum_{\nu=N_*}^{G'} \left(\frac{\exp(1)G'\delta}{\nu}\right)^\nu \le \frac{6}{5}\,6^{-N_*} \le \exp(-N_*).$$

This uses $N_* \ge 6\exp(1)G'\delta$. Multiplying by the at most d_p windows gives at most $2x^{-3}$, since $N_* \ge (\alpha+3)\log_2 x$.

Write $B_p(TT)$ for the number of selected tops on bad orbits at p. The conditional failure probability for each p, t is at most $C_b x^{-2}$ for $C_b = (L_b+2)2^{L_b} + 4$, independent of x. Hence

$$\mathbb{E}\sum_{p\in\mathcal{P}} \omega_p B_p(TT) \le C_b x^{-2}\delta n \sum_{p\in\mathcal{P}} \omega_p.$$

By Markov the chance of exceeding four times this bound is at most $1/4$, with zero weighted count almost surely if the bound is zero. The size has mean δn and variance at most δn. Chebyshev gives probability at most $4/(\delta n)$ of leaving $[\delta n/2, 2\delta n]$. Since $\delta n \to \infty$, one choice satisfies the size interval, the exact ordinary coset shares, and

$$\sum_{p\in\mathcal{P}} \omega_p B_p(TT) \le 4C_b x^{-2}\delta n \sum_{p\in\mathcal{P}} \omega_p.$$

This is the source's BadOK estimate. To obtain the endpoint estimate's bound $|TT|x^{-2+\varepsilon}$, use $\sum_{p\in\mathcal{P}} \omega_p \le 2\log X' = 2x^{\varepsilon/3}$ and $|TT| \ge \delta n/2$. Indeed, with $M_X = \lfloor X' \rfloor$, the prime contributions

$\lfloor M_X/p \rfloor \log p$ to $\log(M_X!)$ give this prime-weight bound using $\lfloor M_X/p \rfloor \geq M_X/(2p)$. Group selection · Weighted condition after transport · Endpoint allowance

The marked tiles supply the endpoints on a good orbit. Each contains $3r$ even selected indices, at least U from the left endpoint, with $i + U$ strictly before the right endpoint and successive gaps $> U > L_2$ for large x. For successive marked tiles, write their increasing lists as f, f', indexed from 0, and use $f(r)$ and $f'(2r-1)$. Lift the latter by d_p on wraparound. Apply test 2 starting at the tile immediately after the first marked tile. The next marked tile is at most $M_0 + 1$ steps away, and there are more than $M_0 + 1$ complete tiles because $G' < d_p$. The final incomplete tile, of length $< W_T$, can occur once in the segment. Thus the forward endpoint distance is $< (M_0 + 3)W_T < G'$, and the inclusive segment contains at most N_1 selected indices. Marked positions · Segment endpoints

Carrying the sets into the geometry

The host installs S_2 at offset $u_0 = 0$ and then chooses TT using the prime and orbit data. Its top places are $((0,0), t)$ for $t \in TT$. At translate u, the thick places are $((0,1), s+u)$ for $s \in T_2$. Their distinct quadratic colors make the two supports disjoint for every u, even when their projections in H intersect. The stated orbit and density bounds apply to the same T_2 at all translates. In §5 the thick partner coordinate $s + u - t$ covers H once as u varies with s, t fixed. Paper VI transports these choices into the kit and family, and uses the polynomial prime bound for the extra level prime q_0. Its weight existence, tame distinctness and common word-code inputs concern the fitted objects. Host places · Initial offset · Installation · Family design · Paper VI, §4

The local theta estimates use the thick-orbit test on the letters $\tau\phi^i$. If the cyclotomic step has order $> 4L_w$, the thick letters satisfy the at-most-two condition in every window of L_w indices. At depth j with $1 \leq j$ and $2j + 2 \leq L_w$, three thick letters linked by successive gaps at most j would lie in one such window. The twisted dual page-word estimate inserts the resulting cluster bound and $j + 1 \leq \mathrm{ord}(\phi)$ into its local hypothesis $\mathrm{HypA}(j)$, along with the numerical weight bounds. Thick-letter window · Cluster and word length · Local theta use

The remaining geometric input, from Papers IV and VI and the companion construction, supplies the fitted objects and period and norm estimates on this host. In particular the comparison uses positive heights a_t for $t \in TT$ and positive weights $w_{u,j}(\rho)$ at $\rho \in \Phi$ for every $u \in H$ and $1 \leq j \leq m$. Write $n_T = |TT|$, $n_2 = |T_2|$, $a_E = xb$, $b' = a_E/X'$, $N = \lceil 2a_E \rceil$, and $\theta = \exp(-c_6 x^{\varepsilon/8})$ with fixed $c_6 > 0$. With $\eta \geq 0$ the fitted weight tolerance, the conditions are

$$\begin{gathered} |a_t/b - x| \leq 1, \qquad a_t/w_{u,j}(\rho) \leq N, \qquad b \leq w_{u,j}(\rho) \leq b + \eta + 1 \text{ at thick partners}, \\ w_{u,j}(\rho) \leq 2b' \text{ at thin partners}, \qquad 2f(x+1)\theta \leq 1/x, \qquad 6f(x+1)(\eta+1) \leq b/x, \\ n_T \geq n x^{-\varepsilon/8}/2, \qquad \delta_2 \geq x^{-\varepsilon}. \end{gathered} \tag{3}$$

The first cost bound follows eventually because $f = 2n = O(x^{2+8\varepsilon})$. The second uses the fitted η. The last line is supplied by the two constructions above, and the fitting and period assertions apply to the same field, sets, lifts, weights and truncation. Scales · Digit density · Costs and ranks · Fitting to the chain · Paper IV

Thin estimate (Paper I, Proposition 4). At all $y \geq X'/3$, for every primitive character χ of the fitted host, including the principal case interpreted as G_z,

$$|G_\chi(y)| \le y\theta. \tag{4}$$

The conductor input in the release is $\log(Bqd) \le 2x^{\varepsilon/5}(\log x)^2$, which follows eventually from $Y \le x$, the bounds on q, d, and $B = O(x^{2+8\varepsilon})$. To verify the real-character input, a real character of the dyadic units is trivial on squares, hence on all units congruent to $1 \bmod 8$. Its primitive conductor divides 8. Multiplying by J_q, J_d gives all real host characters by the Chinese remainder theorem. For a real character ramified at q or d, the box of §2 gives real nonvanishing above $399/400$. The finitely many remaining nonprincipal real characters have a common interval below 1 free of zeros by classical nonvanishing at 1 and continuity. Together with the conductor bound, this common interval verifies the hypotheses of Proposition 4 of Paper I and proves (4). Uniform contour proof Thin input · Host bound and scales · Use on the fitted host

Preserving the domain (proved here from the stated pointwise interfaces). Fix ε, J before fitting and let $\mathcal{W} \subset \mathbb{R}$ be the set of all x admitting `WinDatum` with these constants. Host supply uses exactly $\mathcal{W}$ and requires $x \ge x_0$ for a fixed threshold. The enlargement of the level-prime set, object construction and weight construction use this same scale set and can increase the threshold. The family built from kits has $\mathrm{Win} = \mathcal{W}$ and uses the actual scale x above its threshold for every $x \in \mathcal{W}$. The later closed conditions are requested at the joined parameter bound before the family is chosen. Their conclusion gives one set of constants and a further threshold on that same family. Thus the geometry can be invoked at every sufficiently large $x \in \mathcal{W}$ once these pointwise existence inputs are supplied. On the universal window route, $\mathcal{W}$ contains arbitrarily late power windows. With the refined choice in §3, $\mathcal{W}$ contains an entire tail. The latter property is carried by the concrete family construction through its equality of scale sets. Choice of scale set · Releveling and weights · Weight construction · Kit indexing · Joined conditions · Same-point identity

5. The finite prime-power identity and the thick mean

Fix a fitted cutoff with $x \ge 3$. For any character $\vartheta \in \widehat{G}$ let G_ϑ and ℓ_ϑ use its primitive representative, treating the principal ramp as G_z. For $\rho \in \Phi$ take a positive weight $w(\rho)$, and suppose $a > 0$ and $a/w(\rho) \le N$ for all $\rho \in \Phi$. For a top $t \in H$ write $t_* = ((0,0),t)$ and $P_t(g) = c + g + t_*$, a bijection from $G \setminus \Phi$ to Φ. The finite prime sum ASK uses the factor $(a/p^i - w(P_t(g)))_+$ at a class $g \notin \Phi$. For $p \nmid Bqd$ it takes $g = i\,\mathrm{Frob}_p$ with coefficient $2\log p$. For $p = d$ it sums over the complementary classes in $i\,\mathrm{Frob}_d + I_d$ with coefficient $\log d$. At $2, q$ its contribution vanishes. Only $p^i \le N$ can contribute. Define

$$\mathrm{MAIN}(a,w) = a\left(\log a - 1 - \gamma - \frac{1}{f}\sum_{\rho\in\Phi}\log w(\rho) + \ell_{J_d}\right).$$

The finite balance is

$$\mathrm{ASK} - \mathrm{MAIN} = \frac{1}{f}\sum_{g\notin\Phi}\sum_{\vartheta\in\widehat{G}} \vartheta(-g)\, w(P_t(g)) G_\vartheta\big(a/w(P_t(g))\big). \tag{5}$$

Finite definitions and truncation · Released balance

Proof of (5). First use the uncentered prime ramps. Fourier inversion gives, at p^i, the coefficient

$$\frac{1}{f}\sum_{\vartheta\in\widehat{G}}\vartheta(-g)\frac{1}{|I_p|}\sum_{h\in I_p}\vartheta(i\,\mathrm{Frob}_p+h)=\frac{2}{|I_p|}\,\mathbf{1}_{g\in i\,\mathrm{Frob}_p+I_p},$$

since $|G|=2f$. The inertia and splitting calculations of §1 now give all the displayed prime coefficients. The principal center subtracted from the uncentered sum is $a\log(a/w)-(1+\gamma)a$, averaged over the f partners. For nonprincipal characters, the center added to the uncentered sum for a given g is $a\ell_\vartheta\vartheta(-g)/f$, independent of the partner's weight. The sum of $\vartheta(-g)$ on $G\setminus\Phi$ is zero unless ϑ is trivial on Φ. The unique nonprincipal character trivial there is J_d, with value -1 on the complementary coset, so its sum is $-f$. This accounts for $+a\ell_{J_d}$ in MAIN upon subtraction. Balance proof and normalization

For each u,t,j use $a=a_t,\ w=w_{u,j}$ in (5). The thick partners are $((0,1),s+u)$ for $s\in T_2$. Thus their classes g in (5) are $g_s(u,t)=((1,0),s+u-t)$, each occurring once. Put $y_t=a_t/b$, and let Th_{ut} denote the right side of (5) with thin partners omitted and all thick weights replaced by b. Summing over $u\in H$ kills the characters nontrivial on H. For each surviving character ϑ, the phase in every summand is its value at $((1,0),0)$, independent of s,u,t. In order $\mathbf{1},J_q,J_d,J_qJ_d$ the four signs are $+,+,-,-$. Consequently (**proved here**)

$$\frac{1}{n}\sum_{u\in H}\mathrm{Th}_{ut}=\frac{bn_2}{f}\,G_{\mathrm{sig}}(y_t),\qquad G_{\mathrm{sig}}=G_z+G_{J_q}-G_{J_d}-G_{J_qJ_d}.\tag{6}$$

The four terms are real. The factor $n_2/f=\delta_2/2$ persists after taking the mean. Thick partner bijection · Fourier averaging

The comparison error

We first bound the change at a fixed u,t,j, with $a=a_t$. For $w\ge b$, the function $v\mapsto(a/n-v)_+$ is decreasing and 1-Lipschitz. It has zero contribution for $n>a/b$ at either weight. For a nonprincipal character the two centers $a\ell_\vartheta$ cancel in the weighted difference. Thus by the classical Chebyshev estimate $\sum_{n\le y}\Lambda(n)\le 3y$ the change has modulus at most $3(a/b)(w-b)$. In the principal case the uncentered difference from b to w, taken as its value at b minus its value at w, belongs to $[0,3(a/b)(w-b)]$. From it one subtracts $a\log(w/b)$, which lies in the same interval. This proves the same bound in the principal case. At a thin partner, (3) gives $a/w\ge X'/3$ because $a\ge(x-1)b$, $w\le 2b'$, and $x\ge 3$. The omitted term has modulus at most $a\theta$ by (4). Weight estimates · Thin and thick substitutions

There are n_2 thick and $f-n_2$ thin partners and $2f$ characters in each term of (5). Thus the error at u,t,j has modulus at most

$$2\left[n_2\cdot 3y_t(\eta+1)+(f-n_2)a_t\theta\right]\le b/x$$

by (3). The factor 2 comes from the full character sum and the denominator f. Let $E(u)=\sum_{t\in TT}\sum_{j=1}^{m}(\mathrm{ASK}-\Re\mathrm{MAIN})$ at translate u, and let $\overline{E}$ denote its uniform mean. Put $Q=n_T m b\,\delta_2/2>0$. Summing the error and using (6) yields

$$\left|\frac{\overline{E}}{Q}-\frac{1}{n_T}\sum_{t\in TT}G_{\mathrm{sig}}(y_t)\right|\le\frac{2}{x\delta_2}.\tag{7}$$

The release also uses the weaker allowance $4/(x\delta_2)$ by separately budgeting both errors using all partners. Stored cost comparison · Same-host excess and normalization

For the move to the common cutoff, the classical Mertens estimate $\sum_{n\le y}\Lambda(n)/n=\log y+O(1)$ implies $G_z'(y)=O(1)$ away from integers, since the derivative of its center cancels $\log y$. Each nonprincipal ramp has derivative of modulus at most $\log y+O(1)+|\ell_\vartheta|$. The ramps are continuous and piecewise differentiable, including at the newly admitted integer where the new ramp term vanishes. Since $|y_t-x|\le 1$, (1) at x gives, uniformly in the host,

$$|G_{\mathrm{sig}}(y_t)-G_{\mathrm{sig}}(x)|\le C(1+\log x)|y_t-x|.$$

Together with (1) and (7) this proves the ordinary comparison (**proved here**)

$$\boxed{\left|\frac{\overline{E}}{Q}-G_z(x)\right|\le 3x^{1/2+\varepsilon/2}+C(1+\log x)+\frac{2}{x\delta_2}.} \tag{8}$$

The sharper signed comparison uses (7) and the cutoff move alone. The bounds at x suffice even when a particular y_t lies outside the selected interval of analytic bounds. Derivative comparison · Three bounded characters · Principal term

The geometric upper bound needed at the interface

Supplied by Paper III and the geometric papers. On the fitted host, for fixed geometric constants and sufficiently large eligible x, the same excess satisfies

$$E(u)\le x^{(e_D+e_L+3)\varepsilon}\mathcal{B}\quad\text{for every } u\in H,\qquad \mathcal{B}=n_Ta_E(1+m^2/k)+m\left(2nmbx^{\varepsilon}+\frac{n_2kbx^{2\varepsilon}}{n^2}\right) \tag{9}$$

where $e_D,e_L\ge 0$ are the allowances for the finite divisibility and norm losses. Here the geometric estimates control the finite divisor gain and its gap from the main term. A bound on $\overline{E}$ with the same right side also suffices for (8). Error bracket · Host-specific chain bound · Paper III

To check the size explicitly, divide (9) by Q and use (3). The four terms in $\mathcal{B}/Q$, in order, are

$$O\left(\frac{x^{1/2+\varepsilon/16}}{\delta_2}+\frac{x^{-3-161\varepsilon/16}}{\delta_2}+\frac{x^{1/2+17\varepsilon/16}}{\delta_2}+x^{1/2-31\varepsilon/8}\right).$$

The last term uses $n_2/\delta_2=n$ before the bound $n/n_T\le 2x^{\varepsilon/8}$. With $\delta_2^{-1}\le x^{\varepsilon}$, the largest exponent after the factor in (9) is $1/2+(e_D+e_L+81/16)\varepsilon$. In particular (8) gives the release's looser bound $G_z(x)\le x^{1/2+(e_D+e_L+7)\varepsilon}$ eventually, with a fixed $\varepsilon>0$ absorbing the constants. It holds on the supplied late windows, or at every sufficiently large real x with the refined selection and the pointwise geometric inputs above. The same x,q,d,e,TT,T_2 and weights occur in (3)–(9). Geometry and arithmetic at the same scale · Released scalar conclusion · Rank calculation