MICROSCOPIC SHRINKING TARGETS FOR CUBIC BIRCH PATHS AND LOCAL KLOOSTERMAN CONTRASTS

SCOTT SCHUBR

Abstract. We derive local probability laws for cubic Birch paths and compare them with Kloosterman paths. Write $H = \lfloor p^{1/3} \rfloor$ and let M_p be the maximum of the first H Birch partial sums divided by $\sqrt{H}$, with the linear coefficient uniform modulo p. For every $r > 0$, we obtain $\sqrt{H}\,\mathbb{P}(M_p > r) \to 2/(\pi r)$. The corresponding two-parameter Kloosterman endpoint exceeds $\sqrt{H/2}$ before normalization with probability at least $1/8 - o(1)$. At the cubic scaling $\Delta = p^{-1/3}$, we count paths within sup-norm radius ε of the oscillatory integral $F_{\lambda_0}(t) = \int_0^t \exp(2\pi i(s^3 + \lambda_0 s))\, ds$. For $\Delta = o(\varepsilon)$ and $\varepsilon \to 0$, the probability is asymptotic to $2\varepsilon\Delta/\|\partial_\lambda F_{\lambda_0}\|_\infty$. At radius $r\Delta$ it has a positive threshold and an explicit interval-length law above threshold. For $\lambda_0 = 0$, an interval certificate and the second Euler–Maclaurin term give the first-hit distance $r_*\Delta + \kappa\Delta^2 + o(\Delta^2)$ with $\kappa > 0$. Finally, with reciprocal coefficient fixed at one, we construct both collapse and noncollapse on prime subsequences for arbitrarily slowly growing initial windows.

1. Introduction and notation

Set $e(u) = \exp(2\pi i u)$ and $e_p(u) = e(u/p)$. For prime p, an integer $H < p$, and $A, B \in \mathbb{F}_p$, write

$$U_{p,A}(j) = \sum_{x=1}^{j} e_p(x^3 + Ax), \qquad V_{p,A,B}(j) = \sum_{x=1}^{j} e_p(Ax + B/x), \qquad 0 \le j \le H,$$

where $1/x$ denotes the inverse in $\mathbb{F}_p$. A coefficient said to be uniform modulo p is sampled from all of $\mathbb{F}_p$. Subscript c denotes the centered representative of a residue, in $[-(p-1)/2, (p-1)/2]$ for odd p.

Kowalski and Sawin proved that the one-parameter Birch and two-parameter Kloosterman polygonal paths have the same full-path weak limit [2, Theorems 1.3 and 1.5]. They observed differences at shorter scales and asked for a probabilistic formulation of that distinction [2, Remark 1.4]. Lamzouri established uniform large-value bounds for the maximum over the full cubic path [3]. The results below describe the beginning of the path, where the cubic phase has a controlled variation on $p^{1/3}$ terms. The shrinking-target formulation follows the finite-prime path-counting viewpoint in [4].

Date: November 24, 2026.

2020 *Mathematics Subject Classification.* Primary 11L05, Secondary 60F05, 41A60.

Key words and phrases. Birch sums, Kloosterman paths, shrinking targets, cubic oscillatory integrals, Euler–Maclaurin formula.

2. A local separation at square-root amplitude

Theorem 2.1. *Let $H = \lfloor p^{1/3} \rfloor$ and $M_p(A) = H^{-1/2} \max_{1 \le j \le H} |U_{p,A}(j)|$. For every fixed $r > 0$,*

$$\lim_{p \to \infty} \sqrt{H}\, \mathbb{P}_A\big(M_p(A) > r\big) = \frac{2}{\pi r}. \tag{1}$$

If A and B are independent and uniform on $\mathbb{F}_p^\times$, then

$$\mathbb{P}_{A,B}\left(|V_{p,A,B}(H)| \ge \sqrt{H/2}\right) \ge \frac{1}{8} - O(p^{-1}). \tag{2}$$

Proof. Put $f_x = e_p(x^3)$ and $z = e(A_{\mathrm{c}}/p)$. The sequence f_x has bounded total variation on $[1, H]$, since its smooth extension has derivative of magnitude $6\pi x^2/p$. Abel summation and the geometric sum bound give, for $A \neq 0$,

$$\max_{j \le H} |U_{p,A}(j)| \le \frac{C}{|1 - z|}. \tag{3}$$

On the same interval the first derivative of the extension is $O(H^{-1})$ and the second derivative is $O(H^{-2})$. Consequently, the first differences of f_x have maximum and total variation $O(H^{-1})$. Multiplication by $1 - z$ followed by Abel summation on the first differences yields, uniformly in $j \le H$,

$$U_{p,A}(j) = \frac{f_1 z - f_j z^{j+1}}{1 - z} + O\left(\frac{1}{H|1 - z|^2}\right). \tag{4}$$

For $j = 1$ the first-difference sum is empty and the same identity holds.

Set $u = \sqrt{H}\, A_{\mathrm{c}}/p$. On compact u-intervals disjoint from zero, $\sqrt{H}|1 - z| \to 2\pi|u|$ uniformly in the sense that their difference tends to zero. The numerator in (4) has modulus at most 2. For the lower bound choose j to be the nearest integer to $\sqrt{H}/(2|u|)$. Then $f_j \to 1$ and $z^j \to -1$, both uniformly on these compact intervals. Consequently

$$M_p(A) = \frac{1}{\pi|u|} + o(1)$$

uniformly there. Since $|1 - z| \gg |A_{\mathrm{c}}|/p$, (3) confines $M_p(A) > r$ to $|u| \le C'/r$, including the residue $A = 0$. The lattice for u has spacing $\sqrt{H}/p$. Remove a fixed neighborhood of zero and neighborhoods of the boundary points $u = \pm(\pi r)^{-1}$, count the remaining lattice points by the uniform approximation, and then let the removed lengths tend to zero. This gives (1).

For the second assertion first average A, B independently over $\mathbb{F}_p$. Orthogonality gives $\mathbb{E}|V_{p,A,B}(H)|^2 = H$. In the fourth moment the surviving ordered quadruples satisfy

$$x + y = u + v, \qquad x^{-1} + y^{-1} = u^{-1} + v^{-1}$$

in $\mathbb{F}_p$. Since $2H < p$ for sufficiently large p, their common sum is nonzero, and the equations imply $xy = uv$. Thus $\{x, y\} = \{u, v\}$ with multiplicity, and the fourth moment is $2H^2 - H$. The Paley–Zygmund inequality applied to $|V_{p,A,B}(H)|^2$ gives probability at least $H^2/[4(2H^2 - H)] \ge 1/8$ for the stated event. Restricting to $A, B \neq 0$ changes its probability by $O(p^{-1})$. □

3. Shrinking targets at cubic amplitude

Let $\Delta = p^{-1/3}$. Interpolate $P_{p,A}(j\Delta) = \Delta U_{p,A}(j)$ linearly for $0 \le j \le \lceil \Delta^{-1} \rceil$ and restrict to $[0,1]$. Associate to A the parameter $\lambda = \Delta^2 A_{\mathrm{c}}$. For $\lambda \in \mathbb{R}$ define

$$F_\lambda(t) = \int_0^t e(s^3 + \lambda s)\, ds, \qquad E_\lambda(t) = \frac{e(t^3 + \lambda t) - 1}{2}.$$

Fix $\lambda_0 \in \mathbb{R}$, and put

$$D = \partial_\lambda F_{\lambda_0}, \qquad E = E_{\lambda_0}, \qquad L_v = vD + E,$$
$$f(v) = \|L_v\|_\infty, \qquad r_* = \min_{v \in \mathbb{R}} f(v).$$

All norms in these definitions are sup norms on $[0,1]$. For $r > 0$, let $I_r = \{v \in \mathbb{R} : f(v) < r\}$ and write $|I_r|$ for its Lebesgue length.

Theorem 3.1. *For every fixed λ_0, one has $r_* > 0$. Suppose $\varepsilon_p \to 0$ and $\Delta = o(\varepsilon_p)$. Then*

$$\mathbb{P}_A\big(\|P_{p,A} - F_{\lambda_0}\|_\infty < \varepsilon_p\big) \sim \frac{2\varepsilon_p \Delta}{\|D\|_\infty}. \tag{5}$$

For every fixed $r > r_*$,

$$\Delta^{-2}\mathbb{P}_A\big(\|P_{p,A} - F_{\lambda_0}\|_\infty < r\Delta\big) \longrightarrow |I_r|. \tag{6}$$

If $r < r_*$, the event in (6) is empty for all sufficiently large primes.

Lemma 3.2. *For $A \neq 0$, $\|P_{p,A}\|_\infty \ll 1/|\lambda|$. For λ in a fixed compact interval,*

$$P_{p,A}(t) = F_\lambda(t) + \Delta E_\lambda(t) + \Delta^2 \left(\frac{g'_\lambda(t) - g'_\lambda(0)}{12} + \frac{\theta(1-\theta)}{2} g'_\lambda(t) \right) + O(\Delta^3), \tag{7}$$

uniformly for $t \in [0,1]$, where $g_\lambda(t) = e(t^3 + \lambda t)$ and $\theta = t/\Delta - \lfloor t/\Delta \rfloor$.

Proof. The variation bound in (3) applies up to $\lceil \Delta^{-1} \rceil$. Multiplying that estimate by Δ and using $|1 - e(A_{\mathrm{c}}/p)| \gg |A_{\mathrm{c}}|/p$ gives the first bound, which persists under interpolation. At $t = j\Delta$ the summands are $g_\lambda(x\Delta)$. The composite trapezoidal expansion gives

$$\Delta \sum_{x=1}^{j} g_\lambda(x\Delta) = F_\lambda(j\Delta) + \Delta E_\lambda(j\Delta) + \frac{\Delta^2}{12}\big(g'_\lambda(j\Delta) - g'_\lambda(0)\big) + O(\Delta^3).$$

Indeed, on a step of length Δ the trapezoid excess over the integral is $\Delta^3 g''_\lambda/12 + O(\Delta^4)$ at the left endpoint. Summing and using the Riemann approximation of g''_λ gives the derivative difference. A right-endpoint sum differs from the trapezoidal sum by $\Delta(g_\lambda(j\Delta) - g_\lambda(0))/2$. Linear interpolation of F_λ adds $\Delta^2\theta(1-\theta)F''_\lambda(t)/2 + O(\Delta^3)$. Interpolation of the other displayed terms changes them by $O(\Delta^3)$. The needed derivatives are bounded on compact λ sets uniformly through the last mesh vertex, and $F''_\lambda = g'_\lambda$. □

Proof of Theorem 3.1. As $F'_{\lambda_0}(0) = 1$, the target has positive norm. The first part of Lemma 3.2 therefore confines all hits at sufficiently small radii to a fixed compact interval of λ. On that interval the map $\lambda \mapsto F_\lambda$ is continuous and injective. For injectivity, differentiate equality of the two curves in t and obtain $e((\lambda - \lambda_0)t) = 1$ for every $t \in [0,1]$. Compactness and (7) now imply $\lambda \to \lambda_0$ uniformly among hits as the radius tends to zero.

The Taylor expansion in the sup norm gives

$$\|F_\lambda - F_{\lambda_0}\|_\infty = |\lambda - \lambda_0| \|D\|_\infty + O(|\lambda - \lambda_0|^2). \tag{8}$$

Here $D(t) = 2\pi i \int_0^t se(s^3 + \lambda_0 s)\, ds$ is nonzero. At radius ε_p with $\Delta = o(\varepsilon_p)$, the hit set is sandwiched by intervals about λ_0 of radii $(1 + o(1))\varepsilon_p / \|D\|_\infty$, by (8) and the uniform $O(\Delta)$ error. The λ-lattice has spacing Δ^2, and division by $p = \Delta^{-3}$ gives (5).

At radius $r\Delta$, estimates (7) and (8) confine hits to bounded $v = (\lambda - \lambda_0)/\Delta$. In this range

$$\Delta^{-1}(P_{p,A} - F_{\lambda_0}) = vD + E + O(\Delta) \tag{9}$$

uniformly in the sup norm. The function f is convex and tends to infinity as $|v| \to \infty$, since $D \neq 0$. To see that $r_* > 0$, observe that E is outside the real span of D. The leading term of $D(t)$ at zero is $\pi i t^2$. The leading term of $E(t)$ is $\pi i \lambda_0 t$ if $\lambda_0 \neq 0$, and $\pi i t^3$ if $\lambda_0 = 0$. Thus their orders of vanishing differ. The real span of D is closed.

For $r > r_*$ the set I_r is a nonempty bounded open interval and its boundary consists of two points. The v-lattice has spacing Δ, with an offset depending on p and λ_0. Comparison in (9) sandwiches the event between the sublevel intervals for $r - C\Delta$ and $r + C\Delta$. Their lengths converge to $|I_r|$ by continuity and convexity, since a convex coercive function has no plateau at a level strictly above its minimum. Dividing the lattice count by p yields (6). For $r < r_*$, (9) excludes all hits for sufficiently small Δ. □

4. The first hit for the zero-parameter target

We now set $\lambda_0 = 0$. In this section $D = \partial_\lambda F_0$ and $E = E_0$. Let v_* denote the minimizer of f and let t_* be the maximizing argument of $|L_{v_*}|$. Their uniqueness is proved below.

Theorem 4.1. *For $\lambda_0 = 0$, one has $0.536 < r_* < 0.541$, and the distance to the closest arithmetic path has the expansion*

$$\min_{A \in \mathbb{F}_p} \|P_{p,A} - F_0\|_\infty = r_*\Delta + \kappa\Delta^2 + o(\Delta^2), \qquad \kappa = -\frac{\pi v_*^2}{3r_*} \operatorname{Im} L_{v_*}(t_*) > 0. \tag{10}$$

In particular, the open ball of radius $r_*\Delta$ contains no arithmetic path for all sufficiently large primes.

Lemma 4.2. *Let $J = [-0.75, -0.54]$. The minimizer v_* is unique and belongs to the interior of J. For each $v \in J$, the maximizing argument $t(v)$ of $|L_v(t)|$ is unique, lies in $(0.87, 0.94)$, and satisfies $\operatorname{Im} L_v(t(v)) < -0.35$. The function f is differentiable at v_* and $f'(v_*) = 0$.*

Proof. We give the finite inequalities and their continuation off the grid. Let G be the grid $\{k/1000 : 0 \le k \le 1000\}$, and put

$$U_v(t) + iW_v(t) = L_v(t)e(-t^3).$$

The following bounds hold on G. Rows with $v \in \partial J$ are checked at both endpoints of J.

Quantity	Arguments	Bound
$\lvert L_v(t)\rvert$	$v = -0.645$, all t	< 0.537
$\lvert L_v(t)\rvert$	$v \in \partial J$, $t \notin (0.83, 0.97)$	< 0.516
$\operatorname{Im} L_v(t)$	$v \in \partial J$, $t \in [0.83, 0.97]$	< -0.365
$\lvert L_v(t)\rvert$	$v \in \partial J$, $t = 0.904$	> 0.542
$U_v(t)$	$v \in \partial J$, $t = 0.904$	> 0.536
$W_v(t)$	$v \in \partial J$, $t \in [0.83, 0.87]$	> 0.044
$W_v(t)$	$v \in \partial J$, $t \in [0.94, 0.97]$	< -0.099
$W_v'(t)$	$v \in \partial J$, $t \in [0.87, 0.94]$	< -2.4

To verify the table, truncate $e(t^3)$ and the series

$$D(t) = 2\pi i t^2 \sum_{n=0}^{\infty} \frac{(2\pi i t^3)^n}{n!(3n+2)}$$

after $n = 40$. The error in the exponential and in the series before the prefactor $2\pi i t^2$ is at most $7^{41}/41!$. For the series for D, integrate the exponential remainder against s before multiplying by $2\pi i$. Enclose π between $N/10^{40}$ and $(N+1)/10^{40}$ with $N = 31415926535897932384626433832795028841971$. These bounds follow from Machin's identity and the alternating arctangent series. Check the norm bounds by squaring the positive comparison values. The accompanying script `verify-birch-boundary.mjs` performs the checks using integer interval endpoints, outward rounding after every product and quotient, and the rational Taylor error above. It derives the enclosure of π using 80 terms in each of the two arctangent series in Machin's identity.

For passage off the grid, differentiate to get

$$L_v'(t) = \pi i t(2v+3t)e(t^3),$$
$$U_v'(t) = 6\pi t^2 W_v(t), \qquad W_v'(t) = \pi t(2v+3t) - 6\pi t^2 U_v(t). \tag{11}$$

On $J \times [0,1]$, $|L_v| < 3.36$ and $|L_v'| < 7$. The former follows from $|D(t)| \le \pi t^2$ and $|E| \le 1$. These formulas also give $|W_v'| < 72$ and $|W_v''| < 1500$. Each closed interval in the table includes its grid endpoints. The distance to its nearest grid point is at most 0.0005.

The map $v \mapsto |L_v(t)|$ is convex for fixed t. Thus for all $v \in J$ and $t \notin (0.83, 0.97)$ the first time-region bound gives $|L_v(t)| < 0.516 + 0.0035$. At $t = 0.904$, affinity of U_v gives $f(v) \ge U_v(0.904) > 0.536$ throughout J. It follows that every maximizing t for $v \in J$ is in $(0.83, 0.97)$. Affinity of $\operatorname{Im} L_v$ then gives $\operatorname{Im} L_v(t) < -0.365 + 0.0035 < -0.35$ there. The estimate $f(-0.645) < 0.537 + 0.0035 < 0.542$ and $f(v) > 0.542$ at both endpoints of J place every global minimizer in its interior by convexity.

The three inequalities for W_v and W_v' extend between the endpoints of J by affinity. Between grid points, the derivative bounds show that W_v stays positive on $[0.83, 0.87]$, negative on $[0.94, 0.97]$, and strictly decreasing on $[0.87, 0.94]$. The identity

$$\frac{d}{dt}|L_v(t)|^2 = 2\pi t(2v+3t)W_v(t)$$

and positivity of $2v + 3t$ on $J \times [0.83, 0.97]$ prove uniqueness of $t(v)$. The implicit function theorem applied to $W_v(t(v)) = 0$ gives smooth dependence on v in the interior of J.

Finally $D(t) \neq 0$ for $t \in (0,1]$. In fact,

$$\operatorname{Re} D(t) = -\frac{2\pi}{3} \int_0^{t^3} u^{-1/3} \sin(2\pi u)\, du < 0.$$

The integral is positive through the first half-period. On the second half-period it decreases and remains positive at $t = 1$, by pairing each u in the first half-period with $u + 1/2$ using the decreasing weight. For $v \in J$, the squared maximum is the supremum over $t \in [0.83, 0.97]$ of the quadratics $|vD(t) + E(t)|^2$, whose leading coefficients have a strictly positive common lower bound. Subtracting that lower bound times v^2 leaves a convex supremum. Hence f^2 is strictly convex on J and the minimizer is unique. Differentiability and $f'(v_*) = 0$ follow from the smooth maximizing argument and the interior minimum. □

Proof of Theorem 4.1. At $\lambda_0 = 0$ the rescaled lattice is $v = A_c\Delta$. For bounded v, put $\lambda = \Delta v$ in (7) and use Taylor expansion in λ. This gives

$$\Delta^{-1}(P_{p,A} - F_0) = L_v + \Delta T_v + O(\Delta^2), \tag{12}$$

where

$$T_v(t) = \pi i \left[\frac{v^2}{3} \big(e(t^3) - 1\big) + \big(vt + \big(\tfrac{1}{2} + 3\theta(1-\theta)\big)t^2\big)\, e(t^3) \right].$$

To see the coefficient of v^2, use $\partial_\lambda^2 F_0(t) = \frac{2\pi i}{3}(e(t^3) - 1)$. The other terms come from $\partial_\lambda E_0(t) = \pi i t e(t^3)$ and $g_0'(t) = 6\pi i t^2 e(t^3)$ in (7).

At a maximizing argument $t = t(v)$ for $v \in J$, one has $W_v(t) = 0$. Thus $\operatorname{Re}(\overline{L_v(t)}\, ie(t^3)) = 0$, and the radial correction in (12) is exactly

$$\frac{\operatorname{Re}(\overline{L_v(t(v))} T_v(t(v)))}{|L_v(t(v))|} = -\frac{\pi v^2 \operatorname{Im} L_v(t(v))}{3f(v)}.$$

This expression is positive on J by Lemma 4.2, and it is continuous in a neighborhood of v_*. Its value at v_* is κ.

Choose a lattice point v_p nearest to v_*. Then $f(v_p) = r_* + o(\Delta)$ by $f'(v_*) = 0$. A maximizing argument for the norm on the right of (12) tends to t_*, by uniform convergence and uniqueness of the maximum at v_*. At all such arguments, the radial term involving $ie(t^3)$ tends to zero uniformly in the fractional position $\theta \in [0,1]$. The remaining radial term tends to κ. This gives the upper bound $r_* + \kappa\Delta + o(\Delta)$ for the minimum of the scaled norm.

Conversely let v_p correspond to a minimizing arithmetic path. The upper bound ensures a distance $O(\Delta)$, so the localization in Theorem 3.1 confines v_p to a compact set. By (9), $f(v_p) \to r_*$. Uniqueness implies $v_p \to v_*$, hence $v_p \in J$ for all sufficiently large p. Evaluate (12) at $t(v_p)$. The radial formula gives a lower bound $f(v_p) + \kappa\Delta + o(\Delta) \geq r_* + \kappa\Delta + o(\Delta)$. This proves (10) and its consequence at the critical radius. □

5. Prime subsequences with fixed reciprocal coefficient

The two-parameter bound in Theorem 2.1 uses averaging over B. With $B = 1$, the next theorem exhibits both collapse and noncollapse for initial windows whose length grows arbitrarily slowly relative to the chosen primes.

Theorem 5.1. *There are constants $c, \delta > 0$ and a choice of two prime sequences p_H^- and p_H^+ indexed by all sufficiently large integers H, with both primes at each index exceeding any prescribed bound in H, such that*

$$H^{-1/2} \max_{j \le H} |V_{p_H^-, A, 1}(j)| \longrightarrow 0 \qquad \text{in probability over } A,$$

$$\mathbb{P}_A \left(|V_{p_H^+, A, 1}(H)| > \delta \sqrt{H} \right) \ge c \qquad \text{for all sufficiently large } H.$$

The primes can be chosen so that $H^3/p_H^{\pm} \to 0$. On both sequences the corresponding normalized Birch maximum then converges in probability to zero.

Proof. Let $Q = H!$ and let q denote the inverse of p modulo Q for $p > H$. For $1 \le x \le H$,

$$e_p(x^{-1}) = e\left(-\frac{q}{x} + \frac{1}{px} \right). \tag{13}$$

Indeed, $(1 + kp)/x$ represents x^{-1} modulo p when $k \equiv -q \pmod{x}$. For $p \equiv 1 \pmod{Q}$ the phase in (13) differs from 1 by $O(1/x)$. The initial partial sums therefore differ from the geometric sums $\sum_{x \le j} e_p(Ax)$ by $O(\log H)$ uniformly in j. The geometric bound and uniform distribution of A imply collapse after division by $\sqrt{H}$. Dirichlet's theorem allows the primes in this class to exceed any prescribed bound for each H.

For noncollapse, let S be the set of integers in $[H/2, H]$ with a prime factor exceeding $\sqrt{H}$. We have $\#S \gg H$. To see this, count the multiples in that interval of primes $\ell \in (\sqrt{H}, H^{3/4}]$. These sets of multiples are disjoint, since an integer at most H has at most one prime factor exceeding $\sqrt{H}$. Mertens' theorem for primes gives $\sum_{\sqrt{H} < \ell \le H^{3/4}} 1/\ell = \log(3/2) + o(1)$, and the floor errors are $O(H^{3/4})$. We use Dirichlet's and Mertens' theorems in their standard forms [1].

Take q uniform on $(\mathbb{Z}/Q\mathbb{Z})^\times$ and ϑ uniform on $[0, 1]$, independently, and define

$$g_q(\vartheta) = H^{-1/2} \sum_{x \in S} e(\vartheta x - q/x).$$

We claim that $\mathbb{E}_{q,\vartheta} |g_q(\vartheta)|^4 \ll 1$. Integration in ϑ forces $x + y = u + v$ among quadruples from S. For the remaining average, put $R = 1/x + 1/y - 1/u - 1/v$ and let d be its denominator in lowest terms. Since $d \mid Q$, the reduction of q is uniform in $(\mathbb{Z}/d\mathbb{Z})^\times$, giving the normalized Ramanujan sum $\mu(d)/\varphi(d)$ when $R \ne 0$. If $R = 0$, equality of sums gives equality of the unordered pairs, by equality of products. These pairs contribute $O(H^2)$ to the unnormalized fourth moment.

When all four distinguished prime factors above $\sqrt{H}$ differ, each divides d since it occurs in exactly one reciprocal denominator. Thus $\varphi(d) \gg H^2$, and the $O(H^3)$ triples of free coordinates give $O(H)$ in absolute value. When at least two of the factors agree, choose their two positions and a common prime $\ell > \sqrt{H}$. There are at most $O((H/\ell)^2 H)$ choices since three coordinates determine the fourth. Summing over ℓ gives $O(H^{5/2})$ quadruples. For each nondiagonal one, $0 < |R| \ll 1/H$ so $d \gg H$. The elementary bound $\varphi(d) \ge \sqrt{d/2}$ follows from the prime-power formula for φ. The contribution of these quadruples in absolute value is $O(H^2)$. Division by H^2 proves the fourth-moment claim.

Consequently, for each H some unit q_H modulo Q has $\int_0^1 |g_{q_H}(\vartheta)|^4 \, d\vartheta \ll 1$ with an absolute constant. Choose $p \equiv q_H^{-1} \pmod{Q}$, arbitrarily large, by Dirichlet's theorem. Let $f(A) = H^{-1/2} V_{p,A,1}(H)$ and let $g(A)$ be the same sum restricted

to S. By (13) the difference $g(A) - g_{q_H}(A/p)$ is $O(\log H/(p\sqrt{H}))$ uniformly in A. The fourth moment of $g_{q_H}(A/p)$ agrees with its continuous fourth moment when $p > 2H$, by orthogonality of the integer frequencies. Thus $\mathbb{E}_A|g(A)|^4 \ll 1$. Also

$$\mathbb{E}_A|f(A)|^2 = 1, \qquad \mathbb{E}_A\big(f(A)\overline{g(A)}\big) = \#S/H \gg 1.$$

Splitting the latter expectation into $|f| \le \delta$ and $|f| > \delta$, Cauchy–Schwarz and Hölder give

$$\#S/H \le \delta\,\mathbb{E}_A|g| + \big(\mathbb{E}_A|f|^2\big)^{1/2}\big(\mathbb{E}_A|g|^4\big)^{1/4}\mathbb{P}_A(|f| > \delta)^{1/4}.$$

Since $\mathbb{E}_A|g| \le 1$, taking δ below half the absolute lower bound for $\#S/H$ proves noncollapse. Choose both prime sequences with $p_H^{\pm} \ge H^4$ if needed. In this regime the variation bound leading to (3) applies with $p_H^{\pm}, H$ and gives the concluded Birch collapse. □

6. Further questions

Question 6.1. With $B = 1$ fixed and $H = \lfloor p^{1/3} \rfloor$, do constants $c, \delta > 0$ exist such that $\mathbb{P}_A(|V_{p,A,1}(H)| > \delta\sqrt{H}) \ge c$ for all sufficiently large primes? A sufficient condition is the uniform fourth-moment bound

$$\sum_{\substack{1\le x,y,u,v\le H\\ x+y=u+v}} e_p(x^{-1} + y^{-1} - u^{-1} - v^{-1}) \ll H^2,$$

by orthogonality in A, the exact second moment, and Paley–Zygmund. The congruence construction in Theorem 5.1 controls windows H chosen before the primes.

Question 6.2. For general λ_0, determine the first correction to $\min_A \|P_{p,A} - F_{\lambda_0}\|_\infty$ beyond $r_*\Delta$, and characterize the hit count at radius $r_*\Delta$. The zero-parameter proof identifies two quantities governing this correction, the maximizing set of $|L_v(t)|$ and the radial part of the second-order interpolation term.

Disclosure of AI assistance

ChatGPT assisted with the derivation and checking of proof arguments, implementation and execution of the interval certificate, and preparation and editing of the manuscript.

References

[1] H. Davenport, *Multiplicative Number Theory*, 3rd ed., revised by H. L. Montgomery, Grad. Texts in Math., vol. 74, Springer-Verlag, New York, 2000.

[2] E. Kowalski and W. F. Sawin, *Kloosterman paths and the shape of exponential sums*, Compos. Math. **152** (2016), no. 7, 1489–1516, doi:10.1112/S0010437X16007351.

[3] Y. Lamzouri, *On the distribution of the maximum of cubic exponential sums*, J. Inst. Math. Jussieu **19** (2020), no. 4, 1259–1286, doi:10.1017/S1474748018000385.

[4] S. Schubr, *Kloosterman shrinking targets and cyclotomic trace geometry*, Luminary Research Archive preprint, edition 1, 2026, doi:10.5281/zenodo.24444480.

Luminary Oddities Math Lab