Translated-Window Statistics for Curves over Finite Fields

Critical Covariance Laws, Beyond-Poisson Tails, and a Structure–Randomness Program

Xiyu Hu

17 July 2026

Research note and status report

Abstract

Let $C \subset \mathbb{F}_p^2$ be a plane curve and let $\Omega \subset \mathbb{F}_p^2$ be obtained by reducing a bounded Archimedean window of linear scale H. For a uniformly translated window, write

$$N_\Omega(a) = \#\big(C \cap (a + \Omega)\big), \qquad a \in \mathbb{F}_p^2.$$

The critical two-dimensional regime is $H \asymp \sqrt{p}$, equivalently $|\Omega| \asymp p$.

This note records the current state of a program to combine local Fourier analysis, bilinear trace-function estimates, a non-tangential truncation, and Bombieri–Pila type algebraic certification. Results are arranged by mathematical importance and prospective publishability. The strongest result is an FKMS-dependent critical covariance theorem for every even monomial graph $y = x^d$ with $d \geq 6$: for arbitrary localized window shapes and local displacements,

$$\operatorname{Cov}\big(N_\Omega(a), N_\Omega(a+h)\big) = \frac{|\Omega \cap (\Omega + h)|}{p} + o(1),$$

with a power-saving error. The proof uses an exact positive Fourier factorization of the overlap function on a small auxiliary torus and the 2026 bilinear trace-function theorem of Fouvry–Kowalski–Michel–Sawin for a gallant root-count sheaf.

A second principal result is a beyond-Poisson upper-tail theorem. Fixed integral polynomial arcs occur in a proportion $\asymp p^{-1}$ of critical translations and create clusters of size $\asymp p^{1/(2e)}$ for a degree-e graph. This is enormously larger than a Poisson tail at the same height and forces divergence of every fixed factorial moment of order $k > 2e$. For the parabola and a square window, the fourth factorial moment has liminf at least 14/3, whereas Poisson(1) has fourth factorial moment 1. These statements do not rule out weak Poisson convergence at fixed counts; they prove that any such limit must coexist with non-Poisson moderate deviations and failure of untruncated moment convergence.

The note also gives exact laws for quadratic graphs, a Burgess-based critical law for cubic graphs, a universal displacement-averaged spectral benchmark, a non-tangential truncation whose total-variation cost tends to zero, and a self-contained determinant certificate that sends dense microscopic clusters to integral auxiliary curves where Bombieri–Pila applies. Unclosed parts, especially the critical weak law and separated higher correlations, are stated as precise research problems. This is a research draft rather than a peer-reviewed paper; novelty and the sheaf-theoretic input require specialist verification before submission.

Contents

1 Ranked status of the project

The following table is the intended reading order. "Closed" means that a complete proof is included or reduced to a standard cited theorem. "FKMS-dependent" means that the only deep external input is the March 2026 preprint [7]; the deduction in this note is explicit, but the input itself is not yet a journal theorem.

Rank	Result	Status	Main content
1	Critical covariance for even monomials	Closed, FKMS-dependent	Arbitrary localized shapes; $H \asymp \sqrt{p}$; power-saving covariance and variance; Euclidean overlap kernel.
2	Beyond-Poisson algebraic tails	Closed	Tail probability $\gg p^{-1}$ at cluster height $p^{1/(2e)}$; factorial moments diverge for $k > 2e$; parabola fourth-moment liminf $\geq 14/3$.
3	Cubic and quadratic critical laws	Closed	Exact arbitrary-window formula for quadratics; Burgess power saving for cubic graphs and convex/column-convex windows.
4	Non-tangential and algebraic exceptional stages	Closed	Close-pair exceptional probability $\ll \lvert\Omega\rvert L^2/p^2$; determinant certificate; Bombieri–Pila microbox bound.
5	Universal wave benchmarks	Closed	Exact Fourier/physical correlation identities; displacement-averaged spectral flattening; exact GL_2 orientation average.
6	Weak critical law and higher correlations	Open	Truncated Poisson/non-Poisson law, separated third and higher correlations, general fixed curves and polynomials.

The strongest plausible submission package consists of Rank 1 together with the shape-free factorization, with Ranks 2 and 4 supplying the distributional and structural context. No exhaustive MathSciNet or Zentralblatt priority search has been completed.

2 Setup and the strongest critical covariance theorem

Choose centered representatives

$$\langle x\rangle_p \in \left[-\frac{p-1}{2}, \frac{p-1}{2}\right] \cap \mathbb{Z}$$

for $x \in \mathbb{F}_p$. A bounded set $K \subset \mathbb{R}^2$ and a scale $H = o(p)$ give a finite-field window

$$\Omega_{p,H,K} = (HK \cap \mathbb{Z}^2) \pmod{p}.$$

This is not an intrinsic topology on $\mathbb{F}_p^2$; it is a mixed local condition consisting of a congruence modulo p and an Archimedean height cutoff.

For a finite set $C \subset \mathbb{F}_p^2$, put

$$N_\Omega(a) = \#(C \cap (a + \Omega)), \qquad A = |\Omega|.$$

When $|C| \sim p$,

$$\mathbb{E}_a N_\Omega(a) = \frac{|C|A}{p^2} \sim \frac{A}{p}.$$

For a regular planar window, $A \asymp H^2$, so $H \asymp \sqrt{p}$ is the scale at which the mean is of order one.

A window Ω is called (R,H)-localized if, after a translation, its centered lift is contained in $[-RH, RH]^2$.

Theorem 2.1 (Critical covariance for even monomial graphs)**.** *Fix an even integer $d \geq 6$, constants $R \geq 1$ and $0 < c_0 < C_0$. There are a finite set of primes $\mathcal{P}_d$ and a number $\eta_d > 0$ with the following property. Let $p \notin \mathcal{P}_d$, let*

$$C_d = \{(x, x^d) : x \in \mathbb{F}_p\},$$

let $c_0\sqrt{p} \le H \le C_0\sqrt{p}$*, and let* Ω *be an* (R,H)*-localized window of size* A*. If* $h \in \mathbb{F}_p^2$ *satisfies*

$$\|\langle h\rangle_p\|_\infty \le RH,$$

then

$$\Gamma_{d,\Omega}(h) := \mathbb{E}_a\big(N_\Omega(a) - A/p\big)\big(N_\Omega(a+h) - A/p\big) = \frac{g_\Omega(h)}{p} + O_{d,R,c_0,C_0}\left(\frac{HA}{p^2} + \frac{A}{p^{1+\eta_d}}\right), \qquad (1)$$

where

$$g_\Omega(h) = |\Omega \cap (\Omega + h)|.$$

In particular, if $A \ll p$*, then for some* $\eta'_d > 0$*,*

$$\operatorname{Var} N_\Omega = \frac{A}{p} + O(p^{-\eta'_d}), \qquad \Gamma_{d,\Omega}(h) = \frac{g_\Omega(h)}{p} + O(p^{-\eta'_d}).$$

Remark 2.2 (Status of the theorem)**.** *The proof below is complete after invoking [7, Theorems 1.1 and 1.3]. The latter is a March 2026 preprint. The case* $d = 6$ *is covered by checking the gallant condition directly for the standard representation of* S_5*; the example stated in [7, Section 9.9] is phrased with a slightly more conservative degree threshold.*

Corollary 2.3 (Euclidean overlap law for arbitrary shapes)**.** *Let* $K \subset \mathbb{R}^2$ *be bounded and Jordan measurable, let*

$$\Omega_p = (HK \cap \mathbb{Z}^2) \pmod p, \qquad \frac{H}{\sqrt{p}} \longrightarrow \tau \in (0,\infty),$$

and let $h_p/H \to z \in \mathbb{R}^2$ *in centered coordinates. Under the hypotheses of Theorem 2.1,*

$$\operatorname{Var} N_{\Omega_p} \longrightarrow \tau^2|K|,$$

and

$$\Gamma_{d,\Omega_p}(h_p) \longrightarrow \tau^2|K \cap (K+z)|.$$

Thus boxes, balls, ellipses, disconnected regions, and rough Jordan-measurable shapes obey the same second-order law; the shape enters only through its overlap kernel.

2.1 Exact graph decomposition

For a polynomial graph

$$C_f = \{(x, f(x)) : x \in \mathbb{F}_p\},$$

define

$$r_f(a,b) = \#\{x \in \mathbb{F}_p : f(x+a) - f(x) = b\}.$$

For $a \ne 0$, put $t_a(b) = r_f(a,b) - 1$. Since $\sum_b r_f(a,b) = p$, one has $\sum_b t_a(b) = 0$.

Proposition 2.4 (Exact graph covariance decomposition)**.** *For every set window* Ω *and every* $h \in \mathbb{F}_p^2$*,*

$$\Gamma_{f,\Omega}(h) = \frac{g_\Omega(h)}{p} - \frac{1}{p^2}\sum_{b\in\mathbb{F}_p} g_\Omega((0,b) - h) + \frac{1}{p^2}\sum_{a\ne 0}\sum_{b\in\mathbb{F}_p} t_a(b) g_\Omega((a,b) - h). \qquad (2)$$

Proof. For $a = 0$, the difference equation has p solutions when $b = 0$ and none otherwise. For $a \neq 0$, it has $1 + t_a(b)$ solutions. Insert this difference count into the physical-space correlation identity proved in Proposition 5.1 below. The contribution of the constant function 1 is A^2/p^2 and cancels the squared mean. $\square$

If the lift of Ω lies in an $O(H)$-box and

$$m_\Omega(x) = \#\{y : (x, y) \in \Omega\},$$

then

$$\sum_b g_\Omega((0, b) - h) = \sum_x m_\Omega(x) m_\Omega(x + h_1) \ll HA. \tag{3}$$

All arithmetic difficulty is therefore concentrated in the final term of (2).

2.2 Shape-free factorization of the overlap

A general ball or body does not give a product weight in the two difference variables. Autocorrelation supplies an exact substitute.

Lemma 2.5 (Positive Fourier factorization on a small torus)**.** *Let $\widetilde{\Omega} \subset \mathbb{Z}^2$ be contained in a box of side at most $2RH$, and put $A = |\widetilde{\Omega}|$. Choose an integer $Q > 8RH$ and regard $\widetilde{\Omega}$ as a subset of $(\mathbb{Z}/Q\mathbb{Z})^2$. Its ordinary overlap function*

$$g_{\widetilde{\Omega}}(u) = |\widetilde{\Omega} \cap (\widetilde{\Omega} + u)|$$

satisfies, throughout its non-wrapping difference range,

$$g_{\widetilde{\Omega}}(u) = \sum_{\rho \in (\mathbb{Z}/Q\mathbb{Z})^2} c_\rho e_Q(\rho \cdot u), \qquad c_\rho = \frac{1}{Q^2} \left| \widehat{\mathbf{1}_{\widetilde{\Omega}}}^{(Q)}(\rho) \right|^2 \geq 0, \tag{4}$$

with

$$\sum_\rho c_\rho = A. \tag{5}$$

Consequently, on the single $O_R(H) \times O_R(H)$ support box, the weight

$$(a, b) \longmapsto g_\Omega((a, b) - h)$$

is an exact sum of rank-one weights

$$\sum_\rho c_\rho \alpha_\rho(a) \beta_\rho(b),$$

where $|\alpha_\rho| = |\beta_\rho| = 1$ and the total coefficient mass is A.

Proof. On the auxiliary torus, the Fourier transform of an autocorrelation is the squared modulus of the Fourier transform. Fourier inversion gives (4); Parseval at $u = 0$ gives (5). The choice of Q prevents wraparound on the ordinary difference set. Translating by h multiplies each Fourier coefficient by a phase. The indicator of the single support box is absorbed into the one-variable coefficient sequences, preventing periodic copies from entering the finite-field sum. $\square$

The lemma is stronger than approximating a body by rectangles: the relevant projective tensor norm is at most A, independently of boundary smoothness, convexity, or topology.

2.3 The root-count sheaf

Fix even $d \geq 6$ and put

$$n = d - 1, \qquad \Phi_d(X) = (X+1)^d - X^d.$$

Thus n is odd and $\deg \Phi_d = n$.

Lemma 2.6 (The difference polynomial is supermorse)**.** *Over $\mathbb{C}$, the polynomial Φ_d has simple critical points and pairwise distinct critical values. Consequently, outside a finite set of primes depending only on d, its reduction modulo p is supermorse.*

Proof. The derivative is

$$\Phi_d'(X) = d\big((X+1)^n - X^n\big).$$

Its critical points are indexed by $\zeta \in \mu_n \setminus \{1\}$:

$$x_\zeta = \frac{1}{\zeta - 1}, \qquad \frac{x_\zeta + 1}{x_\zeta} = \zeta.$$

At such a point,

$$\Phi_d''(x_\zeta) = dnx_\zeta^{n-1}(\zeta^{n-1} - 1) = dnx_\zeta^{n-1}(\zeta^{-1} - 1) \neq 0,$$

so every critical point is simple. Its critical value is

$$\Phi_d(x_\zeta) = (\zeta - 1)^{-n}.$$

Suppose the values corresponding to ζ and η are equal. Then

$$\frac{\eta - 1}{\zeta - 1} = \omega, \qquad \omega^n = 1.$$

Taking absolute values gives $|\eta - 1| = |\zeta - 1|$. Two points on the unit circle at the same distance from 1 are equal or complex conjugate, so $\eta = \zeta$ or $\eta = \zeta^{-1}$. In the second case

$$\omega = \frac{\zeta^{-1} - 1}{\zeta - 1} = -\zeta^{-1}, \qquad \omega^n = (-1)^n = -1,$$

contradicting $\omega^n = 1$. Distinctness follows. Reduction preserves this property away from the finite set of primes dividing the relevant discriminants and resultants. □

For $c \in \mathbb{F}_p$, define

$$T_{d,p}(c) = \#\{z \in \mathbb{F}_p : \Phi_d(z) = c\} - 1.$$

Proposition 2.7 (Gallant root-count trace function)**.** *For every even $d \geq 6$ and all primes outside a finite set, $T_{d,p}$ is the trace function of a gallant and light sheaf of complexity bounded in terms of d alone.*

Proof. Consider the finite direct image of the constant sheaf under $\Phi_d : \mathbb{A}^1 \to \mathbb{A}^1$. Its Frobenius trace at c is the number of $\mathbb{F}_p$-points in the fiber; removing the constant summand gives $T_{d,p}$. By Lemma 2.6 and the standard monodromy theorem for supermorse polynomials, the geometric monodromy is S_n acting in its standard $(n-1)$-dimensional representation; see [7, Section 9.9] and [10].

Since $n = d - 1 \geq 5$, $A_n \triangleleft S_n$ is non-abelian simple and perfect. Its action on the standard representation is irreducible. One way to see this is that A_n is doubly transitive, so the endomorphism algebra of its permutation representation has dimension two; after removing the trivial line, the standard module cannot split further. Thus S_n contains a quasisimple normal subgroup acting irreducibly, which is the finite-monodromy condition in the definition of a gallant sheaf. On the maximal open set where Φ_d is finite etale, this summand is lisse and pure of weight zero. Because Φ_d is finite, its middle-extension stalks are mixed of integral weights at most zero. Hence the sheaf is light as well as gallant. Its complexity is bounded in terms of d. $\square$

Lemma 2.8 (Scaling of monomial difference fibers)**.** *For $a \neq 0$,*

$$\#\{x : (x+a)^d - x^d = b\} - 1 = T_{d,p}(ba^{-d}).$$

Proof. Write $x = az$ and factor out a^d. $\square$

2.4 Bilinear closure at the square-root scale

The needed input from [7] gives a power saving for bilinear forms

$$\sum_{m \asymp M} \sum_{n \asymp N} \alpha_m \beta_n K(m^b n^c)$$

when K is gallant and light, one variable has polynomial length, and $MN \geq p^{3/4+\delta}$.

Lemma 2.9 (Critical bilinear bound)**.** *Fix even $d \geq 6$ and $R \geq 1$. There exists $\eta_d > 0$ such that, for all sufficiently large good primes, if I, J are intervals of centered representatives contained in $[-R\sqrt{p}, R\sqrt{p}]$ and $|\alpha_a|, |\beta_b| \leq 1$, then*

$$\sum_{\substack{a \in I \\ a \neq 0}} \sum_{b \in J} \alpha_a \beta_b T_{d,p}(ba^{-d}) \ll_{d,R} p^{1-\eta_d}.$$

Proof. The contribution of $b = 0$ is $O_d(\sqrt{p})$. Split the remaining centered representatives by sign and into dyadic shells

$$M \leq |a| < 2M, \qquad N \leq |b| < 2N.$$

Negative signs give homothetic pullbacks of the same gallant sheaf. If necessary interchange the variables, changing the fixed exponents from $(-d, 1)$ to $(1, -d)$.

Choose $\delta \in (0, 1/4)$. If $MN < p^{3/4+\delta}$, the trivial estimate is $O_d(MN)$. Otherwise the larger of M, N is at least $p^{3/8+\delta/2}$, hence at least p^δ, and the bilinear theorem applies. It gives, for some $\eta_0 > 0$,

$$\sum_{a \asymp M} \sum_{b \asymp N} \alpha_a \beta_b T_{d,p}(ba^{-d}) \ll \|\alpha\|_2 \|\beta\|_2 (MN)^{1/2-\eta_0} \ll (MN)^{1-\eta_0}.$$

Since $MN \ll_R p$, every good block is $O(p^{1-\eta_0})$, while every bad block is $O(p^{3/4+\delta})$. Summing $O_R((\log p)^2)$ blocks and absorbing logarithms into a smaller positive power proves the claim. $\square$

Proof of Theorem 2.1. By Proposition 2.4, Lemma 2.8, and (3), it remains to bound

$$\mathcal{R}(h) = \sum_{a \neq 0} \sum_{b} T_{d,p}(ba^{-d}) g_\Omega((a,b) - h).$$

The support lies in a single $O_R(H) \times O_R(H)$ box of centered representatives. Lemma 2.5 gives

$$g_\Omega((a,b) - h) = \sum_\rho c_\rho \alpha_\rho(a) \beta_\rho(b), \qquad \sum_\rho |c_\rho| = A.$$

Apply Lemma 2.9 term by term:

$$|\mathcal{R}(h)| \leq \sum_{\rho} |c_\rho| \left| \sum_{a\neq 0,b} \alpha_\rho(a)\beta_\rho(b)T_{d,p}(ba^{-d}) \right| \ll_{d,R} Ap^{1-\eta_d}.$$

After division by p^2 and addition of the vertical correction, this is (1). ◻

Proof of Corollary 2.3. Jordan measurability gives

$$\frac{|\Omega_p|}{H^2} \longrightarrow |K|, \qquad \frac{g_{\Omega_p}(h_p)}{H^2} \longrightarrow |K \cap (K+z)|.$$

The error in Theorem 2.1 tends to zero because $A \ll H^2 \asymp p$. ◻

Corollary 2.10 (Affine monomial graphs)**.** *The conclusion remains valid for reductions of*

$$f(X) = \alpha X^d + \beta X + \gamma \in \mathbb{Z}[X], \qquad \alpha \neq 0,$$

with even $d \geq 6$, after excluding primes dividing α and changing the localization constant by a factor depending on β.

Proof. The fixed integral shear $(x, y) \mapsto (x, y - \beta x)$ reduces to $y = \alpha x^d + \gamma$. Translation removes γ, while the scalar α gives a homothetic pullback of the same root-count sheaf. ◻

3 Beyond-Poisson algebraic tails

This section records the strongest distributional information currently available. It is deliberately placed before the lower-degree covariance results because it changes the correct formulation of the limiting-law problem.

For an integer-valued random variable N, write

$$(N)_k = N(N-1)\cdots(N-k+1).$$

Theorem 3.1 (Mesoscopic algebraic tail)**.** *Let $f \in \mathbb{Z}[X]$ have degree $e \geq 2$, let $K \subset \mathbb{R}^2$ be bounded with nonempty interior, and suppose*

$$H \to \infty, \qquad H = o(p), \qquad 0 < c \leq H^2/p \leq C < \infty.$$

Let C_f be the reduction of the graph of f and let

$$\Omega = (HK \cap \mathbb{Z}^2) \pmod{p}.$$

Then there are constants $c_1, c_2, c_3 > 0$, depending on f and K, such that

$$\mathbb{P}_a(N_\Omega(a) \geq c_1 H^{1/e}) \geq c_2 \frac{H^2}{p^2} \geq \frac{c_3}{p}. \tag{6}$$

Moreover, for every fixed $k \geq 1$,

$$\mathbb{E}_a(N_\Omega(a))_k \gg_{f,K,k} \frac{H^{2+k/e}}{p^2}. \tag{7}$$

Consequently, for every fixed $k > 2e$,

$$\mathbb{E}(N_\Omega)_k \gg H^{k/e-2} \longrightarrow \infty.$$

Proof. Choose a square $z_0 + [-3\rho, 3\rho]^2$ contained in the interior of K. For a sufficiently small constant $c_* = c_*(f, K) > 0$, the integral arc

$$\mathcal{A}_H = \{(n, f(n)) : 0 \le n \le c_* H^{1/e}\}$$

has displacement from $(0, f(0))$ contained in $[-\rho H, \rho H]^2$. Every lattice point

$$q \in H\big(z_0 + [-\rho, \rho]^2\big)$$

produces a translation $a = (0, f(0)) - q$ for which $\mathcal{A}_H \subset a + HK$. There are $\gg_K H^2$ distinct translations and each contains $\gg_f H^{1/e}$ curve points. Dividing by p^2 gives both (6) and (7). $\square$

Corollary 3.2 (Quantitative separation from Poisson moderate deviations)**.** *Assume $H^2/p \to \lambda_K \in (0, \infty)$ and let $Z \sim \mathrm{Poisson}(\lambda)$ with fixed $\lambda > 0$. At*

$$m_p = c_1 H^{1/e} \asymp p^{1/(2e)},$$

Theorem 3.1 gives

$$\mathbb{P}(N_\Omega \ge m_p) \gg p^{-1}.$$

By the standard Chernoff bound,

$$\mathbb{P}(Z \ge m_p) \le \left(\frac{e\lambda}{m_p}\right)^{m_p} = \exp\big(-\Theta(p^{1/(2e)} \log p)\big).$$

Thus the critical finite-field process has a genuinely non-Poisson moderate-deviation tail, even if its probabilities at each fixed count were eventually to converge to Poisson probabilities.

Remark 3.3. *The tail-producing set has probability $\asymp p^{-1}$, so it vanishes in the weak topology. Therefore Theorem 3.1 does not disprove weak Poisson convergence. It proves instead that weak convergence, high-moment convergence, and moderate-deviation asymptotics are different questions.*

3.1 The parabola: Poisson first and second order, non-Poisson fourth order

Let $C : y = x^2$ and $\Omega = [0, H)^2$, with $H < p/4$. A translated count has the same law as

$$N_{s,t} = \#\{u \in [0, H) : u^2 + su + t \pmod p \in [0, H)\},$$

where (s, t) is uniform in $\mathbb{F}_p^2$.

Proposition 3.4 (Exact first two factorial moments)**.** *One has*

$$\mathbb{E}N = \frac{H^2}{p}, \qquad \mathbb{E}(N)_2 = \frac{H^3(H-1)}{p^2},$$

and

$$\operatorname{Var} N = \frac{H^2}{p} - \frac{H^3}{p^2}.$$

Thus, if $H^2/p \to \lambda$, then the first two factorial moments tend to λ and λ^2. In particular, at $\lambda = 1$ the mean and variance both tend to one.

Proof. The mean is the general translation average. For distinct $u_1, u_2 \in [0, H)$ and arbitrary $v_1, v_2 \in [0, H)$, the two equations

$$u_i^2 + su_i + t = v_i \qquad (i = 1, 2)$$

have a unique solution (s, t) because $u_1 - u_2 \neq 0$. There are $H(H-1)$ ordered pairs (u_1, u_2) and H^2 ordered pairs (v_1, v_2), giving the second factorial moment. The variance formula also follows from Theorem 4.1. $\square$

Proposition 3.5 (Parabola fourth moment)**.** *Let $C : y = x^2$, let $\Omega = [0, H)^2$, and assume $H^2/p \to 1$. If N is the number of curve points in a uniformly translated copy of Ω, then*

$$\liminf_{p\to\infty} \mathbb{E}(N)_4 \geq \frac{14}{3}.$$

For every fixed $k > 4$,

$$\mathbb{E}(N)_k \gg_k H^{k/2-2} \longrightarrow \infty.$$

Since Poisson(1) has $\mathbb{E}(Z)_4 = 1$, untruncated convergence of factorial moments already fails in order four.

Proof. After a bijective reparameterization of the translation, N has the same law as

$$N_{s,t} = \#\{u \in [0, H) : u^2 + su + t \pmod p \in [0, H)\},$$

with (s, t) uniform. For $0 \leq c \leq H/2$, set

$$L_c = \left\lfloor \sqrt{H - 1 - c} \right\rfloor.$$

For every x_0 with $L_c \leq x_0 \leq H - 1 - L_c$, take

$$s = -2x_0, \qquad t = x_0^2 + c.$$

Then

$$u^2 + su + t = (u - x_0)^2 + c,$$

so $N_{s,t} \geq 2L_c + 1$. These parameter pairs are distinct for $2H < p$. Therefore

$$\mathbb{E}(N)_k \geq \frac{1}{p^2} \sum_{c=0}^{\lfloor H/2 \rfloor} (H - 2L_c)(2L_c + 1)_k.$$

For $k = 4$, the numerator is

$$16H \sum_{c \leq H/2} (H - c)^2 + o(H^4) = \frac{14}{3} H^4 + o(H^4),$$

which proves the asserted liminf. For $k > 4$, the numerator is $\gg H^{k/2+2}$; division by $p^2 \asymp H^4$ gives the divergence. $\square$

The combination of Proposition 3.4 and Proposition 3.5 is especially instructive: the process is Poisson-like at first and second order but fails at fourth order. Since $\mathbb{E}N^2 = \mathbb{E}(N)_2 + \mathbb{E}N = O(1)$ at the critical scale,

$$\mathbb{P}(N > M) \leq \frac{\mathbb{E}N^2}{M^2} \ll M^{-2},$$

so the family is tight even though its higher moments are not uniformly integrable.

For three graph points there is an exact resonance relation. If $u_2 = u_1 + a$, $u_3 = u_1 + b$, and $v_i \in [0, H)$, then a common pair (s, t) exists precisely when

$$(a - b)v_1 + bv_2 - av_3 \equiv ab(a - b) \pmod{p}.$$

The disappearance of u_1 from this compatibility condition is one concrete source of geometry-of-numbers structure in higher correlations.

4 Closed low-degree covariance laws

4.1 Quadratic graphs: an exact arbitrary-window identity

Let

$$q(x) = \alpha x^2 + \beta x + \gamma, \qquad \alpha \neq 0,$$

and put $C_q = \{(x, q(x)) : x \in \mathbb{F}_p\}$. For a weight $w : \mathbb{F}_p^2 \to \mathbb{C}$, define

$$X_w(a) = \sum_{P \in C_q} w(P - a), \qquad m_w(x) = \sum_y w(x, y),$$

and

$$g_w(h) = \sum_z w(z)\overline{w(z - h)}, \qquad k_w(t) = \sum_x m_w(x)\overline{m_w(x - t)}.$$

Theorem 4.1 (Exact quadratic variance and covariance)**.** *For every weight w and every $h = (h_1, h_2)$,*

$$\mathbb{E}X_w = \frac{1}{p}\sum_z w(z), \tag{8}$$

$$\operatorname{Cov}(X_w(a), X_w(a + h)) = \frac{1}{p}g_w(h) - \frac{1}{p^2}k_w(h_1), \tag{9}$$

$$\operatorname{Var}(X_w) = \frac{1}{p}\sum_z |w(z)|^2 - \frac{1}{p^2}\sum_x |m_w(x)|^2. \tag{10}$$

For a set window Ω,

$$\operatorname{Var} N_\Omega = \frac{A}{p} - \frac{1}{p^2}\sum_x m_\Omega(x)^2, \tag{11}$$

$$\Gamma_{q,\Omega}(h) = \frac{|\Omega \cap (\Omega + h)|}{p} - \frac{1}{p^2}\sum_x m_\Omega(x)m_\Omega(x - h_1). \tag{12}$$

No shape or regularity assumption is required.

Proof. For $(u, v) \in \mathbb{F}_p^2$,

$$\widehat{\mathbf{1}_{C_q}}(u, v) = \sum_x e_p(-ux - vq(x)).$$

If $v = 0$ and $u \neq 0$, this vanishes. If $v \neq 0$, it is a non-degenerate quadratic Gauss sum and has squared modulus p. Insert this spectrum into Proposition 5.1. The full inverse Fourier sum gives $p^2 g_w(h)$; the line $v = 0$ gives $pk_w(h_1)$ by one-dimensional Parseval. This proves the formulas. □

For a fixed compact convex body K, write

$$\ell_K(t) = |\{y : (t, y) \in K\}|, \quad A_K(z) = |K \cap (K + z)|, \quad J_K(s) = \int \ell_K(t)\ell_K(t - s)\, dt.$$

Standard lattice-point and Riemann-sum estimates give

$$|HK \cap \mathbb{Z}^2| = H^2|K| + O_K(H),$$

$$\sum_x m_\Omega(x)m_\Omega(x - \lfloor Hs \rfloor) = H^3 J_K(s) + O_K(H^2),$$

and

$$|\Omega \cap (\Omega + \lfloor Hz \rfloor)| = H^2 A_K(z) + O_K(H).$$

Consequently, if $H/\sqrt{p} \to \tau$,

$$\operatorname{Var} N_\Omega \to \tau^2|K|, \qquad \Gamma_{q,\Omega}(\lfloor Hz \rfloor) \to \tau^2|K \cap (K + z)|.$$

4.2 A general one-variable short-sum criterion

A lifted window is called column-convex of scale H if it lies in an $H \times H$ integer square and each vertical fiber is an interval. This includes lattice points in a fixed convex body, after changing H by a constant.

For a polynomial graph C_f, define

$$B_f(p; L) = \max_{a \in \mathbb{F}_p^\times} \max_{\substack{I \subset \mathbb{F}_p \text{ an interval} \\ |I| \le L}} \left| \sum_{b \in I} t_a(b) \right|.$$

Theorem 4.2 (Short-sum criterion)**.** *Let Ω be a lifted column-convex window of scale H and size A. Uniformly in h,*

$$\Gamma_{f,\Omega}(h) = \frac{g_\Omega(h)}{p} + O\left(\frac{HA}{p^2}(1 + B_f(p; 2H + 1))\right).$$

For $A \asymp H^2$ and $H \asymp \sqrt{p}$, the precise covariance law follows from

$$B_f(p; 2H + 1) = o(\sqrt{p}).$$

Proof. For fixed horizontal difference d, the function $b \mapsto g_\Omega(d, b)$ is supported on an interval of length at most $2H + 1$. As a sum of cross-correlations of interval fibers, its supremum plus total variation is bounded by

$$M_d = \sum_x \min(m_\Omega(x), m_\Omega(x - d)).$$

Moreover $\sum_d M_d \le HA$. Discrete summation by parts therefore gives

$$\left| \sum_b t_a(b) g_\Omega((a, b) - h) \right| \ll B_f(p; 2H + 1) M_{a - h_1}.$$

Sum over a, use $\sum M_d \le HA$, and include the vertical correction (3). □

4.3 Cubic graphs: Burgess closes the critical scale

Let

$$f(x) = \alpha x^3 + \beta x^2 + \gamma x + \delta, \qquad p \nmid 6\alpha,$$

and let χ denote the quadratic character of $\mathbb{F}_p$, with $\chi(0) = 0$.

Lemma 4.3 (Cubic difference fibers)**.** *For every $a \neq 0$,*

$$r_f(a, b) = 1 + \chi(\Delta_f(a, b)),$$

where

$$\Delta_f(a, b) = a\left(12\alpha b + a(4\beta^2 - 12\alpha\gamma - 3\alpha^2 a^2)\right).$$

Thus $t_a(b)$ is a quadratic character of a nonconstant affine-linear function of b.

Proof. The polynomial $f(x + a) - f(x) - b$ is quadratic in x with leading coefficient $3\alpha a$. The displayed expression is its discriminant. A nonzero quadratic polynomial over $\mathbb{F}_p$ has $1 + \chi(\Delta)$ roots, including the double-root case. □

The Burgess estimate in the form used here is

$$\left|\sum_{n \in I} \chi(n)\right| \ll_{r,\varepsilon} |I|^{1-1/r} p^{(r+1)/(4r^2)+\varepsilon};$$

see [9].

Theorem 4.4 (Cubic critical covariance)**.** *Let Ω be a lifted column-convex window of scale H and size A. For every fixed $r \geq 1$ and $\varepsilon > 0$,*

$$\Gamma_{f,\Omega}(h) = \frac{g_\Omega(h)}{p} + O_{r,\varepsilon}\left(\frac{HA}{p^2} + \frac{HA}{p^2} H^{1-1/r} p^{(r+1)/(4r^2)+\varepsilon}\right).$$

For $r = 2$ and $A \asymp H^2$,

$$\Gamma_{f,\Omega}(h) = \frac{g_\Omega(h)}{p} + O_\varepsilon\left(\frac{H^3}{p^2} + H^{7/2} p^{-29/16+\varepsilon}\right).$$

At $H \asymp \sqrt{p}$, the non-geometric error is $O_\varepsilon(p^{-1/16+\varepsilon})$. Hence fixed convex shapes have the same critical variance and overlap-kernel covariance limits as in Corollary 2.3.

Proof. Lemma 4.3 and Burgess imply

$$B_f(p; L) \ll_{r,\varepsilon} L^{1-1/r} p^{(r+1)/(4r^2)+\varepsilon}.$$

Insert this into Theorem 4.2. With $r = 2$, one obtains $B_f(p; L) \ll L^{1/2} p^{3/16+\varepsilon}$ and the stated error. □

5 Universal wave identities and averaged spectral benchmarks

Let $G = \mathbb{F}_p^2$ and normalize

$$e_p(t) = \exp(2\pi i t/p), \qquad \widehat{f}(\xi) = \sum_{x\in G} f(x)e_p(-\xi\cdot x).$$

Then

$$f(x) = \frac{1}{p^2}\sum_{\xi\in G}\widehat{f}(\xi)e_p(\xi\cdot x), \qquad \sum_x |f(x)|^2 = \frac{1}{p^2}\sum_\xi |\widehat{f}(\xi)|^2.$$

For a finite set $C \subset G$ and a weight $w : G \to \mathbb{C}$, put

$$X_w(a) = \sum_{P\in C} w(P-a), \qquad \mu_w = \frac{|C|}{p^2}\sum_x w(x).$$

Define

$$r_C(u) = \#\{(P,Q)\in C^2 : Q-P = u\}, \qquad g_w(u) = \sum_x w(x)\overline{w(x-u)}.$$

Proposition 5.1 (Exact Fourier and physical-space correlation formulas)**.** *For every $h \in G$,*

$$\begin{aligned}\Gamma_{C,w}(h) &:= \mathbb{E}_a\big(X_w(a)-\mu_w\big)\overline{\big(X_w(a+h)-\mu_w\big)}\\ &= \frac{1}{p^4}\sum_{\xi\neq 0}|\widehat{\mathbf{1}_C}(\xi)|^2|\widehat{w}(\xi)|^2 e_p(-\xi\cdot h), \end{aligned}\tag{13}$$

$$\mathbb{E}_a X_w(a)\overline{X_w(a+h)} = \frac{1}{p^2}\sum_{u\in G} r_C(u)g_w(u-h). \tag{14}$$

Proof. The function X_w is the convolution of $\mathbf{1}_C$ with the reflected weight $x \mapsto w(-x)$. Its Fourier transform is $\widehat{\mathbf{1}_C}(\xi)\overline{\widehat{w}(\xi)}$. Parseval, with the zero frequency removed after centering, gives (13). Expanding in physical space and grouping pairs by $Q-P$ gives (14). $\square$

For a fixed-degree absolutely irreducible non-linear plane curve, the standard Bombieri–Weil additive-character estimate [1, 18] gives

$$|\widehat{\mathbf{1}_C}(\xi)| \ll_D \sqrt{p} \qquad (\xi \neq 0),$$

and hence

$$\operatorname{Var} N_\Omega \ll_D \frac{|\Omega|}{p}.$$

This has the correct scale but not the leading constant.

5.1 Displacement-averaged spectral flattening

Theorem 5.2 (Global spectral flattening over displacement)**.** *Let $C \subset \mathbb{F}_p^2$ have n points and assume*

$$|\widehat{\mathbf{1}_C}(\xi)| \ll_D \sqrt{p} \qquad (\xi \neq 0).$$

Set

$$\mathfrak{s}_C = \frac{p^2 n - n^2}{p^2-1}$$

and

$$\Gamma^{\mathrm{unif}}_{C,\Omega}(h) = \frac{\mathfrak{s}_C}{p^4}(p^2 g_\Omega(h) - A^2).$$

Then

$$\frac{1}{p^2}\sum_{h\in G}\left|\Gamma_{C,\Omega}(h) - \Gamma^{\mathrm{unif}}_{C,\Omega}(h)\right|^2 \ll_D \frac{E_+(\Omega)}{p^4},$$

where

$$E_+(\Omega) = \#\{x_1 + x_2 = x_3 + x_4 : x_i \in \Omega\}.$$

Since $E_+(\Omega) \le A^3$, the mean square is $O_D(p^{-1})$ when $A \asymp p$.

Proof. Parseval gives $\sum_{\xi\neq 0} |\widehat{\mathbf{1}_C}(\xi)|^2 = p^2 n - n^2$, so $\mathfrak{s}_C$ is the mean nonzero spectral power. After subtracting the uniform-spectrum term, orthogonality in h gives

$$\mathbb{E}_h |R(h)|^2 \ll_D p^{-8} p^2 \sum_{\xi} |\widehat{\mathbf{1}_\Omega}(\xi)|^4.$$

The identity $\sum_\xi |\widehat{\mathbf{1}_\Omega}(\xi)|^4 = p^2 E_+(\Omega)$ proves the result. □

At $A \asymp p$, this theorem controls all but $O(p)$ of the p^2 displacements at constant error scale. The local block $\|h\| \ll \sqrt{p}$ also has $O(p)$ elements, so the theorem cannot by itself rule out the possibility that every locally relevant displacement is exceptional. Theorem 2.1 is precisely a local closure of this gap for an infinite family.

5.2 Exact orientation-averaged benchmark

For $L \in \mathrm{GL}_2(\mathbb{F}_p)$ define $w_L(x) = w(L^{-1}x)$.

Proposition 5.3 (Exact GL_2-averaged law)**.** *Let $S \subset \mathbb{F}_p^2$ have size n. Then*

$$\mathbb{E}_{L\in \mathrm{GL}_2(\mathbb{F}_p)} \mathrm{Var}(X_{w_L}) = \frac{(p^2 n - n^2)(p^2\|w\|_2^2 - |\sum w|^2)}{p^4(p^2-1)}.$$

More generally, for $h \neq 0$,

$$\mathbb{E}_{L\in \mathrm{GL}_2(\mathbb{F}_p)} \Gamma_{S,w_L}(Lh) = \frac{p^2 n - n^2}{p^4(p^2-1)} \sum_{\xi\neq 0} |\widehat{w}(\xi)|^2 e_p(-\xi\cdot h).$$

Proof. For each fixed nonzero frequency, $L^T\xi$ is uniformly distributed over $\mathbb{F}_p^2 \setminus \{0\}$. Average the exact Fourier formula and use Parseval. □

This benchmark shows exactly what complete spectral mixing would predict, but averaging over GL_2 destroys the fixed Archimedean orientation of the original problem.

6 Non-tangential truncation and algebraic exceptional structure

6.1 The non-tangential stage

For $1 \le L < p/4$, define

$$D_L = \{u \in \mathbb{F}_p^2 \setminus \{0\} : \|\langle u\rangle_p\|_\infty \le L\}.$$

Proposition 6.1 (Close-pair exceptional set)**.** *Let $C \subset \mathbb{F}_p^2$ be absolutely irreducible, non-linear, and of degree $e < p$. Let Ω have size A, and define*

$$\mathcal{E}_L = \{a : \exists P \neq Q \in C \cap (a + \Omega),\ Q - P \in D_L\}.$$

Then

$$|\mathcal{E}_L| \leq e^2 A |D_L| \ll_e AL^2.$$

Proof. For $u \neq 0$, C and $C - u$ are distinct; otherwise translation by u would force a full affine line orbit into C, contradicting irreducibility and non-linearity. Bezout therefore gives

$$r_C(u) = \#(C \cap (C - u)) \leq e^2.$$

A fixed ordered pair with difference u is placed in the same translated window in $g_\Omega(u) \leq A$ ways. Count witnesses and sum over $u \in D_L$. $\square$

Corollary 6.2 (Total-variation cost and packing)**.** *Define*

$$N^{\mathrm{nt}}_{\Omega,L}(a) = \begin{cases} N_\Omega(a), & a \notin \mathcal{E}_L, \\ 0, & a \in \mathcal{E}_L. \end{cases}$$

Then

$$d_{\mathrm{TV}}\big(\mathcal{L}(N_\Omega), \mathcal{L}(N^{\mathrm{nt}}_{\Omega,L})\big) \ll_e \frac{AL^2}{p^2}.$$

At $A \asymp p$ and $L = p^{1/2-\sigma}$, the cost is $O(p^{-2\sigma})$. If the lift of Ω lies in an $O(H)$-box, then on the non-tangential event

$$N_\Omega(a) \ll \left(1 + \frac{H}{L}\right)^2.$$

Proof. Couple the two variables by the same translation; they differ only on $\mathcal{E}_L$. The packing bound follows by subdividing the containing box into cells of diameter less than L. $\square$

This definition is stronger and more stable than a derivative-based tangency condition: every tangent arc, low-order contact, and microscopic algebraic cluster contains a close pair and is removed automatically.

6.2 A determinant certificate and Bombieri–Pila

Theorem 6.3 (Local integral lifting)**.** *Let $F \in \mathbb{F}_p[X, Y]$ be absolutely irreducible of degree $e \geq 2$. Let B be a box of side $L < p/2$, and put*

$$M_e = \binom{e+2}{2}, \qquad S_e = \frac{e(e+1)(e+2)}{3}.$$

Assume

$$M_e! L^{S_e} < p.$$

Then

$$\#(V(F)(\mathbb{F}_p) \cap B) \ll_{e,\varepsilon} L^{1/e+\varepsilon}.$$

More precisely, if the number of points exceeds $\max(M_e - 1, e^2)$*, all lifted points lie on a primitive, absolutely irreducible $G \in \mathbb{Z}[X, Y]$ of degree e whose reduction is a nonzero scalar multiple of the translated equation of F.*

Proof. Translate the lifted box into $[0,L]^2$. Form the point-by-monomial evaluation matrix using all monomials X^iY^j with $i+j\le e$. There are M_e columns. The coefficient vector of the translated equation F gives a nonzero column relation modulo p, so every $M_e\times M_e$ minor is divisible by p. A Leibniz expansion contains at most $M_e!$ terms, each of size at most

$$L^{\sum_{i+j\le e}(i+j)}=L^{S_e}.$$

The determinant is therefore both divisible by p and strictly smaller than p, hence zero over $\mathbb{Z}$. The matrix has rational rank below M_e, yielding a nonzero primitive $G\in\mathbb{Z}[X,Y]$ of degree at most e through all lifted points.

If there are more than e^2 points, Bezout forces F and $\overline{G}$ to have a common component. Absolute irreducibility of F and $\deg\overline{G}\le e$ imply $\overline{G}=cF$, so $\deg G=e$ and G is absolutely irreducible. The coefficient-uniform Bombieri–Pila theorem [2] then gives the stated integer-point bound. If there are only boundedly many points, enlarge the constant. □

Sharper positive-characteristic determinant and geometry-of-numbers methods extend the scale range; see [3, 11, 16]. The value of Theorem 6.3 is its explicit certificate: a dense exceptional microbox is not merely non-random but lies on an exact integral curve reducing to the original finite-field curve.

Corollary 6.4 (Two-scale structure–randomness decomposition)**.** *Let $A\asymp p$ and $H\asymp\sqrt{p}$. Choose*

$$L_{\mathrm{alg}}=p^{\theta},\qquad 0<\theta<1/S_e,$$

and

$$L_{\mathrm{nt}}=p^{1/2-\sigma},\qquad \sigma>0.$$

Then outside a set of translations of probability $O_e(p^{-2\sigma})$, all points in the critical window are L_{nt}-separated. On the exceptional set, every L_{alg}-microbox with more than $\max(M_e-1,e^2)$ points has an exact integral degree-e certificate and contains at most

$$O_{e,\varepsilon}(L_{\mathrm{alg}}^{1/e+\varepsilon})$$

points.

This is the rigorous version of the proposed architecture:

$$\boxed{\text{wave-dispersed part}\quad\cup\quad\text{mesoscopic close-pair part}\quad\cup\quad\text{microscopic algebraic part.}}$$

7 Unclosed ideas and precise next problems

7.1 The critical weak law

The complete limiting distribution is open even for the parabola. The beyond-Poisson results imply that an untruncated all-moment proof cannot work. A viable formulation is:

(i) choose $L=p^{1/2-\sigma}$ and study $N^{\mathrm{nt}}_{\Omega,L}$;

(ii) prove a weak law for the truncated process;

(iii) transfer the law to N_Ω using Corollary 6.2;

(iv) treat the algebraic tail separately as a moderate-deviation phenomenon.

Question 7.1 (Truncated critical law)**.** *For the parabola, a cubic graph, or an even monomial graph, does $N^{\mathrm{nt}}_{\Omega,L}$ converge to Poisson(λ) for a fixed regular body of area parameter λ? If not, identify the arithmetic limiting law.*

7.2 Separated higher correlations

For $k \geq 3$, define

$$R_{C,k}(u_2, \dots, u_k) = \#\{P : P, P + u_2, \dots, P + u_k \in C\}$$

and the k-fold overlap

$$G_{\Omega,k}(u_2, \dots, u_k) = |\Omega \cap (\Omega - u_2) \cap \cdots \cap (\Omega - u_k)|.$$

Then

$$\mathbb{E}_a(N_\Omega(a))_k = \frac{1}{p^2} \sum_{u_2,\dots,u_k} R_{C,k}(u_2, \dots, u_k) G_{\Omega,k}(u_2, \dots, u_k).$$

After the non-tangential truncation, the singular small diagonals are removed, but several coupled difference variables remain. The trilinear estimates in [7] are relevant, but do not immediately yield the full hierarchy required by a Poisson method.

Question 7.2 (Separated third correlation)**.** *For $C_d : y = x^d$ with even $d \geq 6$, prove a power-saving asymptotic for the third factorial moment after deleting $\mathcal{E}_L$, uniformly for $p^\varepsilon \ll L \ll p^{1/2-\varepsilon}$.*

7.3 General fixed polynomials and plane curves

For a general polynomial graph, the unresolved second-order error is

$$\mathcal{R}_f(h) = \sum_{a \neq 0, b} (r_f(a, b) - 1) g_\Omega((a, b) - h).$$

Bounding each a-slice separately discards the cancellation needed at the square-root scale. The correct target is a genuinely bilinear or two-parameter trace-function estimate for the full sum.

Question 7.3 (General polynomial bilinearization)**.** *For a fixed $f \in \mathbb{Z}[X]$, reorganize $\mathcal{R}_f(h)$ into bilinear trace kernels with a projective tensor norm controlled by $|\Omega|$, and prove $\mathcal{R}_f(h) = o(p^2)$ at $H \asymp \sqrt{p}$.*

For a general implicit curve, one would like a local Hasse-derivative expansion, followed by a wave-packet or Newton-polygon decomposition. The obstacle is to retain small-height coefficients while passing from a finite-field branch to an Archimedean-localized model.

7.4 Possible analytic mechanisms

The following tools remain plausible, but none is presently a complete proof.

(1) **Multilinear trace functions.** Preserve all difference variables rather than applying absolute values one slice at a time.

(2) **Localized restriction or decoupling.** Develop a finite-field restriction theorem with an additional Archimedean frequency cutoff $\|\xi\| \lesssim p/H$.

(3) **Wave packets and tubes.** Relevant if window orientations vary or tangent tubes are studied; fixed balls alone do not naturally create a Kakeya problem.

(4) **Higher p-adic divisibility.** Replace $p \mid \Delta$ by $p^r \mid \Delta$ in a determinant argument using Hasse derivatives or repeated clustering.

(5) **Newton support.** Select weighted monomials adapted to the Newton polygon and anisotropic windows, rather than all monomials of bounded total degree.

(6) **Structure–randomness iteration.** Use trace-function estimates on dispersed configurations and determinant methods only on density increments.

8 Literature map and publication assessment

8.1 Where the present results sit

The Archimedean antecedent is the Bombieri–Pila determinant philosophy [2], whose motivating $1/2$ exponent was discussed in Sarnak's 1988 manuscript [17]; later determinant methods include Heath-Brown [8]. The large finite-field window regime comes from Fourier completion and Bombieri–Weil square-root cancellation. Classical small-box work includes Chang–Cilleruelo–Garaev–Hernandez–Shparlinski–Zumalacarregui [3], Cilleruelo–Shparlinski [4], Kerr–Mohammadi [11], and Merai–Shparlinski [16]. These works use geometry of numbers, Vinogradov mean value estimates, and positive-characteristic analogues of Bombieri–Pila to control individual very small boxes.

Average and distributional results include Mak [12], Mak–Zaharescu [15, 13, 14], and the short trace-function literature [5, 6]. The present project differs in focusing on arbitrary Archimedean window shapes at the critical area $|\Omega| \asymp p$, local displacement covariance, and a combined spectral/non-tangential/algebraic decomposition.

8.2 Prospective publishable components

Primary candidate. Theorem 2.1 and Lemma 2.5 form the clearest short-paper package. They give a concrete arithmetic application of the root-count bilinear forms highlighted in [7, Section 9.9], reach the critical square-root scale, and retain arbitrary localized window shapes. A public literature search did not reveal the same translated-window covariance theorem, but this is not a proof of priority.

Secondary candidate. Theorem 3.1, Corollary 3.2, and Proposition 3.5 form a concise "beyond-Poisson" note. The arguments are elementary but the formulation is useful: a possible weak Poisson law coexists with polynomially rare algebraic clusters, factorial-moment divergence, and radically non-Poisson moderate deviations.

Supporting architecture. The close-pair truncation, determinant certificate, global spectral mean square, and exact GL_2 benchmark are probably not independent headline theorems. Their value is to make the full proof strategy rigorous and to isolate exactly where multilinear trace-function input is still missing.

8.3 Checks required before submission

Before any submission, the following are essential:

(i) an exhaustive MathSciNet/Zentralblatt and author-level priority search;

(ii) independent verification by a specialist of the gallant/light sheaf conventions, branch values, and bad-prime exclusions;

(iii) confirmation that the exact version of [7] used here applies uniformly to all dyadic pullbacks $T_{d,p}(\pm ba^{-d})$;

(iv) a line-by-line check of the small-torus support cutoff in Lemma 2.5;

(v) numerical experiments for fixed-count distributions after non-tangential truncation;

(vi) separation of journal-ready theorems from conjectural programmatic sections.

Acknowledgment of status

This document is a technical research note prepared to consolidate an ongoing investigation. It is not a claim of accepted novelty or peer-reviewed correctness. Results marked closed have complete arguments in the note subject to their cited inputs; the strongest result depends on a recent preprint. Open questions and speculative mechanisms are explicitly separated from theorems.

References

[1] E. Bombieri, *On exponential sums in finite fields*, Amer. J. Math. 88 (1966), 71–105.

[2] E. Bombieri and J. Pila, *The number of integral points on arcs and ovals*, Duke Math. J. 59 (1989), 337–357.

[3] M.-C. Chang, J. Cilleruelo, M. Z. Garaev, J. Hernandez, I. E. Shparlinski, and A. Zumalacarregui, *Points on curves in small boxes and applications*, Michigan Math. J. 63 (2014), 503–534.

[4] J. Cilleruelo and I. E. Shparlinski, *Concentration of points on curves in finite fields*, Monatsh. Math. 171 (2013), 315–327.

[5] E. Fouvry, E. Kowalski, and P. Michel, *The sliding-sum method for short exponential sums*, arXiv:1307.0135.

[6] E. Fouvry, E. Kowalski, P. Michel, C. S. Raju, J. Rivat, and K. Soundararajan, *On short sums of trace functions*, Ann. Inst. Fourier 67 (2017), no. 1, 423–449.

[7] E. Fouvry, E. Kowalski, P. Michel, and W. Sawin, *Bilinear forms with trace functions*, arXiv:2511.09459, version 3, 11 March 2026.

[8] D. R. Heath-Brown, *The density of rational points on curves and surfaces*, Ann. of Math. 155 (2002), 553–595.

[9] D. R. Heath-Brown and L. B. Pierce, *Burgess bounds for short mixed character sums*, J. Lond. Math. Soc. 91 (2015), no. 3, 693–708.

[10] N. M. Katz, *Exponential Sums and Differential Equations*, Annals of Mathematics Studies 124, Princeton University Press, 1990.

[11] B. Kerr and A. Mohammadi, *Points on polynomial curves in small boxes modulo an integer*, J. Number Theory 223 (2021), 64–78.

[12] K.-H. Mak, *The distribution of points on curves over finite fields in some small rectangles*, Monatsh. Math. 174 (2014), no. 4, 577–586.

[13] K.-H. Mak and A. Zaharescu, *Poisson type phenomena for points on hyperelliptic curves modulo p*, Funct. Approx. Comment. Math. 47 (2012), no. 1, 65–78.

[14] K.-H. Mak and A. Zaharescu, *On the distribution of the number of points on a family of curves over finite fields*, Int. J. Number Theory 8 (2012), no. 7, 1789–1811.

[15] K.-H. Mak and A. Zaharescu, *The distribution of rational points and polynomial maps on an affine variety over a finite field on average*, arXiv:1301.1359.

[16] L. Merai and I. E. Shparlinski, *Sparsity of curves and additive and multiplicative expansion of rational maps over finite fields*, Acta Arith. 188 (2019), 401–411.

[17] P. Sarnak, *Torsion points on varieties and homology of abelian covers*, unpublished manuscript, 1988.

[18] A. Weil, *On some exponential sums*, Proc. Natl. Acad. Sci. USA 34 (1948), 204–207.