Proceedings of the Royal Society of Edinburgh: Section A Mathematics, 1–12, *2026*
DOI:10.1017/prm.2026.10130

Incidence geometry and polynomial expansion over finite fields

Nuno Arala
Welfengarten 1, 30167 Hannover, Germany
(arala@math.uni-hannover.de)

Sam Chow
Mathematics Institute, Zeeman Building, University of Warwick, Coventry CV4 7AL, United Kingdom (sam.chow@warwick.ac.uk) (corresponding author)

(Received 14 February 2025; revised 22 January 2026; accepted 22 January 2026)

We use spectral theory and algebraic geometry to establish a higher-degree analogue of a Szemerédi–Trotter-type theorem over finite fields, with an application to polynomial expansion.

Keywords: algebraic varieties; finite fields; incidence geometry; polynomial expansion; spectral theory

2020 *Mathematics Subject Classification:* Primary 11T06; 11B30; 11G25; Secondary 51B05

1. Introduction

1.1. Incidence geometry

The Szemerédi–Trotter theorem asserts that if $\mathcal{P}$ is a set of points and $\mathcal{L}$ is a set of lines in $\mathbb{R}^2$ then

$$\#\{(\mathfrak{p},\mathfrak{l}) \in \mathcal{P}\times\mathcal{L} : \mathfrak{p}\in\mathfrak{l}\} \ll |\mathcal{P}|^{2/3}|\mathcal{L}|^{2/3} + |\mathcal{P}| + |\mathcal{L}|.$$

Such estimates have implications for the sum–product problem of Erdős and Szemerédi, see [6]. Vinh established the following finite-field analogue.

THEOREM 1.1 ([19]). *Let $\mathcal{P}$ be a set of points and $\mathcal{L}$ a set of lines in $\mathbb{F}_q^2$. Then*

$$\left|\#\{(\mathfrak{p},\mathfrak{l}) \in \mathcal{P}\times\mathcal{L} : \mathfrak{p}\in\mathfrak{l}\} - \frac{|\mathcal{P}|\cdot|\mathcal{L}|}{q}\right| \leqslant q^{1/2}\sqrt{|\mathcal{P}|\cdot|\mathcal{L}|}.$$

© The Author(s), 2026. Published by Cambridge University Press on behalf of The Royal Society of Edinburgh. This is an Open Access article, distributed under the terms of the Creative Commons Attribution licence (http://creativecommons.org/licenses/by/4.0), which permits unrestricted re-use, distribution and reproduction, provided the original article is properly cited.

https://doi.org/10.1017/prm.2026.10130 Published online by Cambridge University Press

In our previous paper, we bootstrapped Vinh's theorem to infer the following generalization to higher degrees.

THEOREM 1.2 ([2]). *Let $n \in \mathbb{N}$, and let $q > n$ be a prime power. Let $\mathcal{P}$ be a set of points in $\mathbb{F}_q^2$, and $\mathcal{Q}$ a collection of subsets of $\mathbb{F}_q^2$ of the shape*

$$\{(x, y) \in \mathbb{F}_q^2 : y = a_n x^n + \cdots + a_0\},$$

where $a_0, \ldots, a_n \in \mathbb{F}_q$. Then

$$\left| \#\{(\mathfrak{p}, \mathfrak{q}) \in \mathcal{P} \times \mathcal{Q} : \mathfrak{p} \in \mathfrak{q}\} - \frac{|\mathcal{P}| \cdot |\mathcal{Q}|}{q} \right| \leqslant q^{n/2} \sqrt{|\mathcal{P}| \cdot |\mathcal{Q}|}.$$

We had initially sought to generalize Vinh's result to sets

$$\{(x, y) \in \mathbb{F}_q^2 : y = (x + a)^n + b\}.$$

Although Theorem 1.2 could reasonably be considered to achieve this, its error term $q^{n/2}\sqrt{|\mathcal{P}| \cdot |\mathcal{Q}|}$ exceeds the main term $|\mathcal{P}| \cdot |\mathcal{Q}|/q$ in this context, if $n \geqslant 3$. We now present such a generalization in which the error term is smaller than the main term, when $|\mathcal{P}|, |\mathcal{Q}| \gg q^2$ and q is sufficiently large. Our results apply to symmetric polynomials, i.e.

$$F(a, x) = F(x, a).$$

THEOREM 1.3. *Let $F(a, x) \in \mathbb{F}_q[a, x]$ be a non-diagonal, symmetric polynomial such that $\deg(F) < \operatorname{char}(\mathbb{F}_q)$. Let $\mathcal{P}$ be a set of points and $\mathcal{Q}$ a collection of subsets of $\mathbb{F}_q^2$ of the shape*

$$\mathfrak{q}_{a,b} = \{(x, y) \in \mathbb{F}_q^2 : F(a, x) + b + y = 0\} \tag{1.1}$$

in $\mathbb{F}_q^2$. Then

$$\#\{(\mathfrak{p}, \mathfrak{q}) \in \mathcal{P} \times \mathcal{Q} : \mathfrak{p} \in \mathfrak{q}\} = \frac{|\mathcal{P}| \cdot |\mathcal{Q}|}{q} + O_{\deg(F)}(q^{5/6}\sqrt{|\mathcal{P}| \cdot |\mathcal{Q}|}).$$

1.2. Expanding polynomials in finite fields

For a polynomial $P \in \mathbb{F}_q[x_1, \ldots, x_k]$ and sets $A_1, \ldots, A_k \subseteq \mathbb{F}_q$, we would usually expect

$$P(A_1, \ldots, A_k) := \{P(x_1, \ldots, x_k) : x_1 \in A_1, \ldots, x_k \in A_k\}$$

to be much larger than $A_1, \ldots, A_k$, i.e. P should *expand*. There are some algebraic obstructions, e.g. F could be additively or multiplicatively structured, or some A_i could be dense. The *Elekes–Rónyai problem* [5, 7] is to demonstrate expansion away from such obstructions. The problem is also studied over more general fields, and is closely related to the sum–product problem.

In order to quantify the amount of expansion, Tao [18] proposed the following hierarchy.

https://doi.org/10.1017/prm.2026.10130 Published online by Cambridge University Press

(i) *Weak asymmetric expansion*: there exist $c, C > 0$ such that

$$|P(A_1, \dots, A_k)| \geq C^{-1} \min\{|A_1|, \dots, |A_k|\}^{1-c} q^c$$

whenever $|A_1|, \dots, |A_k| \geq Cq^{1-c}$.

(ii) *Moderate asymmetric expansion*: there exist $c, C > 0$ such that

$$|P(A_1, \dots, A_k)| \geq C^{-1} q$$

whenever $|A_1|, \dots, |A_k| \geq Cq^{1-c}$.

(iii) *Almost strong asymmetric expansion*: there exist $c, C > 0$ such that

$$|P(A_1, \dots, A_k)| \geq q - Cq^{1-c}$$

whenever $|A_1|, \dots, |A_k| \geq Cq^{1-c}$.

(iv) *Strong asymmetric expansion*: there exist $c, C > 0$ such that

$$|P(A_1, \dots, A_k)| \geq q - C$$

whenever $|A_1|, \dots, |A_k| \geq Cq^{1-c}$.

(v) *Very strong asymmetric expansion*: there exist $c, C > 0$ such that

$$P(A_1, \dots, A_k) = \mathbb{F}_q$$

whenever $|A_1|, \dots, |A_k| \geq Cq^{1-c}$.

The qualifier 'asymmetric' in the above refers to the fact that the sets $A_1, \dots, A_k$ are allowed to differ. Typically, q is allowed to vary, and P is allowed to vary in some family, meaning that c, C do not depend on them.

In our previous paper, we established almost strong expansion for a Zariski-dense set of polynomials of degree d in $d+1$ variables, strengthening and generalizing the main result of [13]. We are now able to handle ternary polynomials of more general degree.

THEOREM 1.4. *Let $P \in \mathbb{F}_q[x, y, z]$ be a polynomial of the form*

$$F(x, G(y, z)) + H(y, z) + J(x),$$

where $F, G, H \in \mathbb{F}_q[x, y]$ and $J \in \mathbb{F}_q[x]$ are such that:

(i) *F is non-diagonal and symmetric with $\deg(F) < \mathrm{char}(\mathbb{F}_q)$, and*
(ii) *G and H are algebraically independent.*

https://doi.org/10.1017/prm.2026.10130 Published online by Cambridge University Press

Then there exists a constant $C = C_{\deg(P)}$ *such that if* $X, Y, Z \subseteq \mathbb{F}_q$ *and* $|X|, |Y|, |Z| \geq Cq^{1/2}$ *then*

$$|P(X,Y,Z)| = q - O_{\deg(P)}\left(\frac{q^{11/3}}{|X| \cdot |Y| \cdot |Z|}\right).$$

COROLLARY 1.5. *Let* $\varepsilon > 0$*, let* P *be a polynomial of the form given in Theorem 1.4, and let* $X, Y, Z \subseteq \mathbb{F}_q$ *be such that* $|X|, |Y|, |Z| \gg q^{8/9+\varepsilon}$*. Then*

$$|P(X,Y,Z)| = q - O_{\varepsilon,\deg(P)}(q^{1-3\varepsilon}).$$

COROLLARY 1.6. *Let* $\varepsilon > 0$*, let* $P(x,y,z) = (x-y)^k + z$*, where* $k \geq 2$ *is an integer, and let* q *be a power of a prime* $p > k$*. Let* $X, Y, Z \subseteq \mathbb{F}_q$ *be such that* $|X|, |Y|, |Z| \gg q^{8/9+\varepsilon}$*. Then*

$$|P(X,Y,Z)| = q - O_{\varepsilon,k}(q^{1-3\varepsilon}).$$

A great deal of work has been done on binary expansion [3, 8, 10–12]. The Elekes–Rónyai problem was essentially solved for weak binary expansion by Bukh and Tsimerman [4], and for moderate expansion by Tao [18]. Therein, Tao also obtained a general result for almost strong binary expansion, but with a technical condition that is difficult to check in practice.

These developments led us to consider the Elekes–Rónyai problem for ternary polynomials over finite fields. Aside from the aforementioned results, Shkredov [16] established very strong expansion for x^2+xy+z. As for higher-degree ternary polynomials, we were previously successful at demonstrating almost strong expansion [2, Theorem 1.6] only in the special case

$$P(x,y,z) = xG(y,z) + H(y,z) + ax^k,$$

where G and H are algebraically independent. Theorem 1.4 is more general. For example, it handles

$$P(x,y,z) = (x-y)^k + z,$$

whereas our previous approach does not reach such polynomials.

1.3. Methods

We will see that Theorem 1.4 is easy to deduce from Theorem 1.3. To prove Theorem 1.3, we use Vinh's spectral framework [19], which begins by realizing incidences as edges of a regular graph. The main task is to bound the second largest eigenvalue, in absolute value. Vinh is able to use the square of an adjacency matrix to swiftly achieve this in the case of point–line incidences. In our higher-degree setting, we need to cube the matrix, giving rise to a family of algebraic curves on

https://doi.org/10.1017/prm.2026.10130 Published online by Cambridge University Press

which we need to count $\mathbb{F}_q$ points. The Weil bound is suitable for this purpose, as we are able to show that almost all curves in the family are absolutely irreducible.

Organization. We will prove Theorem 1.3 in § 2. Then, in § 3, we will deduce Theorem 1.4.

Notation. We use the Vinogradov notations $\ll$ and $\gg$, as well as the Bachmann–Landau notation $O(\cdot)$. Let f and g be complex-valued functions. We write $f \ll g$ or $f = O(g)$ if $|f| \leq C|g|$ pointwise, for some constant $C > 0$. We use a subscript within these notations to indicate possible dependence on the implied constant C.

2. An incidence bound

In this section, we prove Theorem 1.3. The first subsection will establish the result assuming a lemma about absolute irreducibility. The latter will then be demonstrated in the second subsection.

2.1. Spectral theory

Definition 2.1. *Let $\mathcal{G}$ be a graph with n vertices. We denote the eigenvalues of the adjacency matrix of $\mathcal{G}$ by $\lambda_1(\mathcal{G}), \dots, \lambda_n(\mathcal{G})$, ordered in such a way that*

$$|\lambda_1(\mathcal{G})| \geq \cdots \geq |\lambda_n(\mathcal{G})|.$$

Definition 2.2. *For a symmetric polynomial $F \in \mathbb{F}_q[x, y]$, we define the graph $\mathcal{G}_F$ as the graph with set of vertices $\mathbb{F}_q^2$ and an edge between (a, b) and (x, y) if and only if*

$$F(a, x) + b + y = 0.$$

We will prove the following spectral bound for the graphs $\mathcal{G}_F$.

Lemma 2.3. *Let $F \in \mathbb{F}_q[x, y]$ be a non-diagonal, symmetric polynomial such that $\deg(F) < \operatorname{char}(\mathbb{F}_q)$. Then $\lambda_2(\mathcal{G}_F) \ll_{\deg(F)} q^{5/6}$.*

Proof [Proof of Lemma 2.3 assuming Lemma 2.4]. Let A be the adjacency matrix of $\mathcal{G} = \mathcal{G}_F$, the rows and columns of which are indexed by points in $\mathbb{F}_q^2$. We consider the matrix A^3.

The entry of A^3 in the row corresponding to (a, b) and the column corresponding to (c, d) is the number of paths of length 3 in $\mathcal{G}$ connecting (a, b) to (c, d). This is the number of solutions $(x_1, x_2, y_1, y_2) \in \mathbb{F}_q^4$ to the system of equations

$$\begin{cases} F(a, x_1) + b + y_1 = 0 \\ F(x_1, x_2) + y_1 + y_2 = 0 \\ F(x_2, c) + y_2 + d = 0. \end{cases}$$

https://doi.org/10.1017/prm.2026.10130 Published online by Cambridge University Press

Solutions to this system are in bijection with pairs $(x_1, x_2) \in \mathbb{F}_q^2$ satisfying

$$F(a, x_1) + F(x_2, c) - F(x_1, x_2) + b + d = 0, \tag{2.1}$$

with mutually inverse bijections being given by

$$(x_1, x_2, y_1, y_2) \mapsto (x_1, x_2)$$

and

$$(x_1, x_2) \mapsto (x_1, x_2, -F(a, x_1) - b, -F(x_2, c) - d).$$

Let $C_{\mathbf{a}} = C/\mathbb{F}_q$ be the affine plane curve defined in coordinates (x_1, x_2) by (2.1), where $\mathbf{a} = (a, b, c, d)$, so that, as argued above,

$$A^3_{(a,b),(c,d)} = \#C(\mathbb{F}_q).$$

Observe that (2.1) indeed defines a curve; the polynomial on the left-hand side of (2.1) does not vanish, since F is assumed not to be diagonal.

By the Weil bound,

$$\#C(\mathbb{F}_q) = \begin{cases} q + O(\sqrt{q}) & \text{if C is absolutely irreducible} \\ O(q) & \text{otherwise.} \end{cases}$$

Denote by U the $q^2 \times q^2$ matrix all the entries of which equal 1. The upshot is that

$$A^3 = qU + M,$$

where M is a $q^2 \times q^2$ real matrix, the entries of which are indexed by points $\mathbf{a} \in \mathbb{F}_q^4$, such that

$$M_{\mathbf{a}} = \begin{cases} O(\sqrt{q}) & \text{if } C_{\mathbf{a}} \text{ is absolutely irreducible} \\ O(q) & \text{otherwise.} \end{cases}$$

Lemma 2.4, which we will prove in § 2.2, asserts that the points $\mathbf{a}$ such that $C_{\mathbf{a}}$ is not absolutely irreducible are contained in a proper subvariety $Z \subset \mathbb{A}^4$.

We have $\lambda_1(\mathcal{G}) = q$, since $\mathcal{G}$ is clearly q-regular. Note that the eigenspace corresponding to the eigenvalue 1 is spanned by the vector $\mathbf{u} = (1, \ldots, 1)$. Since A is symmetric, the spectral theorem implies that

$$|\lambda_2(\mathcal{G})|^3 = \max\{\lambda \in \mathbb{R} : \|A^3\mathbf{x}\|_2 \leq \lambda\|\mathbf{x}\|_2 \text{ for any } \mathbf{x} \text{ with } \mathbf{x} \cdot \mathbf{u} = 0\}.$$

For $\mathbf{x}$ with $\mathbf{x} \cdot \mathbf{u} = 0$, we have $U\mathbf{x} = \mathbf{0}$ and hence

$$A^3\mathbf{x} = (qU + M)\mathbf{x} = M\mathbf{x}.$$

Now

$$\|A^3\mathbf{x}\|_2 = \|M\mathbf{x}\|_2 \leq \|M\|_2\|\mathbf{x}\|_2,$$

https://doi.org/10.1017/prm.2026.10130 Published online by Cambridge University Press

by Cauchy-Schwarz, where

$$\|M\|_2 = \sqrt{\sum_{\mathbf{a}\in\mathbb{F}_q^4} M_{\mathbf{a}}^2}.$$

It follows that

$$|\lambda_2(\mathcal{G})| \leq \left(\sum_{\mathbf{a}\in\mathbb{F}_q^4} M_{\mathbf{a}}^2\right)^{1/6}. \tag{2.2}$$

Using that $\#Z(\mathbb{F}_q) = O(q^3)$, since Z is a proper subvariety of $\mathbb{A}^4$,

$$\begin{aligned}\sum_{\mathbf{a}\in\mathbb{F}_q^4} M_{\mathbf{a}}^2 &= \sum_{\mathbf{a}\in Z(\mathbb{F}_q)} M_{\mathbf{a}}^2 + \sum_{\mathbf{a}\notin Z(\mathbb{F}_q)} M_{\mathbf{a}}^2 \\ &= O(q^3\cdot q^2) + O(q^4\cdot q) = O(q^5).\end{aligned}$$

Inserting this into (2.2) completes the proof. □

Proof of Theorem 1.3 assuming Lemma 2.3. The result follows immediately by applying the expander mixing lemma [1, Corollary 9.2.5] to the graph $\mathcal{G}_F$, together with Lemma 2.3. □

Recall that we proved Lemma 2.3 subject to Lemma 2.4. We now turn to the latter.

2.2. Absolute irreducibility

In this subsection, we establish the following lemma used in the proof of Lemma 2.3.

LEMMA 2.4. *Let $F \in \mathbb{F}_q[x,y]$ be a non-diagonal, symmetric polynomial such that $\deg(F) < \mathrm{char}(\mathbb{F}_q)$. Then there exists a proper subvariety $Z \subset \mathbb{A}^4$ of degree $O_{\deg(F)}(1)$ such that, for $\mathbf{a} = (a,b,c,d) \notin Z(\mathbb{F}_q)$, the curve $C_{\mathbf{a}}$ defined by (2.1) is absolutely irreducible.*

To this end, we will require the following auxiliary result, the proof of which we defer to Appendix A.

LEMMA 2.5. *Let k be a field and let $P,Q,R,S \in k[x]$ be polynomials such that $0 < \deg(P) = \deg(R) < \mathrm{char}(k)$ and the compositions $P\circ Q$ and $R\circ S$ are equal. Then there exists a linear polynomial $L \neq 0$ such that $S = L\circ Q$.*

We will also require the following characterization of additive polynomials [9, Chapter 1].

https://doi.org/10.1017/prm.2026.10130 Published online by Cambridge University Press

LEMMA 2.6. *Let k be an infinite field of finite characteristic p, and let $P \in k[X]$ satisfy*

$$P(x+y) = P(x) + P(y) \qquad (x, y \in k).$$

Then

$$P = \sum_{j=0}^{K} a_j X^{p^j}$$

for some $a_0, \ldots, a_K \in k$.

Proof [Proof of Lemma 2.4]. We begin by observing that the reducibility of $C_{\mathbf{a}}$ can be characterized in terms of polynomial equations in the coefficients of the polynomial defining $C_{\mathbf{a}}$, see [15, Chapter V, Theorem 2A]. Thus, there exists a subvariety $Z \subset \mathbb{A}^4$ of degree $O_{\deg(F)}(1)$ such that $C_{\mathbf{a}}$ is absolutely irreducible if and only if $\mathbf{a} \notin Z$. It remains to prove that Z is proper, in other words, that $C_{\mathbf{a}}$ is irreducible over $\overline{\mathbb{F}_q}$ for some $\mathbf{a} \in \overline{\mathbb{F}_q}^4$. Henceforth, we assume for a contradiction that $C_{\mathbf{a}}$ is reducible over $\overline{\mathbb{F}_q}$ for every $\mathbf{a} \in \overline{\mathbb{F}_q}^4$.

Our assumption implies that the polynomial

$$F(a, x_1) + F(x_2, c) - F(x_1, x_2) + t$$

is reducible for any $a, c, t \in \overline{\mathbb{F}_q}$. By [14, §3.3, Corollary 1], this forces the polynomial $F(a, x_1) + F(x_2, c) - F(x_1, x_2)$ to be composite—namely it has the form $P(Q(x_1, x_2))$ with $\deg(P) \geq 2$—for any $a, c \in \overline{\mathbb{F}_q}$. Specializing $a = 0$, we deduce that the polynomial

$$F(0, x_1) + F(x_2, c) - F(x_1, x_2) \in \mathbb{F}_q(c)[x_1, x_2]$$

has the form $P(Q(x_1, x_2))$, for some polynomials P and Q with coefficients in $\overline{\mathbb{F}_q(c)}$ and $\deg(P) \geq 2$. Note that the coefficients of P and Q lie in some finite extension of $\mathbb{F}_q(c)$. Furthermore, any such extension is the function field of a curve $X/\mathbb{F}_q$, see [17, Chapter II, Remark 2.5]. This gives rise to a polynomial identity

$$F(0, x_1) + F(x_2, \pi(p)) - F(x_1, x_2) = P_p(Q_p(x_1, x_2)), \tag{2.3}$$

where $\pi : X \to \mathbb{A}^1$ is a non-constant morphism, and P_p, Q_p are polynomials whose coefficients are algebraically parametrized by a point $p \in X$.

For $p, p' \in X$ with

$$F(x, \pi(p)) \neq F(x, \pi(p')) \in \overline{\mathbb{F}_q}[x],$$

the above yields

$$F(x_2, \pi(p)) - F(x_2, \pi(p')) = P_p(Q_p(x_1, x_2)) - P_{p'}(Q_{p'}(x_1, x_2)). \tag{2.4}$$

If we fix p, p' and regard the above as an identity of polynomials in the variable x_1 with coefficients in $\overline{\mathbb{F}_q}(x_2)$, then we have the equality of polynomial compositions

$$(P_{p'} + F(x_2, \pi(p)) - F(x_2, \pi(p'))) \circ Q_{p'}(\cdot, x_2) = P_p \circ Q_p(\cdot, x_2).$$

https://doi.org/10.1017/prm.2026.10130 Published online by Cambridge University Press

For generic $p, p' \in X$, the degrees in x_1 of

$$P_p, \qquad P_{p'} + F(x_2, \pi(p)) - F(x_2, \pi(p'))$$

are equal. Now Lemma 2.5 implies that for generic—and hence for all—$p, p' \in X$ the polynomials $Q_{p'}(\cdot, x_2)$ and $Q_p(\cdot, x_2)$ are related by left-composition with a linear polynomial with coefficients in $\overline{\mathbb{F}_q}(x_2)$. In other words, the family of polynomials $Q_p(x_1, x_2)$ has the form

$$Q_p(x_1, x_2) = f_p(x_2)h(x_1, x_2) + g_p(x_2),$$

for some families of polynomials $f_p, g_p \in \overline{\mathbb{F}_q}[x_2]$ with $f_p \neq 0$, and some $h \in \mathbb{F}_q[x_1, x_2]$. We assume as we may that the polynomials f_p do not have a common factor, for such a factor can be absorbed into h.

Inserting the above into (2.4) yields

$$\begin{aligned} &F(x_2, \pi(p)) - F(x_2, \pi(p')) \\ &= P_p(f_p(x_2)h(x_1, x_2) + g_p(x_2)) - P_{p'}(f_{p'}(x_2)h(x_1, x_2) + g_{p'}(x_2)). \end{aligned}$$

Since h is non-constant—as otherwise the left-hand side of (2.3) would not depend on x_1, which is clearly a contradiction—we obtain the polynomial identity

$$F(x_2, \pi(p)) - F(x_2, \pi(p')) = P_p(f_p(x_2)z + g_p(x_2)) - P_{p'}(f_{p'}(x_2)z + g_{p'}(x_2)).$$

Let $m \geq 2$ be the generic degree of P_p, and let a_p be its leading coefficient. Comparing leading coefficients in z yields

$$a_p f_p(x_2)^m = a_{p'} f_{p'}(x_2)^m,$$

which implies that any two f_p and $f_{p'}$ are related by multiplication by a root of unity. Since the f_p were assumed not to share a common factor, this implies that $(f_p)_p$ is a family of constant polynomials. We thus obtain

$$P_p\left(f_p z + g_p(x_2)\right) - P_{p'}(f_{p'} z + g_{p'}(x_2)) = F(x_2, \pi(p)) - F(x_2, \pi(p')).$$

We now make the variable change $w = f_{p'} z + g_{p'}(x_2)$ and observe that

$$P_p\left(\frac{f_p}{f_{p'}}(w - g_{p'}(x_2))\right) = P_{p'}(w) + F(x_2, \pi(p)) - F(x_2, \pi(p')).$$

The right-hand side is a polynomial in x_2 and w where no monomial involving both x_2 and w shows up with a non-zero coefficient. On the other hand, the left-hand side is of the form

$$P_p(aw + \ell(x_2)),$$

where ℓ is non-constant. This expands as

$$P_p(aw) + P_p(\ell(x_2)) + C + M(w, x_2)$$

https://doi.org/10.1017/prm.2026.10130 Published online by Cambridge University Press

for some constant C, where $M(w, x_2)$ is a linear combination of monomials involving both w and x_2. We must have $M(w, x_2) = 0$, so P_p satisfies the identity

$$P_p(x + y) = P_p(x) + P_p(y) + C.$$

The upshot is that the polynomial $P_p + C$ satisfies the assumption of Lemma 2.6. Therefore, either $\deg(P_p) = 1$ or $\deg(P_p) \geqslant \mathrm{char}(\mathbb{F}_q)$. The latter would contradict our assumption on the degree of F, in light of (2.3). Therefore, P_p has degree 1, which contradicts our assumption that it has degree at least 2. □

3. Expanding polynomials

In this section, we prove Theorem 1.4. We begin by showing that

$$\#\{(G(y, z), H(y, z)) : (y, z) \in Y \times Z\} \gg |Y| \cdot |Z|.$$

To this end, we argue as follows. Identify $\mathbb{F}_q^2$ with affine space $\mathbb{A}^2_{\mathbb{F}_q}$. The condition that G and H are algebraically independent implies that the map φ given by $\varphi(y, z) = (G(y, z), H(y, z))$ is dominant, i.e. its image is not contained in a proper subvariety of $\mathbb{A}^2_{\mathbb{F}_q}$. Therefore, if we denote by $k = O(1)$ the degree of φ, there exists a proper subvariety $\mathcal{X} \subsetneq \mathbb{A}^2_{\mathbb{F}_q}$, of degree $O(1)$ by [2, Lemma 3.1], such that $\varphi^{-1}(\mathbf{a})$ is finite whenever $\mathbf{a} \notin \mathcal{X}$. Since φ is at most k-to-1 outside $\varphi^{-1}(\mathcal{X})$, it follows that

$$\begin{aligned}
|\varphi(Y \times Z)| &\geq |\varphi((\mathbb{A}^2_{\mathbb{F}_q} \setminus \varphi^{-1}(\mathcal{X}))(\mathbb{F}_q) \cap (Y \times Z))| \\
&\geq \frac{1}{k} |(\mathbb{A}^2_{\mathbb{F}_q} \setminus \varphi^{-1}(\mathcal{X}))(\mathbb{F}_q) \cap (Y \times Z)| \\
&\geq \frac{1}{k} (|Y| \cdot |Z| - |\varphi^{-1}(\mathcal{X})(\mathbb{F}_q)|) \\
&= \frac{1}{k} (|Y| \cdot |Z| + O(q)).
\end{aligned}$$

In the final line, we used that $\varphi^{-1}(\mathcal{X})$, being a proper subvariety of $\mathbb{A}^2_{\mathbb{F}_q}$, has dimension at most 1, and its degree is bounded in terms of the degrees of $\mathcal{X}$ and φ. It is now clear that, for sufficiently large C, if $|Y|, |Z| \geq Cq^{1/2}$, then the above is at least $c|Y| \cdot |Z|$ for a suitably small constant $c > 0$.

Recall (1.1), and consider the set

$$\mathcal{Q} = \{\mathfrak{q}_{G(y,z),H(y,z)} : (y, z) \in Y \times Z\}.$$

We just proved that $|\mathcal{Q}| \gg |Y| \cdot |Z|$. Define also $W = \mathbb{F}_q \setminus P(X, Y, Z)$, and set

$$\mathcal{P} = \{(x, J(x) - w) : x \in X, w \in W\},$$

noting that $|\mathcal{P}| = |X| \cdot |W|$. It follows from the definitions of $\mathcal{P}$, $\mathcal{Q}$, P and W that if $\mathfrak{p} \in \mathcal{P}$ and $\mathfrak{q} \in \mathcal{Q}$ then $\mathfrak{p} \notin \mathfrak{q}$. Therefore, Theorem 1.3 yields

$$\frac{|\mathcal{P}| \cdot |\mathcal{Q}|}{q} = O(q^{5/6} \sqrt{|\mathcal{P}| \cdot |\mathcal{Q}|}),$$

https://doi.org/10.1017/prm.2026.10130 Published online by Cambridge University Press

whence

$$|\mathcal{P}| \cdot |\mathcal{Q}| \ll q^{11/3}.$$

On using that

$$|\mathcal{Q}| \gg |Y| \cdot |Z|, \quad |\mathcal{P}| = |X| \cdot |W|, \quad |W| = q - |P(X, Y, Z)|,$$

the result follows.

Acknowledgement
We thank an anonymous referee for a careful reading and for helpful comments.

Funding
NA was funded through the Engineering and Physical Sciences Research Council Doctoral Training Partnership at the University of Warwick.

Rights
For the purpose of open access, the authors have applied a Creative Commons Attribution (CC-BY) licence to any Author Accepted Manuscript version arising from this submission.

References

1. N. Alon and J. H. Spencer. *The probabilistic method. fourth edition, Wiley Ser. Discrete Math. Optim.*, fourth edition, (John Wiley & Sons, Inc, Hoboken, NJ, 2016).
2. N. Arala and S. Chow. Expansion properties of polynomials over finite fields. *Finite Fields Appl.* **108** (2025), 102687, 14.
3. J. Bourgain. More on the sum-product phenomenon in prime fields and its applications. *Int. J. Number Theory*. **1** (2005), 1–32.
4. B. Bukh and J. Tsimerman. Sum–product estimates for rational functions. *Proc. Lond. Math. Soc. (3)* **104** (2012), 1–126.
5. F. de Zeeuw. *A survey of Elekes-Rónyai-type problems*, New trends in intuitive geometry, 95–124. Bolyai Soc. Math. Stud. Vol.**27**, (János Bolyai Mathematical Society, Budapest, 2018).
6. G. Elekes. On the number of sums and products. *Acta Arith.* **81** (1997), 365–367.
7. G. Elekes and L. Rónyai. A combinatorial problem on polynomials and rational functions. *J. Comb. Theory Ser. A* **89** (2000), 1–20.
8. M. Garaev and C.-Y. Shen. On the size of the set $A(A+1)$. *Math. Z.* **265** (2010), 125–132.
9. D. Goss. *Basic structures of function field arithmetic* (Springer–Verlag, Berlin, 1996).
10. D. Hart, L. Li and C.-Y. Shen. Fourier analysis and expanding phenomena in finite fields. *Proc. Amer. Math. Soc.* **141** (2013), 461–473.
11. N. Hegyvári and F. Hennecart. Explicit constructions of extractors and expanders. *Acta Arith.* **140** (2009), 233–249.
12. A. Mohammadi. On growth of the set $A(A+1)$ in arbitrary finite fields. *Electron. J. Combin.* **27** (2020), #P 4.3.
13. T. Pham, L. A. Vinh and F. de Zeeuw. Three-variable expanding polynomials and higher-dimensional distinct distances. *Combinatorica*. **39** (2019), 411–426.
14. A. Schinzel. *Polynomials with special regard to reducibility* (Cambridge University Press, Cambridge, 2000).
15. W. M. Schmidt. *Equations over finite fields: an elementary approach. Lecture Notes in Mathematics*, Vol.**536** (Springer–Verlag, Berlin–New York, 1976).
16. I Shkredov. On monochromatic solutions of some nonlinear equations in $\mathbb{Z}/p\mathbb{Z}$. *Mat. Zametki*. **88** (2010), 625–634.

https://doi.org/10.1017/prm.2026.10130 Published online by Cambridge University Press

17. J. H. Silverman. *The arithmetic of elliptic curves*, second edition, Grad. Texts in Math. Vol.**106**, (Springer, Dordrecht, 2009).
18. T. Tao. Expanding polynomials over finite fields of large characteristic, and a regularity lemma for definable sets. *Contrib. Discrete Math.* **10** (2015), 22–98.
19. L. A. Vinh. The Szemerédi-Trotter type theorem and the sum-product estimate in finite fields. *European J. Combin.* **32** (2011), 1177–1181.

Appendix A. On polynomial composition

Here, we prove Lemma 2.5. A variant of this featured as Problem 5 at the 14th edition of the Romanian Masters of Mathematics, and the proof that we present here is adapted from the official solutions to that problem.

Proof of Lemma 2.5. Since $P \circ Q = R \circ S$, one has

$$\deg(P)\deg(Q) = \deg(R)\deg(S)$$

which, since $\deg(P) = \deg(R)$, implies that $\deg(Q) = \deg(S)$. Let d be this common degree. If α and β denote the leading coefficients of Q and S, respectively, and $c = \alpha\beta^{-1}$, then $S - cQ$ is a polynomial of degree smaller than d. If this polynomial is constant, we are done. So we assume henceforth that it is non-constant, i.e. that $S = cQ + M$ where M is a polynomial with $0 < \deg(M) < d$.

We then have

$$R(cQ(x) + M(x)) = P(Q(x)),$$

and the left-hand side expands as a linear combination of products of powers of $Q(x)$ and $M(x)$. Among the terms that contain positive powers of $M(x)$, the one with the largest degree is $\gamma r c^{r-1} Q(x)^{r-1} M(x)$, where r is the common degree of R and P and γ is the leading coefficient of R. Recall that $0 < r < \mathrm{char}(k)$. Therefore, $rc^{r-1}Q^{r-1}M$ differs from $T \circ Q$ by a polynomial of lower degree, for some polynomial T, and in particular

$$\deg(Q^{r-1}M) = \deg(T \circ Q).$$

The degree of the right-hand side is divisible by d, whereas the degree of the left-hand side is $d(r-1) + \deg(M)$, which is not divisible by d because $0 < \deg(M) < d$. This contradiction completes the proof of the lemma. □

https://doi.org/10.1017/prm.2026.10130 Published online by Cambridge University Press