Mathematical Proceedings of the Cambridge Philosophical Society

VOL. 180 MAY 2026 PART 3

Math. Proc. Camb. Phil. Soc. (2026), **180**, 477–489 on behalf of The Cambridge Philosophical Society.
© The Author(s), 2025. Published by Cambridge University Press

doi: 10.1017/S0305004125101771

First published online 27 November 2025

A note on transverse sets and bilinear varieties

BY LUKA MILIĆEVIĆ

Mathematical Institute of the Serbian Academy of Sciences and Arts, Kneza Mihaila 36, 11000 *Belgrade, Serbia.*

e-mail: luka.milicevic@turing.mi.sanu.ac.rs

(*Received* 01 *September* 2023*; revised* 09 *July* 2025*; accepted* 12 *July* 2025)

Abstract

Let G and H be finite-dimensional vector spaces over $\mathbb{F}_p$. A subset $A \subseteq G \times H$ is said to be *transverse* if all of its rows $\{x \in G\colon (x, y) \in A\}$, $y \in H$, are subspaces of G and all of its columns $\{y \in H\colon (x, y) \in A\}$, $x \in G$, are subspaces of H. As a corollary of a bilinear version of the Bogolyubov argument, Gowers and the author proved that dense transverse sets contain bilinear varieties of bounded codimension. In this paper, we provide a direct combinatorial proof of this fact. In particular, we improve the bounds and evade the use of Fourier analysis and Freiman's theorem and its variants.

2020 Mathematics Subject Classification: 05E18 (Primary); 11T06 (Secondary)

1. *Introduction*

In this paper, we fix a prime field $\mathbb{F}_p$ and consider finite-dimensional vector spaces over $\mathbb{F}_p$, usually denoted by G or H. A classical result, due to Bogolyubov, adapted to the finite vector space setting, reads as follows.

THEOREM 1·1 (Bogolyubov argument [4]). *Suppose that G is a finite-dimensional vector space over $\mathbb{F}_p$. Let $A \subseteq G$ be a set of size at least $\delta|G|$. Then $2A - 2A = \{a_1 + a_2 - a_3 - a_4\colon a_1, a_2, a_3, a_4 \in A\}$ contains a subspace of G of codimension $O_\delta(1)$.*

Motivated by applications in number theory [2] and higher order Fourier analysis [6], Bienvenu and Lê [1], and Gowers and the author [5], respectively, independently proved a

https://doi.org/10.1017/S0305004125101771 Published online by Cambridge University Press

bilinear generalisation of the Bogolyubov argument. It states that if we are given a dense set A inside $G \times H$, then replacing rows or columns by their difference sets several times, we may find bilinear structure in the resulting set. To be precise, let us write $\mathsf{D}_{\mathsf{hor}}A$ for the *horizontal difference set* of A which is defined as $\mathsf{D}_{\mathsf{hor}}A = \{(x_1 - x_2, y) \colon (x_1, y), (x_2, y) \in A\}$. Similarly, we write $\mathsf{D}_{\mathsf{ver}}A$ for the *vertical difference set* of A which is defined as $\mathsf{D}_{\mathsf{ver}}A = \{(x, y_1 - y_2) \colon (x, y_1), (x, y_2) \in A\}$.

THEOREM 1·2 (Bilinear Bogolyubov argument [**1, 5**]). *Let $A \subseteq G \times H$ be a set of size at least $\delta|G||H|$. Then*

$$\mathsf{D}_{\mathsf{ver}}\mathsf{D}_{\mathsf{ver}}\mathsf{D}_{\mathsf{hor}}\mathsf{D}_{\mathsf{hor}}\mathsf{D}_{\mathsf{ver}}A$$

contains a set of the form $\{(x, y) \in U \times V \colon \beta(x, y) = 0\}$, where $U \leq G$ and $V \leq H$ are subspaces of codimensions $O_\delta(1)$ and $\beta \colon G \times H \to \mathbb{F}_p^r$ is a bilinear map, where $r = O_\delta(1)$.

We refer to the sets of the form $\{(x, y) \in U \times V \colon \beta(x, y) = 0\}$, like the ones appearing in the conclusion above, as *bilinear varieties*. We also refer to the quantity $\operatorname{codim}_G U + \operatorname{codim}_H V + r$ as the *codimension* of the bilinear variety.

The bounds in Theorem 1·2 were later improved by Hosseini and Lovett [**8**] to nearly optimal ones, namely $\operatorname{codim}_G U, \operatorname{codim}_H V, r \leq O(\log^{80}(2\delta^{-1}))$, at the small cost of considering a slightly longer sequence of directional difference maps.

In [**5**], it was observed that if a dense set A is invariant under maps $\mathsf{D}_{\mathsf{ver}}$ and $\mathsf{D}_{\mathsf{hor}}$ then Theorem 1·2 implies that A contains a large bilinear variety of small codimension. Such sets were named *transverse* by Bienvenu, González–Sánchez and Martínez [**3**].

Definition 1·3. A subset $A \subseteq G \times H$ is said to be transverse if $\{y \in H \colon (x, y) \in A\}$ is a non-empty[1] subspace of H for each $x \in G$ and $\{x \in G \colon (x, y) \in A\}$ is a non-empty subspace of G for each $y \in H$.

In this language, Theorem 1·2, with bounds due to Hosseini and Lovett, implies the following corollary.

COROLLARY 1·4. *Let $A \subseteq G \times H$ be a transverse set of size at least $\delta|G||H|$. Then A contains a bilinear variety of codimension $O(\log^{80}(2\delta^{-1}))$.*

In [**3**], Bienvenu, González-Sánchez and Martínez considered the problem of understanding the relationship between transverse sets and bilinear varieties. Interestingly, they constructed examples of transitive sets that are not bilinear varieties, indicating that some subtleties are likely to arise in the proof of Corollary 1·4.

Despite the significant algebraic structure present in the assumptions, the only known proof of Corollary 1·4 proceeds via the bilinear Bogolyubov argument, which is capable of generating algebraic structure starting from sets that are merely dense. Here we provide a different, more direct proof of Theorem 1·3, giving improved bounds. Our proof also identifies the relevant bilinear map combinatorially, rather than by using Fourier analysis or

[1] The restriction that the subspaces are non-empty is a minor one, as if we allowed empty subspaces, we could have passed to product $\{x \in G \colon (x, 0) \in A\} \times \{y \in H \colon (0, y) \in A\}$.

https://doi.org/10.1017/S0305004125101771 Published online by Cambridge University Press

Freiman's theorem and related results about approximate homomorphisms. The main result of this paper is the following.

THEOREM 1·5. *Let $A \subseteq G \times H$ be a transverse set of size at least $\delta|G||H|$. Then A contains a set of the form $\{(x,y) \in U \times V\colon \beta(x,y) = 0\}$, where $U \leq G$ and $V \leq H$ are subspaces and $\beta\colon G \times H \to \mathbb{F}_p^r$ is a bilinear map, obeying the bounds $\mathrm{codim}_G U = O(\log_p^3 \delta^{-1})$, $\mathrm{codim}_H V = O(\log_p^2 \delta^{-1})$ and $r = O(\log_p \delta^{-1})$.*

Let us remark that the bound $r = O(\log_p \delta^{-1})$ on the codimension of the bilinear map β above is optimal up to constant. We also remark that the proof by Hosseini and Lovett implies that $\mathrm{codim}_G U = O(\log^{80}(2\delta^{-1}))$, $\mathrm{codim}_H V = O(\log_p^{28}(2\delta^{-1}))$ and $r = O(\log^4(2\delta^{-1}))$.

Finally, a multidimensional generalisation of Corollary 1·4 was obtained in [7], again depending on results in the spirit of the bilinear Bogolyubov argument (appropriately generalised). We expect that arguments of this paper could be useful in finding a direct proof of the multidimensional generalisation of Corollary 1·4, but, as we shall remark later in the paper, this approach would no longer be elementary in the higher order setting.

Proof overview. The key concept in the proof is that of a *linear system of subspaces* inside a vector space H. This is a collection of subspaces $(V_x)_{x \in G}$, where $V_x \leq H$, and the indices x range over another vector space G, with the property that $V_{x_1+x_2} \subseteq V_{x_1} + V_{x_2}$ for all $x_1, x_2 \in G$ (see Definition 2·1). There is a natural way to pass from a transverse set $A \subseteq G \times H$ to a linear system of subspaces, namely by fixing a dot product $\cdot$ on H and taking orthogonal complements: $V_x = \{z \in H\colon (\forall y \in H)\ (x,y) \in A \implies y \cdot z = 0\}$.

The first important step of the proof is to prove structural result for quasirandom linear systems of subspaces (Proposition 2·7), by which we mean that most V_x have same dimension d, and most intersections $V_{x_1} \cap V_{x_2}$ are the zero subspace. The conclusion is that then the subspaces V_x can be *generated linearly*, namely that there exist a bilinear map $\Phi\colon G \times \mathbb{F}_p^d \to H$ and a linear map $\Psi\colon \mathbb{F}_p^d \to H$ such that $\{\Phi(x,\lambda) + \Psi(\lambda)\colon \lambda \in \mathbb{F}_p^d\} = V_x$ holds for most x. Notice the occurrence of $\mathbb{F}_p^d$ in the domains of the maps, which is optimal as the dimension of V_x is typically d. Unlike the previous results concerning the bilinear setting, where the structure was obtained by picking an element $\phi(x)$ from each V_x and then exploiting approximate linearity in such a function ϕ, we use the full power of quasirandomness of a linear systems of subspaces and find maps $\phi_x : \mathbb{F}_p^d \to V_x$, which parametrise subspaces V_x, which respect many additive quadruples, in the sense that

$$(\forall \lambda \in \mathbb{F}_p^d)\ \ \phi_{x+a}(\lambda) - \phi_x(\lambda) - \phi_{y+a}(\lambda) + \phi_y(\lambda) = 0.$$

In other words, we may obtain a near homomorphism from G to $\mathrm{Hom}(\mathbb{F}_p^d, H)$. To find such maps ϕ_x, we need some additional linear-algebraic results (Lemmas 2·4 and 2·5). We shall say more about the proof after the statement of the structural result.

Once we have the structural result for quasirandom linear systems of subspaces, it remains to find such a system. To this end, we prove an algebraic regularity lemma for transverse sets (Lemma 2·10). The proof is based on algebraic dependent random choice.

Apart from evading results such as Freiman's theorem and the Bogolyubov argument, another surprising feature of the proof is that, despite the property of being a transverse set

https://doi.org/10.1017/S0305004125101771 Published online by Cambridge University Press

is exact, or '100%' kind of property in the usual language of additive combinatorics, the proof relies on results which concern the so-called almost exact or '99%' regime.

2. *Proof of Theorem 1·5*

Notation. In this paper, we shall frequently consider subsets of products of two vectors spaces G and H. Given $X \subset G \times H$ and an element $x \in G$, we write $X_{x\bullet} = \{y \in H\colon (x,y) \in X\}$ for the *(vertical) slice* of X in the column indexed by x. Likewise, for an element $y \in H$, we write $X_{\bullet y} = \{x \in G\colon (x,y) \in X\}$ for the *(horizontal) slice* of X in the row indexed by y. Intuitively, we take a geometric perspective and informally use a Cartesian coordinate system, where G plays the role of the x-axis and H plays the role of the y-axis.

In order to study a transverse set A, we move to the set of orthogonal complements of its columns. Namely, for each $x \in G$, we consider the subspace $V_x = A_{x\bullet}^{\perp}$, where $^{\perp}$ is taken with respect to some fixed inner product on G. Observe that these subspaces satisfy $V_{x_1+x_2} \subseteq V_{x_1} + V_{x_2}$. Indeed, since A is a transverse set, its rows are subspaces so $A_{x_1\bullet} \cap A_{x_2\bullet} \subseteq A_{x_1+x_2\bullet}$. Therefore

$$V_{x_1+x_2} = A_{x_1+x_2\bullet}^{\perp} \subseteq (A_{x_1\bullet} \cap A_{x_2\bullet})^{\perp} = A_{x_1\bullet}^{\perp} + A_{x_2\bullet}^{\perp} = V_{x_1} + V_{x_2},$$

as claimed. This motivates the following definition.

Definition 2·1. Let V_x be a non-empty subspace of H for each $x \in G$ and assume that $V_0 = \{0\}$. We say that $(V_x)_{x\in G}$ is a linear system of subspaces inside H if $V_{x_1+x_2} \subseteq V_{x_1} + V_{x_2}$ holds for all $x_1, x_2 \in G$.

Observe that in a transverse set A we have $A_{0\bullet} = H$, so $A_{0\bullet}^{\perp} = \{0\}$. Thus, every transverse set gives rise to a linear system of subspaces in a natural way. Likewise, it turns out that every linear system of subspaces produces a transverse set.

LEMMA 2·2. *Let $(V_x)_{x\in G}$ be a linear system of subspaces inside H. Let $A = \cup_{x\in G}\{x\} \times V_x^{\perp}$. Then A is a transverse set.*

Proof. Note that all columns of A are non-empty as each V_x is non-empty, and that all rows are non-empty since $V_0 = \{0\}$. Let $x_1, x_2, y \in G$ be elements such that $(x_1, y), (x_2, y) \in A$. Then $y \in V_{x_1}^{\perp} \cap V_{x_2}^{\perp} = (V_{x_1} + V_{x_2})^{\perp}$. Since $(V_x)_{x\in G}$ is a linear system of subspaces, we also have that $y \in V_{x_1+x_2}^{\perp}$, so $(x_1 + x_2, y) \in A$, proving that the rows of A are subspaces. On the other hand, the columns of A are subspaces by definition.

We remark also the following properties of a linear system of subspaces.

LEMMA 2·3. *Let $(V_x)_{x\in G}$ be a linear system of subspaces inside H.*

(i) *For all $\lambda \in \mathbb{F}_p \setminus \{0\}$ and $x \in G$ we have $V_{\lambda x} = V_x$.*
(ii) *Whenever $x_1, \dots, x_r \in G$ add up to 0, we have $V_{x_1} + \cdots + V_{x_r} = V_{x_1} + \cdots + V_{x_{r-1}}$.*

Proof.

(i) Since $(V_x)_{x\in G}$ is a linear system of subspaces, we have

$$V_{\lambda x} \subseteq V_{(\lambda-1)x} + V_x \subseteq (V_{(\lambda-2)x} + V_x) + V_x \subseteq \cdots \subseteq \lambda V_x = V_x.$$

proving the first claim.

https://doi.org/10.1017/S0305004125101771 Published online by Cambridge University Press

(ii) Suppose that $x_1 + \cdots + x_r = 0$. Then, by the first part, $V_{x_r} = V_{-x_1-\cdots-x_{r-1}} = V_{x_1+\cdots+x_{r-1}}$. Since $(V_x)_{x\in G}$ is a linear system of subspaces, we have

$$V_{x_1+\cdots+x_{r-1}} \subseteq V_{x_1+\cdots+x_{r-2}} + V_{x_{r-1}} \subseteq \cdots \subseteq V_{x_1} + \ldots V_{x_{r-1}},$$

so $V_{x_r} \subseteq V_{x_1} + \ldots V_{x_{r-1}}$, from which the second claim follows.

The following lemma allows us to pass from subspaces to linear maps. It is an exact version of [9, lemma 34], obtained by taking $K = 1$ in that paper. To make the paper self-contained, we provide a proof.

LEMMA 2·4. *Let $U_1, U_2, U_3, U_4 \leq H$ be subspaces of dimension d such that $|U_{i_1} + U_{i_2} + U_{i_3}| = p^{3d}$ for any three distinct indices $i_1, i_2, i_3 \in [4]$ and such that $|U_1 + U_2 + U_3 + U_4| = p^{3d}$. Let $\phi_4 \colon \mathbb{F}_p^d \to U_4$ be a linear isomorphism. Then there exist linear isomorphisms $\phi_i \colon \mathbb{F}_p^d \to U_i$, $i \in [3]$, such that $\phi_1 + \phi_2 = \phi_3 + \phi_4$.*

Proof. Let $e_1, \ldots, e_d$ be the standard basis of $\mathbb{F}_p^d$. Since $U_1 + U_2 + U_3 \subseteq U_1 + U_2 + U_3 + U_4$ and $|U_1 + U_2 + U_3| = p^{3d} = |U_1 + U_2 + U_3 + U_4|$, we have $U_1 + U_2 + U_3 = U_1 + U_2 + U_3 + U_4$, so $U_4 \subseteq U_1 + U_2 + U_3$. In particular, every $u \in U_4$ can be written as a sum $u = u_1 + u_2 + u_3$ for $u_i \in U_i$, $i \in [3]$. Using this fact, for each $i \in [d]$, define $\phi_1(e_i)$, $\phi_2(e_i)$ and $\phi_3(e_i)$ as elements U_1, U_2 and U_3, respectively, such that $\phi_4(e_i) = \phi_1(e_i) + \phi_2(e_i) - \phi_3(e_i)$. Extend each ϕ_i linearly to the whole space $\mathbb{F}_p^d$. It is easy to check that these maps satisfy the conclusion.

It is important that we get further maps ϕ_1, ϕ_2 and ϕ_3 for any given ϕ_4, rather than only being able to find a single quadruple of maps $(\phi_1, \ldots, \phi_4)$ satisfying the properties in the conclusion of the lemma.

Furthermore, we need the related uniqueness result, which is an exact version of [9, lemma 35], obtained by taking $K = 1$ and $r = 0$ in that paper. Like above, we offer a short proof.

LEMMA 2·5. *Let $U_1, U_2, V_1, V_2, W \leq H$ be subspaces of dimension d such that $W \cap (U_1 + U_2 + V_1 + V_2) = \{0\}$. Suppose that linear maps $\phi_i \colon \mathbb{F}_p^d \to U_i$ for $i \in [2]$, $\psi_i \colon \mathbb{F}_p^d \to V_i$ for $i \in [2]$ and $\theta \colon \mathbb{F}_p^d \to W$ satisfy $\phi_1 + \phi_2 + \psi_1 + \psi_2 + \theta = 0$. Then $\theta = 0$.*

Proof. Let $x \in \mathbb{F}_p^d$. Then $\phi_1(x) + \phi_2(x) + \psi_1(x) + \psi_2(x) = -\theta(x)$. This implies that $\theta(x) \in U_1 + U_2 + V_1 + V_2 \cap W$, which is the zero subspace by assumptions. Hence $\theta = 0$.

We shall also make use of elementary fact that near-homomorphisms of vector spaces come from (exact) homomorphisms. The formulation below is a slight rephrasing of [7, corollary 26]. As with the two lemmas above, we give a proof for the sake of self-completeness.

LEMMA 2·6. *Let G_1 and G_2 be two finite-dimensional vector spaces over $\mathbb{F}_p$. Let $A \subseteq G_1$ be a subset and let $\phi \colon A \to G_2$ be a map such that $x + y - z \in A$ and $\phi(x) + \phi(y) = \phi(z) + \phi(x + y - z)$ both hold for at least $(1 - \varepsilon)|G_1|^3$ triples $(x, y, z) \in A^3$. Provided $\varepsilon < 10^{-5}$, there exists an affine map $\Phi \colon G_1 \to G_2$ such that $\Phi(x) = \phi(x)$ holds for at least $(1 - 2\sqrt{\varepsilon})|G_1|$ of elements $x \in G_1$.*

Remark. Returning to the last remark from the introduction, in order to prove the multi-dimensional generalisation of Corollary 1·4 using the approach in this paper, we would

https://doi.org/10.1017/S0305004125101771 Published online by Cambridge University Press

need a higher-order generalisation of Lemma 2·6. Such a generalisation concerns extending multilinear maps defined on $1-o(1)$ proportion of multilinear varieties. Unlike the lemma above, such a higher-order result is highly non-trivial, and, in fact, is one of the main results from [7].

Proof. For this proof, we need the notation for convolutions defined by differences instead of sums. Namely, if $f, g : G_1 \to \mathbb{R}$ are functions, then $f \overline{*} g$ is another function from G_1 to $\mathbb{R}$, defined by $f \overline{*} g(x) = \sum_{y\in G_1} f(x+y)g(y)$.

Let $\tilde{A} = \{(a, \phi(a)) : a \in A\} \subseteq G_1 \times G_2$ be the graph of the function ϕ. Write also $\tilde{G} = G_1 \times G_2$. Our primary aim is to show that $\tilde{A}$ is almost a subgroup of $\tilde{G}$, from which we shall deduce the functional conclusion fairly quickly. Note that $|\tilde{A}| = |A|$ and that $(1-\varepsilon)|G_1| \le |A| \le |G_1|$. The assumption in the lemma means that $\tilde{A}$ has at least $(1-\varepsilon)|\tilde{A}|^3$ additive quadruples. Using the convolution notation, we have $\sum_{x\in\tilde{G}} \mathbb{1}_{\tilde{A}} \overline{*} \mathbb{1}_{\tilde{A}}(x)^2 \ge (1-\varepsilon)|\tilde{A}|^3$. Let S be the set of all $x \in \tilde{G}$ such that $\mathbb{1}_{\tilde{A}} \overline{*} \mathbb{1}_{\tilde{A}}(x) \ge (1-\sqrt{\varepsilon})|\tilde{A}|$. Note that $\|\mathbb{1}_{\tilde{A}} \overline{*} \mathbb{1}_{\tilde{A}}\|_\infty \le |\tilde{A}|$. Observe that

$$\begin{aligned}\sqrt{\varepsilon}|\tilde{A}| \sum_{x\in\tilde{G}\setminus S} \mathbb{1}_{\tilde{A}} \overline{*} \mathbb{1}_{\tilde{A}}(x) &\le \sum_{x\in\tilde{G}\setminus S} (|\tilde{A}| - \mathbb{1}_{\tilde{A}} \overline{*} \mathbb{1}_{\tilde{A}}(x))\mathbb{1}_{\tilde{A}} \overline{*} \mathbb{1}_{\tilde{A}}(x)\\ &\le \sum_{x\in\tilde{G}} (|\tilde{A}| - \mathbb{1}_{\tilde{A}} \overline{*} \mathbb{1}_{\tilde{A}}(x))\mathbb{1}_{\tilde{A}} \overline{*} \mathbb{1}_{\tilde{A}}(x)\\ &= |\tilde{A}|\Big(\sum_{x\in\tilde{G}} \mathbb{1}_{\tilde{A}} \overline{*} \mathbb{1}_{\tilde{A}}(x)\Big) - \Big(\sum_{x\in\tilde{G}} \mathbb{1}_{\tilde{A}} \overline{*} \mathbb{1}_{\tilde{A}}(x)^2\Big)\\ &\le \varepsilon|\tilde{A}|^3,\end{aligned}$$

so $\sum_{x\in S} \mathbb{1}_{\tilde{A}} \overline{*} \mathbb{1}_{\tilde{A}}(x) \ge (1-\sqrt{\varepsilon})|\tilde{A}|^2$, implying also $|S| \ge (1-\sqrt{\varepsilon})|\tilde{A}|$.

We now show that $|S-S| \le (1+7\sqrt{\varepsilon})|S|$. Let $x, y \in S$. Thus $|\tilde{A} \cap (\tilde{A}-x)| = \mathbb{1}_{\tilde{A}} \overline{*} \mathbb{1}_{\tilde{A}}(x) \ge (1-\sqrt{\varepsilon})|\tilde{A}|$ and $|\tilde{A} \cap (\tilde{A}-y)| = \mathbb{1}_{\tilde{A}} \overline{*} \mathbb{1}_{\tilde{A}}(y) \ge (1-\sqrt{\varepsilon})|\tilde{A}|$. Hence,

$$\begin{aligned}\mathbb{1}_{\tilde{A}} \overline{*} \mathbb{1}_{\tilde{A}}(x-y) &= |(\tilde{A}-x) \cap (\tilde{A}-y)| \ge |\tilde{A} \cap (\tilde{A}-x) \cap (\tilde{A}-y)|\\ &\ge |\tilde{A} \cap (\tilde{A}-x)| + |\tilde{A} \cap (\tilde{A}-y)| - |\tilde{A}| \ge (1-2\sqrt{\varepsilon})|\tilde{A}|.\end{aligned}$$

However, using $\sum_{d\in\tilde{G}} \mathbb{1}_{\tilde{A}} \overline{*} \mathbb{1}_{\tilde{A}}(d)^2 \le |\tilde{A}|^3$ and elementary inequality $(1-2\sqrt{\varepsilon})^{-2} \le 1+5\sqrt{\varepsilon}$ which holds when $\varepsilon \le 2^{-8}$, the inequality $\mathbb{1}_{\tilde{A}} \overline{*} \mathbb{1}_{\tilde{A}}(d) \ge (1-2\sqrt{\varepsilon})|\tilde{A}|$ can hold for at most $|\tilde{A}|^3/(1-2\sqrt{\varepsilon})^2|\tilde{A}|^2 \le (1+5\sqrt{\varepsilon})|\tilde{A}|$ elements $d \in \tilde{G}$, showing that $|S-S| \le (1+5\sqrt{\varepsilon})|\tilde{A}| \le (1+7\sqrt{\varepsilon})|S|$.

Since $\varepsilon \le 28^{-2}$, we have $|S-S| \le (5/4)|S|$. A similar elementary argument to the one above, where one considers intersection of sets $S \cap S-a$ and $S \cap S-b$ for elements $a, b \in S-S$, shows that $S-S$ is a subgroup of $\tilde{G}$, which we denote by H. Fix arbitrary $s \in S$. In particular, since $\varepsilon \le 19^{-2}$, we have

$$|H| = |S-S| \le (1+7\sqrt{\varepsilon})|S| \le (1+7\sqrt{\varepsilon})(1-\sqrt{\varepsilon})^2|\tilde{A}| \le (1+10\sqrt{\varepsilon})|\tilde{A}|.$$

Since $S - s \subseteq H$ and thus $S \subseteq s+H$, we have

$$(1-\sqrt{\varepsilon})|\tilde{A}|^2 \le \sum_{d\in s+H} \mathbb{1}_{\tilde{A}} \overline{*} \mathbb{1}_{\tilde{A}}(d) = \sum_{x,y\in\tilde{G}} \mathbb{1}_{\tilde{A}}(x)\mathbb{1}_{\tilde{A}}(y)\mathbb{1}_{s+H}(x-y).$$

By averaging over $y \in \tilde{A}$, we obtain an element y such that $|(s+y+H) \cap \tilde{A}| \ge (1-\sqrt{\varepsilon})|\tilde{A}|$.

https://doi.org/10.1017/S0305004125101771 Published online by Cambridge University Press

We now return to the initial functional perspective. Let $\pi_1 : G_1 \times G_2 \to G_1$ and $\pi_2 : G_1 \times G_2 \to G_2$ be the natural projections. Note that π_1 is injective on $\tilde{A}$, which implies that

$$(1-\sqrt{\varepsilon})|\tilde{A}| \le |\tilde{A} \cap (s+y+H)| = |\pi_1(\tilde{A} \cap (s+y+H))| \le |\pi_1(H)| \le |H| \le (1+5\sqrt{\varepsilon})|\tilde{A}|.$$

Thus, $|\pi_1(H)| > (1/2)|H|$ and $\pi_1(H)$ and H, being subgroups, both have sizes which are powers of p, so $|\pi_1(H)| = |H|$, and thus π_1 is injective on H. Furthermore, since $|A| \ge (1-\varepsilon)|G_1|$, we must have $\pi_1(H) = G_1$. Therefore, writing $h \in s+y+H$ as $(\pi_1(h), \pi_2(h))$ and then defining a map $\Phi : \pi_1(h) \mapsto \pi_2(h)$, we obtain an affine map from $G_1 = \pi_1(s+y+H)$ to G_2. Note that H is then the graph of Φ. Since $|(s+y+H) \cap \tilde{A}| \ge (1-\sqrt{\varepsilon})|\tilde{A}|$ we have $\phi(a) = \Phi(a)$ for at least $(1-\sqrt{\varepsilon})|A| \ge (1-2\sqrt{\varepsilon})|G_1|$ elements $a \in A$, as desired.

The heart of the proof of Theorem 1·5 is that a quasirandom linear system of subspaces is generated by a bilinear map of expected codimension, which is the content of the next proposition.

PROPOSITION 2·7. *Let $(V_x)_{x\in G}$ be a linear system of subspaces inside H. Let d be a non-negative integer and let $\varepsilon > 0$ be such that*:

(i) *for at least $(1-\varepsilon)|G|$ of elements $x \in G$ we have $|V_x| = p^d$; and*

(ii) *for at least $(1-\varepsilon)|G|^2$ of pairs $(x_1, x_2) \in G^2$ we have $V_{x_1} \cap V_{x_2} = \{0\}$.*

Provided $\varepsilon \le p^{-8d}10^{-32}$, there exist a bilinear map $\Phi\colon G \times \mathbb{F}_p^d \to H$ and a linear map $\Psi\colon \mathbb{F}_p^d \to H$ such that $\{\Phi(x,\lambda) + \Psi(\lambda)\colon \lambda \in \mathbb{F}_p^d\} = V_x$ holds for at least $(1-200p^d\sqrt[8]{\varepsilon})|G|$ elements $x \in G$.

We remark that Proposition 2·7 is similar in spirit to [**9**, theorem 33], and these two results can roughly be viewed as '99%' and '1%' cases of the same phenomenon. Given the additional structure in the setting of the present paper, we are able to obtain a more streamlined proof. An important difference from [**9**] is that throughout that paper one has sufficiently strong quasirandom properties, which is not the case in this paper. For that reason, we need an efficient regularity lemma for transverse sets, which we prove later.

Brief proof overview. As described in the introduction, our aim is to find parameterizing linear maps $\phi_x : \mathbb{F}_p^d \to V_x$ such that $\phi_x + \phi_y = \phi_{x+y-a} + \phi_a$ holds for most of choices of x, y and a. We shall first pick a such that most of pairs $(x,y) \in G$ are well behaved, in the sense that

$$\begin{aligned} V_a + V_x + V_y = V_a + V_x + V_{x+y-a} = V_a + V_y + V_{x+y-a} &= V_x + V_y + V_{x+y-a} \\ &= V_a + V_x + V_y + V_{x+y-a} \end{aligned}$$

and subspaces V_a, V_x, V_y and V_{x+y-a} have dimension d. Fix an arbitrary linear isomorphism $\theta\colon \mathbb{F}_p^d \to V_a$. Write $q = (x,y)$. The displayed property allows us to apply Lemma 2·4 to obtain linear isomorphisms $\theta_q^1\colon \mathbb{F}_p^d \to V_x$, $\theta_q^2\colon \mathbb{F}_p^d \to V_y$ and $\theta_q^3\colon \mathbb{F}_p^d \to V_{x+y-a}$, such that $\theta_q^1 + \theta_q^2 = \theta + \theta_q^3$. In the long run, we want θ_q^3 to depend only on x, and, moreover, in a linear fashion. To show that θ_q^3 depends only on $x+y$ instead of both x and y, we rely on Lemma 2·5; this is achieved in Claim 2·9. To stress this sharper dependence, we write ϕ_{x+y-a}^3 instead of θ_q^3. In order to obtain linearity, we look at well-behaved pairs of the form

https://doi.org/10.1017/S0305004125101771 Published online by Cambridge University Press

(x, y), (x', y), (x, y') and (x', y'). Algebraic manipulation shows that

$$\phi^3_{x+y-a} + \phi^3_{x'+y'-a} - \phi^3_{x'+y-a} - \phi^3_{x+y'-a} = 0.$$

Therefore, we may view ϕ^3 as a map from a very dense subset of G to $\mathrm{Hom}(\mathbb{F}_p^d, H)$ that respects most additive quadruples. An application of Lemma 2·6 completes the proof.

Proof of Proposition 2·7. Let us first observe that our assumptions imply that $V_{x_0} \cap (V_{x_1} + \cdots + V_{x_r}) = \{0\}$ holds for all but at most $\left(p^{rd}\sqrt{\varepsilon} + r\varepsilon\right)|G|^{r+1}$ of $(r+1)$-tuples $(x_0, \ldots, x_r) \in G^{[0,r]}$. To see this, assume the contrary. Let $\mathcal{G}$ be the set of all $x \in G$ such that $|V_x| = p^d$. The total number of $(r+1)$-tuples $(x_0, \ldots, x_r) \in G^{[0,r]}$ such that one of $x_1, \ldots, x_r$ is not in $\mathcal{G}$ is at most $r\varepsilon|G|^{r+1}$. Hence we get more than $p^{rd}\sqrt{\varepsilon}|G|^{r+1}$ of $(r+1)$-tuples $(x_0, \ldots, x_r) \in G \times \mathcal{G}^r$ such that $V_{x_0} \cap (V_{x_1} + \cdots + V_{x_r}) \neq \{0\}$. By averaging, we get some r-tuple $(x_1, \ldots, x_r) \in \mathcal{G}^r$ such that for more than $p^{rd}\sqrt{\varepsilon}|G|$ elements $x' \in G$ we have $V_{x'} \cap S \neq \{0\}$, where $S = V_{x_1} + \cdots + V_{x_r}$. Since $|V_{x_i}| = p^d$, we see that $|S| \leq p^{rd}$. By averaging over elements of $S \setminus \{0\}$, there exists a non-zero $s \in S$ such that $s \in V_{x'}$ for more than $\sqrt{\varepsilon}|G|$ elements $x' \in G$. However, then $s \in V_{x'} \cap V_{x''}$ for more than $\varepsilon|G|^2$ pairs (x', x''), which is a contradiction.

We shall now exploit this fact to get more control over sums of subspaces indexed by configurations of elements of G. Write $\varepsilon' = p^{4d}\sqrt{\varepsilon} + 4\varepsilon$, corresponding to the factor in the bound in the observation above for $r \leq 4$.

CLAIM 2·8. *For all but at most* $8\varepsilon'|G|^4$ *of quadruples* $(a, x, y, z) \in G^4$ *we have*:

(i) $|V_a| = |V_x| = |V_y| = |V_{x+y-a}| = |V_z| = p^d$;
(ii) $|V_a + V_x + V_y| = p^{3d}$; *and*
(iii) $V_{x+y-a} \cap (V_x + V_y + V_z) = \{0\}$.

Proof. Note that $|V_a + V_x + V_y| = p^{3d}$ if $V_a \cap (V_x + V_y) = \{0\}$, $V_x \cap V_y = \{0\}$ and $|V_a| = |V_x| = |V_y| = p^d$. By assumptions of the proposition and the observation above, all these conditions, as well as $|V_{x+y-a}| = |V_z| = p^d$ and property (iii), are satisfied for all but at most $8\varepsilon'|G|^3$ triples $(a, x, y) \in G^3$.

For the remainder of the proof, we fix an element $a \in G$ such that the set $\mathcal{T}$ of triples (x, y, z) such that the quadruple (a, x, y, z) satisfies properties in the conclusion of the claim above is of size at least $(1 - 8\varepsilon')|G|^3$. Let $\mathcal{P}$ be the set of all pairs $(x, y) \in G^2$ such that $|V_x| = |V_y| = |V_{x+y-a}| = p^d$ and the subspace $V_a + V_x + V_y$ has size p^{3d}. By Lemma 2·3, for each $(x, y) \in \mathcal{P}$, we have that

$$\begin{aligned} V_a + V_x + V_y = V_a + V_x + V_{x+y-a} = V_a + V_y + V_{x+y-a} &= V_x + V_y + V_{x+y-a} \\ &= V_a + V_x + V_y + V_{x+y-a}. \quad (2{\cdot}1) \end{aligned}$$

Since $|\mathcal{T}| \geq (1 - 8\varepsilon')|G|^3$, we also have that $|\mathcal{P}| \geq (1 - 8\varepsilon')|G|^2$.

Fix an arbitrary linear isomorphism $\theta\colon \mathbb{F}_p^d \to V_a$. For each $q = (x, y) \in \mathcal{P}$, equalities in (2·1) imply that the subspaces V_a, V_x, V_y and V_{x+y-a} satisfy the conditions of Lemma 2·4. Lemma 2·4 provides us with linear isomorphisms $\theta^1_q\colon \mathbb{F}_p^d \to V_x$, $\theta^2_q\colon \mathbb{F}_p^d \to V_y$ and $\theta^3_q\colon \mathbb{F}_p^d \to V_{x+y-a}$, such that $\theta^1_q + \theta^2_q = \theta + \theta^3_q$. Our next step is to show that these maps do not depend entirely on the pair q, but only on the relevant point in G.

https://doi.org/10.1017/S0305004125101771 Published online by Cambridge University Press

Although $q=(x,y)\in\mathcal{P}$ is a pair, we misuse the notation and write $q_3=x+y-a$ for the third element associated to q.

CLAIM 2·9. *For each $i\in[3]$ there exists a set $X_i\subseteq G$ of size $|X_i|\geq(1-10\sqrt{\varepsilon'})|G|$ and linear isomorphisms $\phi_x^i\colon \mathbb{F}_p^d\to V_x$ for $x\in X_i$ such that the following holds. For all but at most $30\sqrt{\varepsilon'}|G|^2$ pairs $q\in\mathcal{P}$ we have $q_i\in X_i$ for each $i\in[3]$ and $\theta_q^i=\phi_{q_i}^i$.*

Proof. Let us focus first on the case $i=1$; we shall comment on the other cases afterwards. Let us count the number N of pairs $(q,q')\in\mathcal{P}^2$ such that $q_1=q'_1$ and $\theta_q^1=\theta_{q'}^1$. Notice that whenever $q=(x,y),q'=(x,y')\in\mathcal{P}$ are two pairs with the same first element, then

$$\theta_q^1+\theta_q^2-\theta_q^3=\theta=\theta_{q'}^1+\theta_{q'}^2-\theta_{q'}^3,$$

so $\theta_q^2-\theta_q^3-\theta_{q'}^2-\theta_{q'}^3+(\theta_q^1-\theta_{q'}^1)=0$. If the indexing elements satisfy $V_x\cap(V_y+V_{x+y-a}+V_{y'}+V_{x+y'-a})=\{0\}$, then by Lemma 2·5 we have $\theta_q^1=\theta_{q'}^1$. On the other hand, Lemma 2·3 shows that $V_y+V_{x+y-a}+V_{y'}+V_{x+y'-a}=V_y+V_{x+y-a}+V_{y'}$. By the choice of a, $V_x\cap(V_y+V_{x+y-a}+V_{y'})=\{0\}$ holds for all but at most $8\varepsilon'|G|^3$ triples $(x,y,y')\in G^3$. From this and Cauchy-Schwarz inequality, we conclude that

$$N\geq|G|^{-1}|\mathcal{P}|^2-8\varepsilon'|G|^3\geq(1-24\varepsilon')|G|^3.$$

Let X_1 be the set of all $x\in G$ such that the number of pairs (q,q') above with x as the first element is at least $(1-10\sqrt{\varepsilon'})|G|$. Then $|X_1|\geq(1-10\sqrt{\varepsilon'})|G|$ and we may take ϕ_x^1 to be the most common map appearing as θ_q^1 among triples with $q_1=x$.

The case $i=2$ follows by symmetry. For the case $i=3$, consider $z\in G$ and $q=(x,z+a-x),q'=(x',z+a-x')\in\mathcal{P}$ agreeing in the third associated element, that is $q_3=q'_3=z$. Observe that we need $V_z\cap(V_x+V_{z+a-x}+V_{x'}+V_{z+a-x'})=\{0\}$ property to get the equality of the maps coming from different pairs. By Lemma 2·3, $V_x+V_{z+a-x}+V_{x'}+V_{z+a-x'}=V_x+V_{z+a-x}+V_{x'}$, so we just need $V_z\cap(V_x+V_{z+a-x}+V_{x'})=\{0\}$. After a change of variables $y=z+a-x$, this condition has the form guaranteed by Claim 2·8, so the same proof works in this case as well.

Let $\tilde{\mathcal{P}}$ be the set of all pairs $q\in\mathcal{P}$ for which the conclusion of the claim above holds. Let $\mathrm{Hom}(\mathbb{F}_p^d,H)$ be the vector space of all linear homomorphisms from $\mathbb{F}_p^d$ to H. Write $\phi^3\colon X_3\to\mathrm{Hom}(\mathbb{F}_p^d,H)$ for the map given by $\phi^3(x)=\phi_x^3$. Let us observe that ϕ^3 respects most additive quadruples. Consider the set Q of all quadruples $(x,x',y,y')\in G^4$ such that:

(i) $x,x'\in X_1,y,y'\in X_2,x+y-a,x'+y-a,x+y'-a,x'+y'-a\in X_3$;
(ii) $(x,y),(x',y),(x,y'),(x',y')\in\tilde{\mathcal{P}}$.

By the work above, we have that $|Q|\geq(1-300\sqrt{\varepsilon'})|G|^4$. Each such quadruple gives rise to an additive quadruple respected by ϕ^3. Namely,

$$\begin{aligned}\phi^3(x+y-a)+\phi^3(x'+y'-a)-\phi^3(x'+y-a)-\phi^3(x+y'-a)\\ =\left(\phi_x^1+\phi_y^2-\theta\right)+\left(\phi_{x'}^1+\phi_{y'}^2-\theta\right)-\left(\phi_{x'}^1+\phi_y^2-\theta\right)-\left(\phi_x^1+\phi_{y'}^2-\theta\right)=0.\end{aligned}$$

On the other hand, each additive quadruple in X_3 can arise at most $|G|$ times in this way. Hence, ϕ^3 respects at least $(1-300\sqrt{\varepsilon'})|G|^3$ additive quadruples in X_3.

https://doi.org/10.1017/S0305004125101771 Published online by Cambridge University Press

We are now in position to apply Lemma 2·6. It provides us with an affine map $\Gamma\colon G\to \mathrm{Hom}(\mathbb{F}_p^d, H)$ such that $\Gamma(x)=\phi^3(x)$, and thus $\mathrm{Im}\,\Gamma(x)=\mathrm{Im}\,\phi_x^3=V_x$, holds for at least $(1-\varepsilon'')|G|$ of $x\in X_3$, where $\varepsilon''\le 40\sqrt[4]{\varepsilon'}$, provided $\varepsilon'\le 10^{-15}$. Finally, we define maps $\Phi\colon G\times\mathbb{F}_p^d\to H$ and $\Psi\colon \mathbb{F}_p^d\to H$ as $\Phi(x,\lambda)=\Gamma(x)(\lambda)-\Gamma(0)(\lambda)$ and $\Psi=\Gamma(0)$, giving us the desired conclusion.

To prove Theorem 1·5 it thus remains to locate a sufficiently quasirandom piece. The following regularity lemma for transverse sets achieves this. The proof is based on dependent random choice and the algebraic structure is essential for ensuring the good bounds.

LEMMA 2·10. *Let $A\subset G\times H$ be a transverse set of size $\delta|G||H|$ and let $\varepsilon>0$ be a parameter. Then there exist a non-negative integer $d\le\lceil\log_p(10^4\delta^{-4})\rceil$ and subspaces $U\le G$ and $V\le H$ of codimensions $\mathrm{codim}_G U\le O\Big((\log(2\delta^{-1})^2\log(2\varepsilon^{-1})\Big)$ and $\mathrm{codim}_H V\le O(\log(2\delta^{-1})^2)$ such that*:

(i) *for at least $(1-\varepsilon)|U|$ of elements $x\in U$ we have $|A_{x\bullet}\cap V|=p^{-d}|V|$; and*

(ii) *for at least $(1-3\varepsilon)|U|^2$ of pairs $(x_1,x_2)\in U^2$ we have $|A_{x_1\bullet}\cap A_{x_2\bullet}\cap V|=p^{-2d}|V|$.*

Proof. Let Y be the set of all $y\in H$ such that $|A_{\bullet y}|\ge(\delta/2)|G|$. By averaging, we have $|Y|\ge(\delta/2)|H|$. Pick $y\in Y$ uniformly at random. Then $\mathbb{P}(x\in A_{\bullet y})=(|A_{x\bullet}\cap Y|/|Y|)$. Let X_{bad} be the set of elements $x\in A_{\bullet y}$ such that $|A_{x\bullet}|\le(\delta^2/100)|H|$. Then

$$\begin{aligned}\mathbb{E}\Big(|A_{\bullet y}|-10|X_{\mathrm{bad}}|\Big)&\ge\frac{\delta}{2}|G|-10\sum_{x\in G}\mathbb{P}(x\in X_{\mathrm{bad}})\\&=\frac{\delta}{2}|G|-10\sum_{x\in G}\mathbb{1}(|A_{x\bullet}|\le\frac{\delta^2}{100}|H|)\mathbb{P}(x\in A_{\bullet y})\ge\frac{\delta}{4}|G|.\end{aligned}$$

Hence, there is a choice of y_0 such that $|A_{\bullet y_0}|\ge(\delta/2)|G|$ and $|A_{x\bullet}|\ge(\delta^2/100)|H|$ holds for at least $(9/10)|A_{\bullet y_0}|$ elements $x\in A_{\bullet y_0}$. Hence, for an arbitrary $x\in A_{\bullet y_0}$ we may find $x'\in A_{\bullet y_0}$ such that $|A_{x'\bullet}|\ge(\delta^2/100)|H|$ and $|A_{x+x'\bullet}|\ge(\delta^2/100)|H|$ both hold. As A is transverse we have $(\delta^4/10^4)|H|\le|A_{x'\bullet}\cap A_{x+x'\bullet}|\le|A_{x\bullet}|$. Thus, $|A_{x\bullet}|\ge(\delta^4/10^4)|H|$ holds for all elements $x\in A_{\bullet y_0}$.

Similarly, there exists an element $x_0\in G$ such that $|A_{x_0\bullet}|\ge(\delta/2)|H|$ and $|A_{\bullet y}|\ge(\delta^4/10^4)|G|$ holds for all elements $y\in A_{x_0\bullet}$. Set $U_0=A_{\bullet y_0}$, $V_0=A_{x_0\bullet}$ and $d_0=\lceil\log_p(10^4\delta^{-4})\rceil$. Hence, U_0 and V_0 have codimension at most d_0 in G and H respectively, for each $x\in U_0$ we have $\mathrm{codim}_{V_0}(A_{x\bullet}\cap V_0)\le d_0$ and for each $y\in V_0$ we have $\mathrm{codim}_{U_0}(A_{\bullet y}\cap U_0)\le d_0$.

We now perform an iterative procedure in which at step i we find further subspaces $U_i\le U_0$, $V_i\le V_0$, of codimension $\mathrm{codim}_{U_0}U_i\le O(id_0\log(2\varepsilon^{-1}))$ and $\mathrm{codim}_{V_0}V_i\le O(id_0)$, such that the proportion of elements $x\in U_i$ such that

$$\mathrm{codim}_{V_i}(A_{x\bullet}\cap V_i)\le d_0-i \tag{2·2}$$

is at least $1-1/2\varepsilon$. The elements x in (2·2) are those that have $A_{x\bullet}\cap V_i$ at least as large as we would like, thus we guarantee that most subspaces restricted to V_i have at least the desired size, if not equal to the one required by property (i). The procedure terminates when we find subspaces U_i and V_i satisfying the claimed properties in the lemma, with d taken to be d_0-i. Let us emphasize that the property that is kept during the procedure, namely

https://doi.org/10.1017/S0305004125101771 Published online by Cambridge University Press

that (2·2) holds for most $x \in U_i$, is weaker than the properties (i) and (ii) of the lemma. Hence, in each step of the argument, we shall either show that the current U_i and V_i have the properties (i) and (ii) or we shall find further U_{i+1} and V_{i+1} for which the the properties described around (2·2) hold.

Suppose that we have subspaces U_i, V_i for some $i \geq 0$ (when $i = 0$ we know that U_0 and V_0 have the properties described around (2·2)). If both properties (i) and (ii) of the lemma hold for U_i and V_i and integer $d_0 - i$, we are done. We now distinguish between two cases, depending on which one of the two properties fails.

Property (i) *fails.* We thus have at least $\varepsilon|U_i|$ elements $x \in U_i$ such that $\operatorname{codim}_{V_i}(A_{x\bullet} \cap V_i) \neq d_0 - i$. Since $\operatorname{codim}_{V_i}(A_{x\bullet} \cap V_i) \leq d_0 - i$ holds for at least $(1 - \varepsilon/2)|U_i|$ elements $x \in U_i$, we conclude that $\operatorname{codim}_{V_i}(A_{x\bullet} \cap V_i) \leq d_0 - i - 1$ holds for at least $(\varepsilon/2)|U_i|$ elements $x \in U_i$. Let $r = \lceil 2 \log_p (4\varepsilon^{-1}) \rceil$ and take $y_1, \ldots, y_r$ uniformly and independently at random from V_i. Let $U_{i+1} = U_i \cap A_{\bullet y_1} \cap \cdots \cap A_{\bullet y_r}$ and set $V_{i+1} = V_i$. Let X_{bad} be the set of all elements $x \in U_{i+1}$ such that $\operatorname{codim}_{V_i}(V_i \cap A_{x\bullet}) \geq d_0 - i$. Observe that, for each $x \in U_i$, $\mathbb{P}(x \in U_{i+1}) = (|A_{x\bullet} \cap V_i|/|V_i|)^r = p^{-r \operatorname{codim}_{V_i}(V_i \cap A_{x\bullet})}$. Then

$$\begin{aligned}\mathbb{E}\Big(|U_{i+1}| - 2\varepsilon^{-1}|X_{\text{bad}}|\Big) &\geq p^{-r(d_0-i-1)}\frac{\varepsilon}{2}|U_i| - 2\varepsilon^{-1}p^{-r(d_0-i)}|U_i| \\ &= 2\varepsilon^{-1}p^{-r(d_0-i)}|U_i|\Big(p^r\frac{\varepsilon^2}{4} - 1\Big) > 0,\end{aligned}$$

by our choice of r. Hence, there is a choice of $y_1, \ldots, y_r$ producing the desired pair of subspaces U_{i+1} and V_{i+1}. Note that $\operatorname{codim}_{U_i} U_{i+1} \leq rd_0 = O(d_0 \log_p (2\varepsilon^{-1}))$ and $\operatorname{codim}_{V_i} V_{i+1} = 0$.

Property (ii) *fails, but* (i) *holds.* In this case have at least $3\varepsilon|U_i|^2$ pairs of elements $(x_1, x_2) \in U_i^2$ such that $\operatorname{codim}_{V_i}(A_{x_1\bullet} \cap A_{x_2\bullet} \cap V_i) \neq 2(d_0 - i)$. We call such a pair (x_1, x_2) *irregular*. Let R be the set of all $a \in U_i$ such that $\operatorname{codim}_{V_i}(A_{a\bullet} \cap V_i) = d_0 - i$. Since property (i) holds, we have $|R| \geq (1 - \varepsilon)|U_i|$. The number of irregular pairs (x_1, x_2) such that $x_1 \notin R$ is at most $\varepsilon|U_i|^2$. Hence, there at least $2\varepsilon|U_i|^2$ irregular pairs $(x_1, x_2) \in R \times U_i$, so by averaging, there exists $a \in R$ such that (a, x) is irregular for at least $2\varepsilon|U_i|$ choices of $x \in U_i$.

By assumption (2·2), at least $\varepsilon|U_i|$ of such elements $x \in U_i$ also have the property that $\operatorname{codim}_{V_i} A_{x\bullet} \cap V_i \leq d_0 - i$. Note that

$$\begin{aligned}\operatorname{codim}_{V_i}(A_{a\bullet} \cap A_{x\bullet} \cap V_i) &= \dim V_i - \dim (A_{a\bullet} \cap A_{x\bullet} \cap V_i) \\ &= \dim V_i - \dim (A_{a\bullet} \cap V_i) - \dim (A_{x\bullet} \cap V_i) + \dim \Big((A_{a\bullet} \cap V_i) + (A_{x\bullet} \cap V_i)\Big) \\ &= \operatorname{codim}_{V_i} A_{a\bullet} + \operatorname{codim}_{V_i} A_{x\bullet} - \dim V_i + \dim \Big((A_{a\bullet} \cap V_i) + (A_{x\bullet} \cap V_i)\Big) \\ &\leq \operatorname{codim}_{V_i} A_{a\bullet} + \operatorname{codim}_{V_i} A_{x\bullet} \leq 2(d_0 - i),\end{aligned}$$

where we used $(A_{a\bullet} \cap V_i) + (A_{x\bullet} \cap V_i) \subseteq V_i$ in the inequality between the last two lines. Since (a, x) is irregular, we have $\operatorname{codim}_{V_i}(A_{a\bullet} \cap A_{x\bullet} \cap V_i) < 2(d_0 - i)$. Setting $V_{i+1} = V_i \cap A_{a\bullet}$, for such an x we have

$$\begin{aligned}&\operatorname{codim}_{V_{i+1}}(A_{x\bullet} \cap V_{i+1}) \\ &\quad = \dim (V_{i+1}) - \dim (A_{x\bullet} \cap V_{i+1}) = \dim (V_i \cap A_{a\bullet}) - \dim (V_i \cap A_{x\bullet} \cap A_{a\bullet}) \\ &\quad = \Big(\dim V_i - \operatorname{codim}_{V_i}(A_{a\bullet} \cap V_i)\Big) - \Big(\dim V_i - \operatorname{codim}_{V_i}(A_{a\bullet} \cap A_{x\bullet} \cap V_i)\Big) \\ &\quad = \operatorname{codim}_{V_i}(A_{a\bullet} \cap A_{x\bullet} \cap V_i) - \operatorname{codim}_{V_i}(A_{a\bullet} \cap V_i) \leq d_0 - i - 1.\end{aligned}$$

https://doi.org/10.1017/S0305004125101771 Published online by Cambridge University Press

We may now repeat the same final step of the previous case, with the same choice of $r = \lceil 2\log_p(4\varepsilon^{-1})\rceil$. Take $y_1, \dots, y_r$ uniformly and independently at random from V_{i+1} and let $U_{i+1} = U_i \cap A_{\bullet y_1} \cap \dots \cap A_{\bullet y_r}$. By the calculations in the previous step, there exists a choice of $y_1, \dots, y_r$ such that $\operatorname{codim}_{V_{i+1}}(A_{x\bullet} \cap V_{i+1}) \le d_0 - i - 1$ for at least $\left(1 - \frac{1}{2}\varepsilon\right)|U_{i+1}|$ elements $x \in U_{i+1}$. In this case, the codimension bounds are $\operatorname{codim}_{U_i} U_{i+1} \le rd_0 = O(d_0 \log_p(2\varepsilon^{-1}))$ and $\operatorname{codim}_{V_i} V_{i+1} \le d_0$.

We are now in position to prove Theorem 1·5.

Proof of Theorem 1·5. Let A be the given transverse set of density δ. Apply Lemma 2·10 with parameter $\varepsilon = 2^{-80}p^{-16\lceil \log_p(10^4\delta^{-4})\rceil}$ to get a positive integer $d \le \lceil \log_p(10^4\delta^{-4})\rceil$ and subspaces $U \le G$ and $V \le H$ such that:

(i) $\operatorname{codim}_G U \le O(\log_p^3(2\delta^{-1}))$ and $\operatorname{codim}_H V \le O(\log_p^2(2\delta^{-1}))$;

(ii) for at least $(1-\varepsilon/3)|U|$ of elements $x \in U$ we have $|A_{x\bullet} \cap V| = p^{-d}|V|$; and

(iii) for at least $(1-\varepsilon/3)|U|^2$ of pairs $(x_1, x_2) \in U^2$ we have $|A_{x_1\bullet} \cap A_{x_2\bullet} \cap V| = p^{-2d}|V|$.

Note that $\varepsilon \le 2^{-80}p^{-16d} < 10^{-32}p^{-8d}$. Fix an inner product $\cdot$ on V and, for each $x \in U$, define subspace $V_x = (A_{x\bullet} \cap V)^\perp$, where $^\perp$ indicates orthogonal complement in V with respect to $\cdot$. As A is a transverse set, the collection $(V_x)_{x\in U}$ is a linear system of subspaces. By properties (ii) and (iii) we have that $|V_x| = p^d$ holds for at least $(1-\varepsilon/3)|U|$ of elements $x \in U$, and three equalities $|A_{x_1\bullet} \cap V| = p^{-d}|V|$, $|A_{x_2\bullet} \cap V| = p^{-d}|V|$ and $|A_{x_1\bullet} \cap A_{x_2\bullet} \cap V| = p^{-2d}|V|$ hold simultaneously for at least $(1-\varepsilon)|U|^2$ of pairs $(x_1, x_2) \in U^2$. These three equalities together imply that

$$|V_{x_1} \cap V_{x_2}| = \frac{|V_{x_1}||V_{x_2}|}{|V_{x_1} + V_{x_2}|} = \frac{\frac{|V|}{|A_{x_1\bullet}\cap V|}\frac{|V|}{|A_{x_2\bullet}\cap V|}}{\frac{|V|}{|(V_{x_1}+V_{x_2})^\perp|}} = \frac{|V||(V_{x_1}+V_{x_2})^\perp|}{p^{-2d}|V|^2}$$
$$= \frac{|V_{x_1}^\perp \cap V_{x_2}^\perp|}{p^{-2d}|V|} = \frac{|A_{x_1\bullet} \cap A_{x_2\bullet} \cap V|}{p^{-2d}|V|} = 1,$$

hence $V_{x_1} \cap V_{x_2} = \{0\}$.

We may apply Proposition 2·7 to linear system of subspaces $(V_x)_{x\in U}$ inside V to find a bilinear map $\Phi\colon U \times \mathbb{F}_p^d \to V$ and a linear map $\Psi\colon \mathbb{F}_p^d \to V$ such that $\{\Phi(x,\lambda) + \Psi(\lambda)\colon \lambda \in \mathbb{F}_p^d\} = V_x$ holds for at least $(1 - 200p^d\sqrt[8]{\varepsilon})|G|$ elements $x \in U$. By our choice of ε, we have $200p^d\sqrt[8]{\varepsilon} \le \frac{1}{3}p^{-d}$. Hence, there exists a set $X \subseteq U$, of size $|X| \ge (1-p^{-d}/3)|U|$, such that $\{\Phi(x,\lambda) + \Psi(\lambda)\colon \lambda \in \mathbb{F}_p^d\} = (A_{x\bullet} \cap V)^\perp$ holds for all $x \in X$.

Taking orthogonal complements in V, we get $\{\Phi(x,\lambda) + \Psi(\lambda)\colon \lambda \in \mathbb{F}_p^d\}^\perp = A_{x\bullet} \cap V$. Let $V' = V \cap (\operatorname{Im}\Psi)^\perp$, and let $\beta_i(x,y) = \Phi(x,e_i)\cdot y$ for $i \in [d]$, where $e_1, \dots, e_d$ is the standard basis of $\mathbb{F}_p^d$. Then for each $x \in X$, we have $\{y \in V'\colon \beta(x,y) = 0\} \subseteq A_{x\bullet}$. We claim that in fact

$$\{(x,y) \in U \times V'\colon \beta(x,y) = 0\} \subseteq A.$$

To see this, take any $(x,y) \in U \times V'$ such that $\beta(x,y) = 0$. The subspace $\{u \in U\colon \beta(u,y) = 0\}$ has codimension at most d inside U. As $|X| \ge (1-p^{-d}/3)|U|$, we conclude that $|X \cap \{u \in U\colon \beta(u,y) = 0\}| \ge (2/3)|\{u \in U\colon \beta(u,y) = 0\}|$. Thus, as $\beta(x,y) = 0$, we have that $(X \cap \{u \in U\colon \beta(u,y) = 0\}) + x \subseteq \{u \in U\colon \beta(u,y) = 0\}$ as well, so

https://doi.org/10.1017/S0305004125101771 Published online by Cambridge University Press

$(X \cap \{u \in U\colon\ \beta(u, y) = 0\}) + x$ and $(X \cap \{u \in U\colon\ \beta(u, y) = 0\})$ intersect, at some element u, say. Thus, $u, u - x \in X$ and $\beta(u, y) = \beta(u - x, y) = 0$. By properties of elements in X, we have that $(u, y), (u - x, y) \in A$. But A is a transverse set, so $(x, y) \in A$, as desired.

Finally, let us remark that terms of the form $\log_p(2\delta^{-1})$ in the bounds above can be replaced by more esthetically pleasing $\log_p(\delta^{-1})$. This stems from the fact that if the density δ is at least $15/16$, then our transverse set A is in fact the full product $G \times H$. We leave this simple observation as an exercise to the reader.

Acknowledgements. This work was supported by the Ministry of Science, Technological Development and Innovation of the Republic of Serbia through the Mathematical Institute of the Serbian Academy of Sciences and Arts. I would like to thank the anonymous referee for a careful reading and helpful comments.

REFERENCES

[1] P.-Y. Bienvenu and T.H. Lê. A bilinear Bogolyubov theorem. *European J. Combin.* **77** (2019), 102–113.
[2] P.-Y. Bienvenu and T.H. Lê. Linear and quadratic uniformity of the Möbius function over $\mathbb{F}_q[t]$. *Mathematika* **65** (2019), 505–529.
[3] P.-Y. Bienvenu, D. González-Sánchez and Á.D. Martínez. A note on the bilinear Bogolyubov theorem: transverse and bilinear sets. *Proc. Amer. Math. Soc.* **148** (2020), 23–31.
[4] N. Bogolioùboff. Sur quelques propriétés arithmétiques des presque-périodes. *Ann. Chaire Phys. Math. Kiev* **4** (1939), 185–205.
[5] W.T. Gowers and L. Milićević. A bilinear version of Bogolyubov's theorem. *Proc. Amer. Math. Soc.* **148** (2020), 4695–4704.
[6] W.T. Gowers and L. Milićević. A quantitative inverse theorem for the U^4 norm over finite fields. Preprint: ArXiv:1712.00241 (2017).
[7] W.T. Gowers and L. Milićević. An inverse theorem for Freiman multi-homomorphisms. Preprint: ArXiv:2002.11667 (2020).
[8] K. Hosseini and S. Lovett. A bilinear Bogolyubov-Ruzsa lemma with poly-logarithmic bounds. *Discrete Anal.*, paper no. 10 (2019), 1–14.
[9] L. Milićević. Approximate quadratic varieties. Preprint: ArXiv:2308.12881 (2023).

https://doi.org/10.1017/S0305004125101771 Published online by Cambridge University Press