Ranges of polynomials control degree ranks of Green and Tao over finite prime fields[*]

Thomas Karam[†1]

[1]Mathematical Institute, University of Oxford, OX26GG United Kingdom

Abstract

Let p be a prime, let $1 \leq t < d < p$ be integers, and let S be a non-empty subset of $\mathbb{F}_p$. We establish that if a polynomial $P : \mathbb{F}_p^n \to \mathbb{F}_p$ with degree d is such that the image $P(S^n)$ does not contain the full image $A(\mathbb{F}_p)$ of any non-constant polynomial $A : \mathbb{F}_p \to \mathbb{F}_p$ with degree at most t, then P coincides on S^n with a polynomial that in particular has bounded degree-$\lfloor d/(t+1) \rfloor$-rank in the sense of Green and Tao. Similarly, we prove that if the assumption holds even for $t = d$, then P coincides on S^n with a polynomial determined by a bounded number of coordinates.

Throughout this paper, the letter n will always denote a positive integer, and all our statements will be uniform in n. A full version of the present paper may be found at [7].

1 Degree ranks and ranges of polynomials

A landmark result of Green and Tao proved in 2007 [3] states that over a finite prime field $\mathbb{F}_p$ for some prime p, a multivariate polynomial with degree $1 \leq d < p$ that is not approximately equidistributed can be expressed as a function of a bounded number of polynomials each with degree at most $d-1$. More formally, we have the following statement.

Theorem 1 ([3], Theorem 1.7). *Let p be a prime, and let $1 \leq d < p$ be an integer. Then there exists a function $K_{p,d} : (0,1] \to \mathbb{N}$ such that for every $\epsilon > 0$, if $P : \mathbb{F}_p^n \to \mathbb{F}_p$ is a polynomial with degree d satisfying $|\mathbb{E}_{x \in \mathbb{F}_p^n} \omega_p^{sP(x)}| \geq \epsilon$ for some $s \in \mathbb{F}_p^*$, then there exist $k \leq K_{p,d}(\epsilon)$, polynomials $P_1, \dots, P_k : \mathbb{F}_p^n \to \mathbb{F}_p$ with degree at most $d-1$ and a function $F : \mathbb{F}_p^k \to \mathbb{F}_p$ satisfying*

$$P = F(P_1, \dots, P_k).$$

It has been known since at least the works of Janzer [5] and Milićević [9] that the conclusion can be made qualitatively more precise. Before stating this strengthening, let us define a notion of degree-d rank for polynomials.

Definition 2. *Let $\mathbb{F}$ be a field, and let $P : \mathbb{F}^n \to \mathbb{F}$ be a polynomial. Let $d \geq 1$ be an integer.*

We say that a polynomial P has degree-d rank at most 1 *if we can write P as a product of polynomials each with degree at most d.*

The degree-d rank *of P is defined to be the smallest nonnegative integer k such that there exist polynomials $P_1, \dots, P_k$ each with degree-d rank at most 1, with degree at most the degree of P, and satisfying*

$$P = P_1 + \cdots + P_k.$$

We denote this quantity by $\operatorname{rk}_d P$.

[*] The full version of this work can be found in [7].

[†] Email: thomas.karam@maths.ox.ac.uk. Supported by ERC Advanced Grant 883810.

The zero polynomial in particular has degree-d rank equal to 0 for all d, and constant polynomials have degree-d rank at most 1 for all d.

We define this notion of degree-d rank in this way as doing so will be convenient for us, but it is worth pointing out that in the original paper [3] of Green and Tao, the notion referred to as the degree-d rank was slightly different: for instance the degree-$(d-1)$ rank was the largest possible k in Theorem 1. Nonetheless, it follows immediately from the definitions that for every polynomial $P : \mathbb{F}_p^n \to \mathbb{F}_p$ and every positive integer d, the degree-d rank of P in the sense of Green and Tao is at most d times the degree-d rank of P in our sense. Therefore, proving that a polynomial has bounded degree-d rank in our sense implies showing that it has bounded degree-d rank in the sense of Green and Tao.

The main qualitative refinement shown in the papers of Janzer [5] and Milićević [9] is that there exists some function $H_{p,d} : (0,1] \to \mathbb{N}$ such that under the assumptions of Theorem 1, we can find $k \le H_{p,d}(\epsilon)$ and polynomials $Q_1, R_1, \dots, Q_k, R_k$ satisfying

$$\deg Q_i, \deg R_i \le d-1 \text{ and } \deg Q_i + \deg R_i \le d$$

for each $i \in [k]$ and such that

$$P = Q_1 R_1 + \cdots + Q_k R_k.$$

In other words, it was shown that

$$\operatorname{rk}_{d-1} P \le H_{p,d}(\epsilon).$$

This is a bound on the degree-$(d-1)$ rank of P, and the numerous developments which arose out of Theorem 1 have to our knowledge entirely or almost entirely focused on the degree-$(d-1)$ rank of P: some extended the range of validity of the results (Kaufman and Lovett [8], Bhowmick and Lovett [2]), and others improved the quantitative bounds on the degree-$(d-1)$ rank, through the closely related question of comparing the partition rank to the analytic rank of tensors (Janzer [5], Milićević [9], Adiprasito, Kazhdan and Ziegler [1], Moshkovitz and Cohen [10], [11], Moshkovitz and Zhu [12]).

For the purposes of studying approximate equidistribution of polynomials this is unsurprising, since the notion of degree-$(d-1)$ rank is indeed by far the most relevant: for instance a random polynomial of the type

$$x_1 Q(x_2, \dots, x_n)$$

with $\deg Q = d-1$ has high degree-$(d-2)$ rank but is nonetheless not approximately equidistributed, since the probability that it takes the value 0 is approximately $2/p - 1/p^2 > 1/p$.

Rather than focus on the fact that for a degree-d polynomial, lack of equidistribution implies bounded degree-$(d-1)$ rank, we may ask for analogues of this statement involving much stronger properties in the assumption and in the conclusion. Correspondingly, the main motivations of this paper are twofold. In one direction, we ask what can be deduced about polynomials for which we know much more than lack of equidistribution. What can we say if we know that a polynomial does not take every value of $\mathbb{F}_p$, or has a smaller range still, in a sense to be made precise ? In the other direction, we can ask, for a fixed integer $1 \le e \le d-1$, whether there are any properties of the distribution of the values of a polynomial which would guarantee that its degree-e rank is bounded above. We will contribute to both directions simultaneously, by showing that if a polynomial P does not have full range, then it must have bounded degree-e rank, for some integer e that is determined by the degree of P and by the smallest degree of a non-constant *one-variable* polynomial that has a range contained in the range of P.

Theorem 3. *Let p be a prime, and let $1 \le t \le d < p$ be integers. There exists a positive integer $\gamma(p,d,t)$ such that the following holds. Let $P : \mathbb{F}_p^n \to \mathbb{F}_p$ be a polynomial with degree at most d. Assume that the image $P(\mathbb{F}_p^n)$ does not contain the image of $\mathbb{F}_p$ by any non-constant polynomial $\mathbb{F}_p \to \mathbb{F}_p$ with degree at most t.*

1. *If $t \le d-1$, then P has degree-$\lfloor d/(t+1) \rfloor$-rank at most $\gamma(p,d,t)$.*

2. If $t = d$ then P is a constant polynomial.

The value $\lfloor d/(t+1) \rfloor$ in the degree of the rank in Theorem 3 is optimal in general, as the following example shows.

Example 4. *Let p be a prime, let $1 \leq d < p$ be an integer, let $t, u \geq 1$ be integers such that $tu \leq d$. If Q is a random polynomial $Q : \mathbb{F}_p^n \to \mathbb{F}_p$ with degree u, then Q^t has degree at most d, the image $Q(\mathbb{F}_p^n)$ is contained in the set $\{y^t : y \in \mathbb{F}_p\}$ of t-th power residues mod-p, but the degree-$(u-1)$-rank of P is usually arbitrarily large as n tends to infinity, even if it is taken in the sense of Green and Tao.*

The last part follows from a counting argument: as n tends to infinity there are $p^{O(n^{u-1})}$ polynomials $\mathbb{F}_p^n \to \mathbb{F}_p$ with degree at most $u-1$, so for every $k \geq 1$, the number of polynomials of the type $F(P_1, \dots, P_k)$ with $P_1, \dots, P_k$ with degree at most $u-1$ and $F : \mathbb{F}_p^k \to \mathbb{F}_p$ a function is at most $p^{p^k} p^{O(kn^{u-1})} = p^{O(kn^{u-1})}$, whereas there are $p^{\Omega(kn^u)}$ polynomials $\mathbb{F}_p^n \to \mathbb{F}_p$ with degree u and hence at least $1/t$ times as many polynomials of the type Q^t above.

Powers of polynomials are not the only simple examples of polynomials that do not have full range in general. They can instead be viewed as a special case of a broader class of examples that arises from composition with a one-variable polynomial.

Example 5. *Let p be a prime, let $1 \leq d < p$ be an integer, let $t, u \geq 1$ be integers such that $tu \leq d$. If $Q : \mathbb{F}_p^n \to \mathbb{F}_p$ is a polynomial with degree u, and $A : \mathbb{F}_p \to \mathbb{F}_p$ is a polynomial with degree t, then the polynomial $A \circ Q$ has degree at most d, and the image $A \circ Q(\mathbb{F}_p^n)$ is contained in the image $A(\mathbb{F}_p)$.*

We stress that the main result from the approximate equidistribution regime will itself be an important black box that we will use in our proof of Theorem 3.

2 Variables with restricted range

We shall in fact prove results in a more general setting than that of Theorem 3, where we allow the assumption to be on the image $P(S^n)$ for some non-empty subset S of $\mathbb{F}_p$ rather than on the whole image $P(\mathbb{F}_p^n)$. On a first reading the set S may be taken to be $\{0, 1\}$. In the setting of restrictions to S^n, the approximate equidistribution statement was proved by Gowers and the author in [4]. Before stating it, let us recall from that paper two points to be aware of regarding restrictions of polynomials to S^n.

The first is that whereas an affine polynomial is either constant or perfectly equidistributed on $\mathbb{F}_p^n$, there is already something to say about the distribution of an affine polynomial P on S^n for general non-empty S: if $S \neq \mathbb{F}_p$ and P depends only on one coordinate, then $P(S^n)$ is not even the whole of $\mathbb{F}_p$. As a simple Fourier argument however shows ([4], Proposition 2.2), an affine polynomial depending on many coordinates is approximately equidistributed on S^n, provided that S contains at least two elements. The second is that we may no longer hope to conclude in general that a polynomial with degree d which is not approximately equidistributed on S^n must itself have bounded degree-$(d-1)$ rank: for instance, the polynomial

$$\sum_{i=1}^{n} x_i^2 - x_i$$

has degree-1 rank equal to n, but only takes the value 0 on $\{0, 1\}^n$ and is in particular not approximately equidistributed on $\{0, 1\}^n$. Nonetheless, the zero polynomial, with which this polynomial coincides on $\{0, 1\}^n$, itself has degree-1 rank equal to 0.

These two remarks motivate an extension of Definition 2.

Definition 6. *Let $\mathbb{F}$ be a field, and let $P : \mathbb{F}^n \to \mathbb{F}$ be a polynomial.*

The degree-0 rank *of P is defined to be the smallest nonnegative integer k such that we can write P as a linear combination of at most k monomials. We denote this quantity by* $\operatorname{rk}_0 P$.

If d is a nonnegative integer and S is a non-empty subset of $\mathbb{F}$ then we define the degree-d rank of P with respect to S *as the smallest value of* $\mathrm{rk}_d(P-P_0)$*, where the minimum is taken over all polynomials P_0 with degree at most the degree of P and satisfying $P_0(S^n)=\{0\}$. We denote this quantity by* $\mathrm{rk}_{d,S}\,P$*.*

We now recall a slight weakening of the main result of [4], Theorem 1.4 from that paper. (Although the full statement of that theorem is slightly more precise, the formulation below is slightly simpler to use and suffices for the purposes of the present paper.)

Theorem 7. *Let p be a prime, let $1 \leq d < p$ be an integer, and let S be a non-empty subset of $\mathbb{F}_p$. There exists a function $H_{p,d,S} : (0,1] \to \mathbb{N}$ such that for every $\epsilon > 0$, if $P : \mathbb{F}_p^n \to \mathbb{F}_p$ is a polynomial with degree d satisfying $|\mathbb{E}_{x\in S^n}\omega_p^{sP(x)}| \geq \epsilon$ for some $s \in \mathbb{F}_p^*$, then $\mathrm{rk}_{d-1,S}\,P \leq H_{p,d,S}(\epsilon)$.*

We note that if S has size 1, then Theorem 7 as well as many of the new results of the present paper hold for immediate reasons: the set S^n then also has size 1, so every polynomial coincides on S^n with a constant polynomial, so has degree-d rank at most 1 for every d.

When S is not the whole of $\mathbb{F}_p$, one important difference between the sets $\mathbb{F}_p^n$ and S^n is that the former is invariant under linear transformations, whereas the latter is not. We have already discussed one effect on this: the fact that x_1 does not take every value of $\mathbb{F}_p$ whereas $x_1+\cdots+x_n$ is approximately equidistributed for n large. The role of coordinates as opposed to general degree-1 polynomials will manifest itself further in the proofs and in the main results of this paper. For this purpose let us make one last definition.

Definition 8. *Let p be a prime, and let $P : \mathbb{F}_p^n \to \mathbb{F}_p$ be a polynomial.*

For each $i \in [n]$, we say that P depends on x_i if the coordinate x_i arises in some monomial of P.

*For k nonnegative integer, we will say that P is k-*determined *if it depends on at most k coordinates.*

3 Statements of main results

Using Theorem 7 as a black box we will prove the following analogue of Theorem 3, where the assumption on the image is now on $P(S^n)$ rather than on $P(\mathbb{F}_p^n)$. The following theorem is the main result that we shall prove in the present paper.

Theorem 9. *Let p be a prime, let $1 \leq t \leq d < p$ be integers and let S be a non-empty subset of $\mathbb{F}_p$. Then there exists a positive integer $C(p,d,t)$ such that the following holds. Let $P : \mathbb{F}_p^n \to \mathbb{F}_p$ be a polynomial with degree at most d. Assume that $P(S^n)$ does not contain the image of $\mathbb{F}_p$ by any non-constant polynomial $\mathbb{F}_p \to \mathbb{F}_p$ with degree at most t.*

1. *If $t \leq d-1$, then P coincides on S^n with a polynomial that has degree-$\lfloor d/(t+1)\rfloor$-rank at most $C(p,d,t)$ and has degree at most d.*
2. *If $t = d$ then P coincides on S^n with a linear combination of at most $C(p,d,t)$ monomials with degrees at most d.*

Equivalently, in both cases we have

$$\mathrm{rk}_{\lfloor \frac{d}{t+1}\rfloor,S}\,P \leq C(p,d,t).$$

The optimal bounds in Theorem 9 and in several of our other statements involving the set S may depend on the choice of S. However, to avoid heavy notation we will at many places avoid making this dependence explicit. (We may safely do so, since for each prime p there are only finitely many subsets of $\mathbb{F}_p$).

Let us look at the extreme cases of item 1 from Theorem 9, and at a situation where they are both simultaneously realised.

Corollary 10. *Let p be a prime, let $1 \leq t \leq d < p$ be integers and let S be a non-empty subset of $\mathbb{F}_p$. Let $P : \mathbb{F}_p^n \to \mathbb{F}_p$ be a polynomial with degree at most d. Let $C_{(i)}(p,d) = C(p,d,1)$ and let $C_{(ii)}(p,d) = C(p,d,\lfloor d/2 \rfloor)$.*

(i) If $P(S^n) \neq \mathbb{F}_p$ then $\operatorname{rk}_{\lfloor d/2 \rfloor, S} P \leq C_{(i)}(p,d)$.

(ii) If $P(S^n)$ does not contain the image of any non-constant polynomial $\mathbb{F}_p \to \mathbb{F}_p$ with degree at most $\lfloor d/2 \rfloor$ then $\operatorname{rk}_{1,S} P \leq C_{(ii)}(p,d)$.

(iii) If $d = 3$ and $P(S^n) \neq \mathbb{F}_p$, then $\operatorname{rk}_{1,S} P \leq C_{(i)}(p,3) = C_{(ii)}(p,3)$.

Proof. Items (i) and (ii) follow from taking $t = 1$ and $t = \lfloor d/2 \rfloor$ in Theorem 9 respectively. Item (iii) follows from either of the items (i) and (ii). □

We now turn our attention to the case of degree-2 polynomials. Throughout the paper, we will write Q_p for the set $\{y^2 : y \in \mathbb{F}_p\}$ of mod-p quadratic residues. Provided that $p \geq 3$, this set has size $\frac{p+1}{2}$ and is in particular not the whole of $\mathbb{F}_p$. We say that a subset of $\mathbb{F}_p$ is an *affine translate* of Q_p if it can be written as $aQ_p + b$ for some $a \in \mathbb{F}_p^*$ and some $b \in \mathbb{F}_p$. In light of the preceding discussion we can formulate three basic constructions of a degree-2 polynomial P such that $P(S^n) \neq \mathbb{F}_p$.

(i) A polynomial of the type $A \circ L$ for some affine polynomial $L : \mathbb{F}_p^n \to \mathbb{F}_p$ and some degree-2 polynomial $A : \mathbb{F}_p \to \mathbb{F}_p$. (Equivalently, the sum of a multiple of L^2 and of a constant.)

(ii) A polynomial that depends only on a small number $r < \log p / \log |S|$ of coordinates, since $P(S^n)$ then necessarily has size at most $|S|^r$.

(iii) A polynomial that vanishes on S^n and has degree at most 2.

The first item of the following result can be interpreted as a converse which says that every example arises as a sum of these three examples, letting aside the value of the bound on the number of coordinates in the second example.

Proposition 11. *There exists an absolute constant $\kappa > 0$ such that the following holds. Let p be a prime, and let S be a non-empty subset of $\mathbb{F}_p$. Let $P : \mathbb{F}_p^n \to \mathbb{F}_p$ be a polynomial with degree 2.*

1. *If $P(S^n) \neq \mathbb{F}_p$, then there exists an affine polynomial $L : \mathbb{F}_p^n \to \mathbb{F}_p$, a degree-2 polynomial $A : \mathbb{F}_p \to \mathbb{F}_p$, and a κp^{15}-determined polynomial J with degree at most 2 such that P coincides on S^n with $A \circ L + J$. (Equivalently, with $AL^2 + J$ for some $A \in \mathbb{F}_p$, with J changed by a constant.)*

2. *If furthermore $P(S^n)$ does not contain any affine translate of Q_p, then P coincides on S^n with a κp^{15}-determined polynomial that has degree at most 2.*

Item 1 from Proposition 11 is significantly stronger than the conclusion that item 1 from Theorem 9 gives in the corresponding case $d = 2$ and $t = 1$: the latter is merely that P has bounded degree-1 rank with respect to S, which we already know by Theorem 7. The proof of Proposition 11 will instead use different techniques which do not appear to generalise well to higher-degree polynomials.

In the more general case where P has general degree $2 \leq d \leq p-1$, one may ask whether just as with item 1 from Proposition 11, it is the case that provided that $P(S^n) \neq \mathbb{F}_p$ we can always obtain a decomposition $P = A \circ Q + J$ with $Q : \mathbb{F}_p^n \to \mathbb{F}_p$, $A : \mathbb{F}_p \to \mathbb{F}_p$ polynomials satisfying $\deg Q \deg A \leq d$, $\deg A \geq 2$ and with J a polynomial determined by a bounded number of coordinates and with degree at most d. This is however not the case in general, as a wider class of examples comes in: for instance, if $d = p-1$, then the polynomial $A : x \to x^{p-1}$ satisfies $A(\mathbb{F}_p) = \{0,1\}$, so if $L_1, \dots, L_{p-2}$ are arbitrary affine polynomials then the image of $\mathbb{F}_p^n$ by the polynomial

$$P = A \circ L_1 + \dots + A \circ L_{p-2}$$

does not contain $p-1$. For $d=2$, such a situation cannot occur, because the Cauchy-Davenport theorem, which will play some role in the proof of Proposition 11, shows that the sumset of any two affine translates of Q_p is the whole of $\mathbb{F}_p$. (However, this is by no means the only or even the main specificity of the case $d=2$ that allows us to say more there than for general $d<p$.)

4 Techniques for the proof of Theorem 9

The basic strategy which we will use to prove Theorem 9 will be essentially as follows: because P has degree d and is not approximately equidistributed, Theorem 7 shows that P coincides on S^n with some polynomial with degree at most d and of the type

$$M \circ (P_1, \dots, P_k)$$

where k is bounded and M is some polynomial. One of the following is always true: either the polynomials $P_1, \dots, P_k$ are approximately jointly equidistributed, in which case the image $(P_1, \dots, P_k)(S^n)$ is the same as if the polynomials $P_1, \dots, P_k$ were jointly equidistributed, or they are not, in which case at least one non-trivial linear combination of the polynomials $P_1, \dots, P_k$ has bounded degree-$(d'-1)$ rank with respect to S, where

$$d' = \max(\deg P_1, \dots, \deg P_k),$$

and we may hence without loss of generality assume that P coincides on S^n with some polynomial with degree at most d and of the type

$$M' \circ (P_1, \dots, P_{k-1}, Q_1, \dots, Q_{k'})$$

where k' is bounded, $Q_1, \dots, Q_{k'}$ are polynomials with degree strictly smaller than the degree of P_k, and M' is some polynomial. This second step, in turn, can only be performed a bounded number of times, which will conclude the argument.

References

[1] K. Adiprasito, D. Kazhdan, and T. Ziegler, *On the Schmidt and analytic ranks for trilinear forms*, arXiv:2102.03659 (2021).

[2] A. Bhowmick and S. Lovett, *Bias vs structure of polynomials in large fields, and applications in effective algebraic geometry and coding theory*, IEEE Trans. Inf. Theory, arXiv:1506.02047 (2015).

[3] B. Green and T. Tao, *The distribution of polynomials over finite fields, with applications to the Gowers norms.* Contr. Discr. Math., **4** (2009), no. 2, 1-36.

[4] W. T. Gowers and T. Karam, *Equidistribution of high-rank polynomials with variables restricted to subsets of* $\mathbb{F}_p$, arXiv:2209.04932 (2022).

[5] O. Janzer, *Polynomial bound for the partition rank vs the analytic rank of tensors*, Discrete Anal. **7** (2020), 1-18.

[6] T. Karam, *High-rank subtensors of high-rank tensors*, arXiv:2207.08030 (2022).

[7] T. Karam, *Ranges of polynomials control degree ranks of Green and Tao over finite prime fields*, arXiv:2305.11088 (2023).

[8] T. Kaufman and S. Lovett, *Worst case to average case reductions for polynomials*, 49th Annual IEEE Symposium on Foundations of Computer Science (2008), 166-175.

[9] L. Milićević, *Polynomial bound for partition rank in terms of analytic rank*, Geom. Funct. Anal. **29** (2019), 1503-1530.

[10] A. Cohen and G. Moshkovitz, Structure vs. randomness for bilinear maps, Discrete Anal. **12** (2022).

[11] A. Cohen and G. Moshkovitz, Partition and analytic rank are equivalent over large fields, (2022).

[12] G. Moshkovitz and D. G. Zhu, *Quasi-linear relation between partition and analytic rank*, arXiv:2211.05780 (2022).