ALGEBRAIC ASPECTS OF THE POLYNOMIAL LITTLEWOOD–OFFORD PROBLEM

ZHIHAN JIN, MATTHEW KWAN, LISA SAUERMANN, AND YITING WANG, WITH AN APPENDIX BY MATTHEW KWAN, ASHWIN SAH AND MEHTAAB SAWHNEY

ABSTRACT. Consider a degree-*d* polynomial $f(\xi_1, \ldots, \xi_n)$ of independent Rademacher random variables ξ_1, \ldots, ξ_n . To what extent can $f(\xi_1, \ldots, \xi_n)$ concentrate on a single point? This is the so-called *polynomial Littlewood–Offord problem*. A nearly optimal bound was proved by Meka, Nguyen and Vu: the point probabilities are always at most about $1/\sqrt{n}$, unless *f* is "close to the zero polynomial" (having only $o(n^d)$ nonzero coefficients).

In this paper we prove several results supporting the general philosophy that the Meka–Nguyen–Vu bound can be significantly improved unless f is "close to a polynomial with special algebraic structure", drawing some comparisons to phenomena in analytic number theory. In particular, one of our results is a corrected version of a conjecture of Costello on multilinear forms (in an appendix with Ashwin Sah and Mehtaab Sawhney, we disprove Costello's original conjecture).

1. INTRODUCTION

Anticoncentration inequalities are an important class of probabilistic inequalities, giving upper bounds on the probability that a random variable falls in a small interval or is equal to a particular value (i.e., they provide limits on the extent to which a random variable "concentrates"). Such inequalities are ubiquitous, playing an essential behind-the-scenes role in a wide range of different types of probabilistic arguments.

In particular, an important direction is the (polynomial) *Littlewood–Offord problem*, which concerns anticoncentration of low-degree polynomials of independent random variables (making minimal assumptions about the structure of the polynomials in question). A representative theorem, due to Meka, Nguyen and Vu [40], is as follows. (The *Rademacher* distribution is the uniform distribution on $\{-1, 1\}$).

Theorem 1.1. Fix $d \ge 1$ and $\varepsilon > 0$, and let $\mathbb{F} \in \{\mathbb{R}, \mathbb{C}\}$ be the field of real or complex numbers. Let n be sufficiently large (in terms of ε, d). Let $f \in \mathbb{F}[x_1, \ldots, x_n]$ be an n-variable polynomial of degree d. Then, at least one of the following holds:

A1 f has at most εn^d nonzero coefficients, or

A2 letting $\xi_1, \ldots, \xi_n \in \{-1, 1\}$ be *i.i.d.* Rademacher random variables, we have

$$\sup_{z \in \mathbb{F}} \mathbb{P}[f(\xi_1, \dots, \xi_n) = z] \le n^{-1/2 + \varepsilon}.$$

In other words, polynomials of n independent random variables have their point probabilities bounded by about $1/\sqrt{n}$, unless they are "close to the zero polynomial" (note that we may take $\varepsilon > 0$ to be arbitrarily small¹, strengthening both **A1** and **A2**, though this forces us to take large n). To see that one cannot hope for a bound stronger than about $1/\sqrt{n}$, consider the polynomial $f(x_1, \ldots, x_n) =$ $(x_1 + \cdots + x_n)^d$. Indeed, then $\sup_{z \in \mathbb{F}} \mathbb{P}[f(\xi_1, \ldots, \xi_n) = z] \ge \sup_{z \in \mathbb{F}} \mathbb{P}[x_1 + \cdots + x_n = z] = {n \choose \lfloor n/2 \rfloor} \cdot 2^{-n}$ is on the order of $1/\sqrt{n}$ (up to constant factors).

Remark 1.2. Littlewood–Offord theorems are usually stated with ξ_1, \ldots, ξ_n being i.i.d. Rademacher (as in Theorem 1.1 above), but it is often not hard to deduce analogous results where ξ_1, \ldots, ξ_n are allowed to have arbitrary distributions, as long as they are independent. Indeed, in some sense the Rademacher case is the "hardest case": other cases can be reduced to this case with simple coupling/conditioning arguments (see for example [40, Section 6]).

The first author was supported by SNSF grant 200021-19696. The second and fourth authors were supported by ERC Starting Grant "RANDSTRUCT" No. 101076777. The third author was supported by the DFG Heisenberg Program.

¹There are more precise quantitative questions in this direction, which we will not pursue in this paper.

To give some history for Theorem 1.1 (more can be found in the surveys [47, 61]): the d = 1 (linear) case of Theorem 1.1 first appeared as a lemma in a 1943 paper of Littlewood and Offord [38]² on random algebraic equations. Their bound was improved by Erdős [15] in 1945, obtaining *exact* bounds for the d = 1 case (the result of this latter paper is now usually called the *Erdős–Littlewood–Offord theorem*, and has been very influential in combinatorics and random matrix theory). Both these papers were for linear forms with real coefficients; complex coefficients were first explicitly considered by Kleitman [31] (though, in the way we have stated Theorem 1.1, the real and complex cases are equivalent).

The higher-degree $(d \ge 2)$ Littlewood–Offord problem was first considered in a 1996 paper of Rosiński and Samorodnitsky [51] (they proved a weaker form of Theorem 1.1 for applications to Lévy chaos), and this direction of study rose to the forefront after a 2006 paper of Costello, Tao and Vu [10] (who proved a weak form of the d = 2 case of Theorem 1.1, as a key ingredient in their work on random symmetric matrices), and a 2013 paper of Razborov and Viola [49] (who proved a weak form of Theorem 1.1 as a tool to study correlation bounds between Boolean functions and polynomials). The d = 2 case of Theorem 1.1 was proved by Costello [9], and Theorem 1.1 for general d was proved by Meka, Nguyen and Vu [40], using a result of Kane [27]. Very recently, an optimal $O(1/\sqrt{n})$ bound in A2 was obtained in the case d = 2 by the second and third authors [34].

It remains a very interesting open problem to obtain sharp $O(1/\sqrt{n})$ bounds in **A2** for $d \ge 3$, but another important question is to study conditions under which one can prove *much stronger* bounds. A natural conjecture along these lines was first made by Costello [9, Conjecture 3], as follows (his paper was about polynomials with complex coefficients, but we have taken the liberty to state his conjecture in both the real and complex cases³).

Conjecture 1.3. Fix $d \ge 2$ and $\varepsilon > 0$, and let $\mathbb{F} \in \{\mathbb{R}, \mathbb{C}\}$. Let n be sufficiently large (in terms of ε, d), and let $f \in \mathbb{F}[x_1, \ldots, x_n]$ be an n-variable polynomial of degree d. Then, at least one of the following holds:

- **B1** there is a reducible⁴ polynomial $g \in \mathbb{F}[x_1, \ldots, x_n]$ with degree at most d, such that f g has at most εn^d nonzero coefficients, or
- **B2** letting $\xi_1, \ldots, \xi_n \in \{-1, 1\}$ be *i.i.d.* Rademacher random variables, we have

$$\sup_{z\in\mathbb{F}}\mathbb{P}[f(\xi_1,\ldots,\xi_n)=z]\leq n^{-d/2+\varepsilon}.$$

To paraphrase, Costello conjectured that either the point probabilities are bounded by about $n^{-d/2}$, or f is close to a product of two lower-degree polynomials. To motivate the bound " $n^{-d/2}$ ", note that if the coefficients of f are bounded integers (and at least a constant fraction of them are nonzero), then the variance of $f(\xi_1, \ldots, \xi_n)$ has order of magnitude n^d , so its standard deviation has order of magnitude $n^{d/2}$. It is reasonable to imagine that, "generically", the probability mass is roughly evenly spread out on the integers within standard-deviation range of the mean. The most obvious way for this "generic" bound to fail is if f factorises into lower-degree polynomials, since when one of the factors is zero, the whole polynomial is zero. Costello conjectured that that a bound of roughly $n^{-d/2}$ is achieved as long as the polynomial is far from factorising (i.e., if it is "robustly irreducible" in the sense of **B1**).

We remark that Conjecture 1.3 falls in the general direction of the *inverse* Littlewood–Offord problem, which is the study of the structural aspects of f that control the concentration behaviour of $f(\xi_1, \ldots, \xi_n)$ (this problem has a long history, which we discuss in more detail in Section 1.1.3). Comparisons can also be made to the *Elekes–Rónyai–Szabó problem* in combinatorial geometry (see [13] for a survey), and the problem of counting integral points on varieties (which is intensively studied in analytic number theory; see e.g. [6]). Indeed, one can consider the general problem of counting the number of zeroes of an *n*-variable polynomial f which lie in a "combinatorial box" $\prod_{i=1}^{n} A_i$. The polynomial Littlewood–Offord problem is concerned with the case where each A_i has size 2 and the dimension n is large, whereas the Elekes–Rónyai–Szabó problem is concerned with the case where n is small and each A_i is large, and in the study of integral points on varieties one is concerned with the specific case where each A_i is a long interval of integers of the form $\{-B, \ldots, B\}$. In the latter two settings, the fundamental question is how

 $^{^{2}}$ A more general result was claimed in a 1939 paper by Doeblin [14], though this is also the subject of a 1958 paper of Kolmogorov [32], which claims that Doeblin's paper did not provide a full proof.

³Unlike Theorem 1.1, here there is a genuine difference between the real and complex cases, as there are real polynomials that factorise over \mathbb{C} but not \mathbb{R} .

 $^{^{4}}$ A polynomial is reducible if it can be factored as the product of two non-constant polynomials. In this paper, we use the convention that the zero polynomial is reducible.

the algebraic properties of the polynomial f control the asymptotics of the number of zeroes in a large box.

The d = 1 case of Conjecture 1.3 trivially coincides with the d = 1 case of Theorem 1.1, and we are also inclined to believe the d = 2 case (we will justify this in Section 1.1, by comparison to related work in analytic number theory), but unfortunately Conjecture 1.3 is too optimistic for $d \ge 3$, as shown by the following proposition.

Proposition 1.4. Fix an integer constant $d \ge 1$. Consider 2d disjoint subsets $I_1, \ldots, I_{2d} \subseteq \{1, \ldots, n\}$, each of size exactly $2\lfloor n/(4d) \rfloor$. For $j \in \{1, \ldots, 2d\}$ let $L_j(x_1, \ldots, x_n) = \sum_{i \in I_j} x_i$, and define the degree-d polynomial

$$f = L_1 \dots L_d - L_{d+1} \dots L_{2d} \in \mathbb{R}[x_1, \dots, x_n].$$

Then, there is $\varepsilon > 0$ (depending only on d) such that:

- **B1** does not hold, even if we take $\mathbb{F} = \mathbb{C}$ (i.e., the real polynomial f is "robustly irreducible", even if we allow factorisations over the complex numbers), and
- for i.i.d. Rademacher random variables $\xi_1, \ldots, \xi_n \in \{-1, 1\}$, we have $\mathbb{P}[f(\xi_1, \ldots, \xi_n) = 0] \ge \varepsilon/n$ (so if $d \ge 3$, then **B**2 does not hold either).

Proposition 1.4 was observed in a conversation between Ashwin Sah, Mehtaab Sawhney and the second author of this paper. The proof proceeds via a random sampling argument restricting to a small subset of the variables where the polynomial essentially factorises, with some careful Ramsey-theoretic arguments to handle non-multilinear terms. The details are provided in Appendix B.

Remark 1.5. Costello [9] also highlighted a special case of Conjecture 1.3 (appearing as [9, Conjecture 2]) for d-multilinear forms (i.e., degree-d polynomials for which the variables are partitioned into d different "types", and every monomial contains exactly one variable of each type). Costello was able to prove this specialised conjecture for bilinear forms (2-multilinear forms), but Proposition 1.4 shows that when d = 3 even this specialised conjecture is false.

Remark 1.6. It seems plausible that with similar methods as in the proof of Proposition 1.4 (but more technical details) one can actually prove that for any irreducible polynomial $F \in \mathbb{F}[X_1, \ldots, X_k]$, polynomials of the form $F(L_1, \ldots, L_k)$ are robustly irreducible in the sense of **B1** (where we fix disjoint sets $I_1, \ldots, I_k \subseteq \{1, \ldots, n\}$ of size about n/k, and take $L_j(x_1, \ldots, x_n) = \sum_{i \in I_j} x_i$ for $j = 1, \ldots, k$). If the polynomial F has integer coefficients and has many zeroes in $([-\sqrt{n}, \sqrt{n}] \cap \mathbb{Z})^k$, then the point probability $\mathbb{P}[F(L_1(\xi_1, \ldots, \xi_n), \ldots, L_k(\xi_1, \ldots, \xi_n)) = 0]$ is large (see also the discussion in Section 1.1.1). So, this would give a large family of counterexamples to Conjecture 1.3.

The purpose of this paper is to open a systematic investigation into the algebraic aspects of the Littlewood–Offord problem, in accordance with the philosophy of Conjecture 1.3. We see there being two main questions to investigate. First, we believe that if **B1** does not hold (i.e., if f is "robustly irreducible"), it should still be possible to at least prove some "power-saving" beyond the general bound in Theorem 1.1.

Conjecture 1.7. In the setting of Conjecture 1.3, either B1 holds, or else

B2' letting $\xi_1, \ldots, \xi_n \in \{-1, 1\}$ be i.i.d. Rademacher random variables, we have

$$\sup_{z\in\mathbb{F}}\mathbb{P}[f(\xi_1,\ldots,\xi_n)=z]\leq n^{-c+\varepsilon}$$

for some c > 1/2 (which does not depend on ε , n or f).

It seems plausible that in fact Conjecture 1.7 holds with c = 1 (this is the best one could hope for, given Proposition 1.4, and in the d = 2 case it would be equivalent to Conjecture 1.3). However, this is probably a very difficult problem: as we will discuss in Section 1.1.1, the general c = 1 case of Conjecture 1.7 seems to be at least as hard as the so-called *affine dimension growth conjecture* in analytic number theory.

Second, it would be of interest to identify conditions under which **B2** can be salvaged.

Question 1.8. In the setting of Conjecture 1.3, what natural assumptions on f are sufficient to guarantee that **B2** holds?

As our first main result, we prove an optimal version of Conjecture 1.7 for *d*-multilinear forms, which can be viewed as a "repaired" version of [9, Conjecture 2] (recalling the discussion in Remark 1.5).

Theorem 1.9. In the setting of Conjectures 1.3 and 1.7, if f is a d-multilinear form, then either **B1** or **B2**' holds with c = 1.

As our next main result, we prove Conjecture 1.7 in the complex quadratic case $(d = 2 \text{ and } \mathbb{F} = \mathbb{C})$.

Theorem 1.10. When d = 2 and $\mathbb{F} = \mathbb{C}$, Conjecture 1.7 holds with c = 13/24.

As our final main result, we show that in the quadratic case (d = 2), one can approach the "generic bound" $1/n = n^{-d/2}$ with an assumption that the quadratic part of f has high rank (this provides a partial answer to Question 1.8).

Theorem 1.11. Fix any $\varepsilon > 0$ and $k \ge 1$, and let $\mathbb{F} \in \{\mathbb{R}, \mathbb{C}\}$. Let n be sufficiently large (in terms of ε, k), and let $f \in \mathbb{F}[x_1, \ldots, x_n]$ be a quadratic n-variable polynomial. Then, at least one of the following holds:

- **C1** there is a quadratic form $g \in \mathbb{F}[x_1, \ldots, x_n]$ with rank⁵ less than $2k^2$, such that f g has at most εn^2 nonzero coefficients, or
- C2 letting $\xi_1, \ldots, \xi_n \in \{-1, 1\}$ be *i.i.d.* Rademacher random variables, we have

$$\sup \mathbb{P}[f(\xi_1,\ldots,\xi_n)=z] \le n^{-1+2/k}$$

In other words, for any $\alpha > 0$, we can prove a bound of the form $n^{-1+\alpha}$ (for large n), as long as f is not close to a quadratic form with rank less than about α^{-2} .

Remark 1.12. It is worth remarking that Theorem 1.11 is reminiscent of a theorem of Gowers and Karam [21], which (rephrased in the language of this paper) implies that for a fixed prime p, any degreed polynomial $f \in \mathbb{Z}[x_1, \ldots, x_n]$, and i.i.d. Rademacher ξ_1, \ldots, ξ_n , if $f(\xi_1, \ldots, \xi_n)$ has "a significant bias mod p", then in a certain sense it is possible to approximate f by a degree-d polynomial of "low rank". This can be interpreted as a "restricted-domain" version of a similar theorem of Green and Tao [22] (qualitatively improved by Janzer [26] and Milićević [41]), which roughly speaking considers the case where ξ_1, \ldots, ξ_n are uniform mod p. However, we do not see a formal connection between Theorem 1.11 and this direction of study.

Remark 1.13. The second and third author [35] previously proved versions of Theorems 1.10 and 1.11 in the setting where f has bounded coefficients, and instead of measuring the "smallness" of a polynomial by counting nonzero coefficients, one considers the sum of the absolute values of the coefficients. This setting makes the problem very different in character (it makes the problem much more analytic than algebraic), and the proofs in [35] have almost nothing in common with the proofs of Theorems 1.10 and 1.11.

We outline the ideas in the proofs of Theorems 1.9 to 1.11 in Section 2. The proofs contain a number of ideas of general interest, including some lemmas related to *property testing* of low-rank tensors and symmetric matrices.

1.1. Further directions, and speculations. It seems that for further progress on Conjecture 1.7 and Question 1.8 (and the d = 2 case of Conjecture 1.3), it might be necessary to integrate the "combinatorial" ideas in this paper with ideas from analytic number theory. Here, we discuss the connection to analytic number theory and also suggest some other directions for further research.

1.1.1. Connections to analytic number theory. Proposition 1.4 (our counterexample to Costello's conjecture) is really a special case of a general observation that if $P \in \mathbb{F}[y_1, \ldots, y_k]$ is any polynomial with $N_P(B)$ zeroes in $\{-B, \ldots, B\}^k$, then one can construct a polynomial $f_P \in \mathbb{F}[x_1, \ldots, x_n]$ of the form $f_P = P(L_1, \ldots, L_k)$ for linear forms $L_1, \ldots, L_k \in \mathbb{F}[x_1, \ldots, x_n]$ such that $\mathbb{P}[f_P(\xi_1, \ldots, \xi_n) = 0]$ scales roughly like $N_P(\sqrt{n})n^{-k/2}$. So, in order to fully understand the algebraic aspects of the polynomial Littlewood–Offord problem, it seems to be necessary to understand how the algebraic features of a polynomial P affect the rate of growth of $N_P(B)$ (which can be interpreted as a certain "density of integral points" on the variety defined by P). To this end, one is forced to confront deep questions in analytic number theory.

⁵Recall that a quadratic form is a homogeneous quadratic polynomial $h \in \mathbb{F}[x_1, \ldots, x_n]$. Any quadratic form can be expressed in matrix form as $h(x) = \vec{x}^T Q \vec{x}$ for some symmetric matrix $Q \in \mathbb{F}^{n \times n}$. The rank of h is defined to be the rank of this matrix Q. Equivalently, the rank of h is the minimum r such that there is a representation $h = \lambda_1 h_1^2 + \cdots + \lambda_r h_r^2$ as a linear combination of squares of homogeneous linear polynomials $h_1, \ldots, h_r \in \mathbb{F}[x_1, \ldots, x_n]$ with coefficients $\lambda_1, \ldots, \lambda_r \in \mathbb{F}$.

Specifically, if we restrict our attention to polynomials of the form f_P , then the bound of shape approximately $1/\sqrt{n}$ in Theorem 1.1 corresponds to the "trivial bound" $N_P(B) = O(B^{k-1})$, which holds for any nonzero polynomial $P \in \mathbb{F}[y_1, \ldots, y_k]$ of fixed degree (see for example [4, Equation (2.3)]). It seems that in order to prove Conjecture 1.7 for polynomials of the form f_P , one would need a powersaving improvement over this "trivial bound" when P is irreducible. Such estimates are indeed available: for an irreducible d-form P, Cohen [8] proved that $N_P(B) \leq B^{k-3/2+o(1)}$, and this was later improved by Pila [48] to $N_P(B) \leq B^{k-2+1/d+o(1)}$ (see also [62]). It is conjectured that $N_P(B) \leq B^{k-2+o(1)}$ for all $d \geq 2$ (which would be best-possible, if true; this is an affine version of the so-called dimension growth conjecture of Heath-Brown [24] for projective varieties). The affine dimension growth conjecture was very recently proved when $d \neq 3$ by Vermeulen [60], building on earlier work of Browning, Heath-Brown and Salberger [5] and Salberger [53]. The methods in all these works are quite different, and it would be very interesting to investigate if any of them can be extended to a full proof of Conjecture 1.7 (not just for polynomials of the form f_P).

Regarding Question 1.8, if f_P has "generic" anticoncentration $n^{-d/2+o(1)}$, this corresponds to a bound of the form $N_P(B) \leq B^{k-d+o(1)}$. This is sometimes called the "probabilistic heuristic" in analytic number theory. It is a very deep open problem to characterise the polynomials P for which such a bound holds, but one interesting sufficient condition is that P is a *d*-form with sufficiently large *Schmidt rank* (this parameter is also sometimes called *h*-invariant or strength), meaning that P cannot be written as a sum of few products of forms of lower degree [55]. We suspect that a version of Conjecture 1.3 along these lines should also hold, as follows.

Conjecture 1.14. For any d, there is $k \in \mathbb{N}$ such that the following holds. In the setting of Conjecture 1.3, either B2 holds or

B1' For some $r \leq k$ there are homogeneous polynomials $g_1, h_1, g_2, h_2, \ldots, g_r, h_r \in \mathbb{F}[x_1, \ldots, x_n]$, each with positive degree, such that $g_1h_1 + g_2h_2 + \cdots + g_rh_r$ has degree at most d, and $f - (g_1h_1 + g_2h_2 + \cdots + g_rh_r)$ has at most εn^d nonzero coefficients.

Note that Theorem 1.11 almost, but not quite, resolves the d = 2 case of Conjecture 1.14. The difference is that in Theorem 1.11, as one wishes to get closer and closer to the generic bound $n^{-d/2+o(1)}$, the rank requirement continues to increase, whereas in Conjecture 1.14 there is a particular rank above which one immediately obtains the generic bound $n^{-d/2+o(1)}$.

In the case where f is a d-multilinear form, the Schmidt rank of f coincides with the so-called *partition* rank of the tensor of coefficients of f. This notion has recently been intensively studied in theoretical computer science and additive combinatorics (see for example [7, 26, 36, 39, 41-44]), and it would likely be much more tractable (and still very interesting) to prove the special case of Conjecture 1.14 where f is a d-multilinear form.

1.1.2. *Further questions.* There are a huge number of other questions related to the results and conjectures in this paper. For example:

- **Q1** Could an even stronger version of Conjecture 1.7 hold, where in **B1** one demands that g is divisible by a polynomial of degree 1 (i.e., this degree-1 polynomial is "directly responsible" for the poor anticoncentration behaviour)?
- **Q2** Can one consider weaker notions of "smallness" of a polynomial than having few nonzero coefficients? In [34], the second and third authors show that in the d = 2 case of Theorem 1.1 one can replace property **A1** (that f has at most εn^2 nonzero coefficients) with the property that one can fix at most εn of the variables to obtain a constant polynomial. It would be very interesting if this notion of "smallness" could also be used for theorems along the lines of Theorem 1.9 to 1.11.
- **Q3** Can one quantify the " n^{ε} " error term in Theorem 1.9? We note that it is not possible to remove $n^{o(1)}$ terms entirely: for example, if we consider the polynomial $f = L_1L_2 L_3L_4$ in the d = 2 case of Proposition 1.4, known results on the "multiplication table problem" (counting the integers respresentable in the form $a \cdot b$, where $a, b \in \{1, \ldots, N\}$; see [17]) show that

$$\sup_{z \in \mathbb{F}} \mathbb{P}[f(\xi_1, \dots, \xi_n) = z] \ge \frac{(\log n)^{\alpha + o(1)}}{n}.$$

for some absolute constant $\alpha > 0$.

Q4 What can we prove in the "higher-dimensional" case, where we are interested in the probability that a length-k vector $(f_1(\xi_1, \ldots, \xi_n), \ldots, f_k(\xi_1, \ldots, \xi_n))$ of polynomials of independent random variables is equal to a particular vector $z \in \mathbb{F}^k$? In the linear (d = 1) case, essentially optimal

results were famously obtained by Halász [23] (see also the results in [16, 25], which consider the dependence on k), but almost nothing is known for higher degrees.

1.1.3. Inverse theorems. From a broader point of view, all the theorems and conjectures in this paper can be interpreted as falling under the umbrella of the *inverse Littlewood–Offord problem*, which is concerned with understanding the structural properties of f that control the anticoncentration behaviour of $f(\xi_1, \ldots, \xi_n)$. In this paper we have considered only the *algebraic* aspects of this problem; aspirationally it would be of interest to consider the interplay between algebraic and *arithmetic* aspects.

To give some context: in the linear (d = 1) case, there is no relevant algebraic structure; the anticoncentration behaviour of $f(\xi_1, \ldots, \xi_n)$ only depends on the multiset $\{a_1, \ldots, a_n\}$ of (degree-1) coefficients of f. It is now understood that, in fact, the anticoncentration behaviour of $f(\xi_1, \ldots, \xi_n) = a_1\xi_1 + \cdots + a_n\xi_n$ is essentially determined by "how much $\{a_1, \ldots, a_n\}$ looks like a generalised arithmetic progression". A particular theorem along these lines, due to Nguyen and Vu [45], appears as Theorem 2.10 later in the paper; earlier theorems in this spirit were proved by Halász [23], Tao and Vu [58,59] and Rudelson and Vershinyn [52].

For higher degrees, the only result in the literature considering both arithmetic and algebraic aspects of the Littlewood–Offord problem is a theorem of Nguyen [46] in the quadratic case (d = 2). This theorem had important consequences in random matrix theory, but unfortunately it is very crude, and we are still a long way from "optimal" inverse theorems for higher degrees. In particular, Nguyen's theorem is only capable of distinguishing "polynomial" anticoncentration from "sub-polynomial" anticoncentration (i.e., distinguishing whether there is a point probability at least n^{-C} , for some constant C, or not). For bilinear forms (a subfamily of quadratic polynomials), Costello made a more refined conjecture [9, Conjecture 1] but we do not know of a general conjecture for quadratic polynomials.

1.1.4. Property testing for (tensor) rank. As we will discuss further in the proof outlines in Section 2, the proofs of Theorems 1.9 to 1.11 depend on two "local-to-global" lemmas which allow one to study the rank of a matrix or tensor via the ranks of small submatrices or small subtensors. First, Lemma 2.2 says that if a d-dimensional $n \times \cdots \times n$ tensor T has the property that a 1-o(1) fraction of its $2^{d-1} \times \cdots \times 2^{d-1}$ subtensors have partition rank 1, then we can change a o(1)-fraction of the entries of T to obtain a tensor with partition rank 1. Second, Lemma 2.7 says that, if an $n \times n$ symmetric matrix A has the property that a 1-o(1) fraction of its $r \times r$ submatrices are singular, then we can change a o(1)-fraction of the entries of T to obtain a tensor with a 1-o(1) fraction of its $r \times r$ submatrices are singular, then we can change a o(1)-fraction of the entries of A to obtain a symmetric matrix with rank less than r.

Lemmas 2.2 and 2.7 are of independent interest, as they permit efficient property testing for matrices and tensors. For example, suppose we have a d-dimensional tensor T of enormous size $n \times \cdots \times n$, and we want to distinguish between the possibilities that T has partition rank 1 and that T is " ε -far" from having partition rank 1. With high probability, we can accomplish this in *constant* time (i.e., with a number of queries that only depends on ε , d, and the desired success probability), by randomly sampling many subtensors with side length 2^{d-1} and checking⁶ whether they have partition rank 1. There is an enormous literature on property testing of graphs, Boolean functions, probability distributions, etc. (see for example the monographs [3, 20, 50] and the references therein), and in particular there has been some specific attention on property testing for the rank of (not necessarily symmetric) matrices; see for example [2, 33, 37]. However, there does not seem to have been any prior work on property testing for any kind of tensor rank, or for matrices constrained to be symmetric. We believe there is quite some scope to investigate these directions further.

For example, while Lemma 2.7 already has near-optimal quantitative aspects, it would be interesting to investigate quantitative aspects for Lemma 2.2. It would also be interesting to investigate the possibility of more efficient property testing algorithms that do not simply randomly sample small submatrices/subtensors (for not-necessarily-symmetric matrices, there are rank-testing algorithms that do strictly better than sampling submatrices; see [37]).

It would also be very interesting to consider more general property testing for the partition rank (or other notions of tensor rank), beyond the setting of Section 4 (where we are only interested in whether the partition rank is 1 or not). Tensor ranks are surprisingly poorly behaved with respect to subtensors (see e.g. the surprising construction of Gowers in [28, Proposition 3.1]), and a straightforward generalisation of Lemma 2.2 to higher ranks does not actually seem to be possible.

⁶While computing the partition rank of a tensor is difficult in general, it is easy to check whether a tensor has partition rank 1 (this amounts to checking whether some matrix of a certain form has rank 1, see Fact 4.2).

1.2. **Organisation of the paper.** In Section 2 we provide outlines of the proofs of Theorems 1.9 to 1.11, including statements of various key lemmas. Then, in Section 3, we record some basic preliminary lemmas that will be used throughout the proofs.

The details of the proof of Theorem 1.9 appear in Sections 4 to 6. First, in Section 4 we prove a key lemma allowing us to restrict our attention to small subtensors, and in Section 5 we characterise the variety of reducible multilinear forms (i.e., reducible tensors). In Section 6 we combine all the relevant ingredients to prove Theorem 1.9.

Most of the rest of the paper is concerned with the proofs of Theorems 1.10 and 1.11. In Section 7 we prove a key lemma allowing us to restrict our attention to small submatrices, and in Section 8 we prove some special-purpose decoupling inequalities that will allow us to efficiently use tools from linear Littlewood–Offord theory to study quadratic anticoncentration. Then, the proof of Theorem 1.11 is completed in Section 9 and the proof of Theorem 1.10 is completed in Section 10.

Finally, we have two appendices: in Appendix A we show how to prove a certain variant of a "geometric Littlewood–Offord theorem" of Fox, Spink and the second author [18], which is an ingredient in the proof of Theorem 1.9. Then, Appendix B (authored by Matthew Kwan, Ashwin Sah and Mehtaab Sawhney) contains the proof of Proposition 1.4 (i.e., the disproof of Costello's conjecture).

1.3. Notation. We use standard asymptotic notation throughout, as follows. For functions f = f(n) and g = g(n), we write f = O(g) to mean that there is a constant C such that $|f(n)| \leq C|g(n)|$, and we write $f = \Omega(g)$ to mean that there is a constant c > 0 such that $f(n) \geq c|g(n)|$ for sufficiently large n. We also write f = o(g) to mean that $f(n)/g(n) \to 0$ as $n \to \infty$. Subscripts on asymptotic notation indicate quantities that should be treated as constants.

Slightly less standardly, for a matrix $A \in \mathbb{F}^{X \times Y}$ and subsets $I \subseteq X, J \subseteq Y$, we write $A[I, J] \in \mathbb{F}^{I \times J}$ for the $|I| \times |J|$ submatrix of A consisting of the rows with indices in I and the columns with indices in J. In the case that $I = \{i\}$ for some $i \in X$, we slightly abuse notation and write A[i, J] instead of $A[\{i\}, J]$. Similarly, for a vector $\vec{v} \in \mathbb{F}^X$ and a subset $I \subseteq X$ we write $\vec{v}[I] \in \mathbb{F}^I$ for the vector obtained from $\vec{v} \in \mathbb{F}^I$ by only taking the coordinates with indices in I. In addition, for a matrix $A \in \mathbb{F}^{X \times Y}$, we write $||A||_0$ for its number of nonzero entries.

We say that a partition $X = I \cup J$ of a set X into two subsets I and J is "non-trivial", if both of the subsets I and J are non-empty. For a real number x, the floor and ceiling functions are denoted $\lfloor x \rfloor = \max\{i \in \mathbb{Z} : i \leq x\}$ and $\lceil x \rceil = \min\{i \in \mathbb{Z} : i \geq x\}$. We will however sometimes omit floor and ceiling symbols and assume large numbers are integers, wherever divisibility considerations are not important. Finally, some conventions: all logarithms in this paper are in base e, unless specified otherwise, and the natural numbers \mathbb{N} include zero.

Acknowledgments. We would like to thank Tim Browning and Sarah Peluse for informing us about relevant references in the analytic number theory literature.

2. Proof outlines

In this section, we outline the proofs of Theorems 1.9 to 1.11. The proofs of Theorems 1.10 and 1.11 are quite closely related, but the proof strategy for Theorem 1.9 is rather different.

2.1. *d*-multilinear forms. First, we outline the proof of Theorem 1.9. Recall that a polynomial $f \in \mathbb{F}[x_1, \ldots, x_n]$ is a *d*-multilinear form if there is a partition $\{1, \ldots, n\} = I_1 \cup \cdots \cup I_d$ such that every term of f is of the form $c_{i_1,\ldots,i_d}x_{i_1}\ldots x_{i_d}$, where $c_{i_1,\ldots,i_d} \in \mathbb{F}$ and $i_j \in I_j$ for each $j \in \{1,\ldots,d\}$. Such a *d*-multilinear form can be equivalently viewed as a function $\mathbb{F}^{I_1} \times \cdots \times \mathbb{F}^{I_d} \to \mathbb{F}$ which is linear in each of its *d* arguments. It can be encoded by the $|I_1| \times \cdots \times |I_d|$ array of coefficients c_{i_1,\ldots,i_d} , and it can also be viewed as an element of the vector space $\mathbb{F}^{I_1} \otimes \cdots \otimes \mathbb{F}^{I_d}$. This gives rise to three different ways to define the notion of an $|I_1| \times \cdots \times |I_d|$ tensor:

- as a *d*-multilinear form, which can be interpreted as a polynomial in $\mathbb{F}[x_1, \ldots, x_n]$ or a function $\mathbb{F}^{I_1} \times \cdots \times \mathbb{F}^{I_d} \to \mathbb{F}$,
- as a function $I_1 \times \cdots \times I_d \to \mathbb{F}$ (i.e., a $I_1 \times \cdots \times I_d$ array with entries in \mathbb{F}), or
- as an element of the vector space $\mathbb{F}^{I_1} \otimes \cdots \otimes \mathbb{F}^{I_d}$.

We will switch between the these three points of view as convenient. Indeed, viewing a tensor as a function $I_1 \times \cdots \times I_d \to \mathbb{F}$ makes it easier to talk about subtensors (i.e., tensors induced by subsets $I'_1 \subseteq I_1, \ldots, I_d \subseteq I'_d$), while viewing a tensor as an element of $\mathbb{F}^{I_1} \otimes \cdots \otimes \mathbb{F}^{I_d}$ makes it easier to talk about factorisation.

Definition 2.1. A tensor $T \in \mathbb{F}^{I_1} \otimes \cdots \otimes \mathbb{F}^{I_d}$ is reducible or has partition rank at most 1 if there is a non-trivial partition $\{1, \ldots, d\} = J_1 \cup J_2$, and tensors $T_1 \in \bigotimes_{j \in J_1} \mathbb{F}^{I_j}$ and $T_2 \in \bigotimes_{j \in J_2} \mathbb{F}^{I_j}$, such that T can be factored as a tensor product $T = T_1 \otimes T_2$. Equivalently, this says that the corresponding d-multilinear form can be factorised into lower-degree factors.

The first ingredient we need is the following theorem, which states that for any tensor T, if most of the small subtensors of T are reducible, then T itself is close to being reducible. In our proof, we will use this as a sort of local-to-global principle: we can study small subtensors of T individually, and make conclusions about the whole tensor T.

Lemma 2.2. Fix d > 1 and $\varepsilon > 0$, and any field \mathbb{F} . Let $\delta = (\varepsilon/2)^{2^{d-1}}$, and let n_0 be sufficiently large in terms of d and ε . Consider a d-dimensional tensor $T : I_1 \times \cdots \times I_d \to \mathbb{F}$, where $|I_1|, \ldots, |I_d| \ge n_0$, and suppose that all but at most a δ -fraction of the $2^{d-1} \times \cdots \times 2^{d-1}$ subtensors of T are reducible. Then one can make T reducible by changing up to an ε -fraction of its entries.

Remark 2.3. The proof of Lemma 2.2 is also suitable for other notions of "tensor rank at most 1", most notably the so-called *slice rank* (see [54, 57]).

The (short) proof of Lemma 2.2 appears in Section 4. We will prove Theorem 1.9 by induction on d, using Costello's result for d = 2 as the base case. Given Lemma 2.2, the main ingredient for the induction step is the following lemma.

Lemma 2.4. Fix $d \ge 3$ and $r \ge 1$ and $\varepsilon > 0$ and $\mathbb{F} \in \{\mathbb{R}, \mathbb{C}\}$. Let n be sufficiently large in terms of d, r and ε . Consider a tensor $T : I_1 \times \cdots \times I_d \to \mathbb{F}$, where $|I_1| = \cdots = |I_{d-1}| = r$ and $I_d = \{1, \ldots, n\}$. For a vector $\vec{x} \in \mathbb{F}^n$, define the (d-1)-dimensional tensor $T\vec{x} : I_1 \times \cdots \times I_{d-1} \to \mathbb{F}$ by

$$T\vec{x}(i_1,\ldots,i_{d-1}) = \sum_{i=1}^n T(i_1,\ldots,i_{d-1},i)x_i$$

Then at least one of the following holds.

D1 All but an ε -fraction of the $r \times \cdots \times r$ subtensors of T are reducible, or

D2 letting $\vec{\xi} = (\xi_1, \dots, \xi_n) \in \{-1, 1\}^n$ be a vector of *i.i.d.* Rademacher random variables, we have

 $\mathbb{P}[T\vec{\xi} \text{ is reducible}] < n^{-1/2+\varepsilon}.$

Roughly speaking, we will use Lemma 2.4 as follows. Given a *d*-multilinear form f (with variables partitioned into d parts), we wish to understand the behaviour of f when the variables take random values. In the induction step, we reveal all the randomness in only the last of the d parts, which allows us to reinterpret f as a (d-1)-multilinear form in the remaining unrevealed variables. One way to visualise this, from the point of view that the tensor of coefficients is a d-dimensional array, is to view our d-dimensional tensor as a "stack" of (d-1)-dimensional tensors. We assign a random sign to each of the (d-1)-dimensional tensors in the stack, and add them together to obtain the "collapsed" (d-1)-dimensional tensor. For the induction to work, we need to know that if the original d-multilinear form was far from being reducible, then the "collapsed" (d-1)-multilinear form is likely to also be far from being reducible. By Lemma 2.2, it suffices to study this "collapse" on tensors whose side-lengths are 2^{d-1} (except in the d-th dimension), and this is precisely the role of Lemma 2.4.

In the statement of Lemma 2.4, note that $T\vec{\xi}$ can be viewed as a linear combination of Rademacher random variables, with coefficients in the space of $|I_{d-1}| \times \cdots \times |I_{d-1}|$ tensors $\mathbb{F}^{I_1} \otimes \cdots \otimes \mathbb{F}^{I_{d-1}}$. In **D2** we are interested in the event that this random linear combination falls in the variety of reducible tensors. So, to prove Lemma 2.4 we will use the following "geometric Littlewood–Offord" theorem (closely related to results of Fox, Spink and the second author [18, Theorem 1.5]), providing bounds on the probability that a Rademacher sum falls in a given variety.

Theorem 2.5. Fix $d \ge 1$ and $\varepsilon > 0$ and $\mathbb{F} \in \{\mathbb{R}, \mathbb{C}\}$ and let $\mathcal{Z} \subsetneq \mathbb{F}^d$ be a (possibly reducible) affine variety. Let n be sufficiently large (in terms of \mathcal{Z}, ε), and consider vectors $\vec{a}_1, \ldots, \vec{a}_n \in \mathbb{F}^d$. Then, at least one of the following holds.

E1 There is a linear subspace $\mathcal{W} \subseteq \mathbb{F}^d$ containing all but at most εn of the vectors \vec{a}_i , such that $\vec{w} + \mathcal{W} \subseteq \mathcal{Z}$ for some vector $\vec{w} \in \mathbb{F}^d$.

E2 letting $\xi_1, \ldots, \xi_n \in \{-1, 1\}$ be *i.i.d.* Rademacher random variables, we have

$$\mathbb{P}[\xi_1 \vec{a}_1 + \dots + \xi_n \vec{a}_n \in \mathcal{Z}] \le n^{-1/2 + \varepsilon}$$

This result follows from the same proof as in [18] (which considers a more general setting). For completeness, we include a proof of Theorem 2.5 in Appendix A, where we deduce Theorem 2.5 from the Meka–Nguyen–Vu bound for the polynomial Littlewood–Offord problem (see Theorem 1.1).

In order to actually apply Theorem 2.5, our final ingredient is a description of the maximal linear subspaces of the variety of reducible tensors.

Lemma 2.6. Fix disjoint sets I_1, \ldots, I_d , fix $\mathbb{F} \in \{\mathbb{R}, \mathbb{C}\}$ and let $r = \prod_{j=1}^d |I_j|$. Let $\mathcal{Z} \subseteq \mathbb{F}^{I_1} \otimes \cdots \otimes \mathbb{F}^{I_d}$ be the set of reducible $I_1 \times \cdots \times I_d$ tensors. Note that $\mathbb{F}^{I_1} \otimes \cdots \otimes \mathbb{F}^{I_d} \cong \mathbb{F}^r$, so we can view \mathcal{Z} as a subset of \mathbb{F}^r . Then:

- (1) $\mathcal{Z} \subseteq \mathbb{F}^r$ is a (possibly reducible) affine variety.
- (2) The maximal linear subspaces of Z are all sets of tensors of the following form: fix a proper subset Ø ⊆ J ⊆ {1,...,d}, fix a |J|-dimensional tensor T* ∈ Ø_{j∈J} 𝔽^{I_j}, and consider all tensors of the form T* ⊗ T', where T' ranges over all tensors T' ∈ Ø_{j∉J} 𝔽^{I_j}.

We prove Lemma 2.6 in Section 5.

In Section 6, we give the details of how to prove Lemma 2.4 using Lemma 2.6 and Theorem 2.5, and deduce Theorem 1.9 from Lemmas 2.2 and 2.4.

2.2. Quadratic polynomials. In this subsection we outline the proofs of Theorems 1.10 and 1.11. We start with some elements common to both proofs, and then split into separate subsections for each of Theorems 1.10 and 1.11.

First, recalling the role that Lemma 2.2 played for the proof of Theorem 1.9, we need a "local-toglobal" lemma for the rank of a symmetric matrix (while *d*-multilinear forms are naturally encoded by *d*-dimensional tensors, quadratic forms are naturally encoded by symmetric matrices). Recall from Section 1.3 that $||A||_0$ denotes the number of nonzero entries in a matrix A.

Lemma 2.7. Let $n \ge r \ge 1$, let $\alpha \in [0,1]$ and let \mathbb{F} be a field. Consider a symmetric matrix $A \in \mathbb{F}^{n \times n}$. If all but at most an α -fraction of the $r \times r$ submatrices of A are singular, then there exists a symmetric matrix $A' \in \mathbb{F}^{n \times n}$ of rank less than r such that $||A - A'||_0 \le O(r^4 \alpha^{1/r} n^2)$.

We remark that Lemma 2.7 is very easy to prove if we drop the requirement that the approximating matrix A' is symmetric. The challenge is that we want to modify A to have rank less than r while maintaining symmetry. The details of the proof of Lemma 2.7 appear in Section 7, but to give a brief idea: first, we show that we can adjust a small number of entries of A to obtain a matrix B which has low rank (but is not necessarily symmetric). Then, we identify a principal submatrix of B which is symmetric and "robustly has full rank", and use this as a "guide" to construct A'.

Remark 2.8. Throughout the paper we do not pay much attention to quantitative aspects, but it seems worth noting that in Lemma 2.7 the dependence on α is best-possible. Indeed, for any fixed integer $r \geq 1$, let $\ell = \lfloor \alpha^{1/2r} n/2 \rfloor$ and consider a binary matrix $A = (a_{ij}) \in \mathbb{R}^{n \times n}$ defined by

$$a_{ij} = \begin{cases} 1 & \text{if } j \ge n - \ell + i \text{ or } i \ge n - \ell + j \\ 0 & \text{otherwise} \end{cases}$$

(this can be informally described as an all-zero matrix, in which an upper-triangular chunk of the topright corner, and a lower-triangular chunk of the bottom-left corner, have been changed to "1"). It is not hard to see that all but at most an α -fraction of the $r \times r$ submatrices are singular, but to make the rank less than r it is necessary to change an $\Omega_r(\alpha^{1/r})$ -fraction of the entries.

Next, one very important aspect of the proofs of Theorems 1.10 and 1.11, which has no counterpart in the proof of Theorem 1.9, is *decoupling*. Namely, in the last subsection (on the proof of Theorem 1.9), we crucially used that *d*-multilinear forms can be inductively broken down into *linear* forms in disjoint sets of variables. For example, a bilinear form in the variables x_1, \ldots, x_m and y_1, \ldots, y_n can be viewed as a linear form in x_1, \ldots, x_m whose coefficients are themselves linear forms in y_1, \ldots, y_n . So, one can study anticoncentration of bilinear forms via tools for *linear* anticoncentration.

In Theorems 1.10 and 1.11, we are interested in general (not necessarily bilinear) quadratic polynomials, where this type of recursive linear structure is not available. However, one of the most important techniques for the polynomial Littlewood–Offord problem, called *decoupling*, nonetheless allows one to deduce information about general quadratic polynomials via tools for linear anticoncentration⁷.

 $^{^{7}}$ The term "decoupling" refers more generally to a class of techniques used to reduce from dependent situations to independent situations, see for example the book-length treatment in [12]. In the context of the polynomial Littlewood–Offord problem, this technique was introduced by Rosiński and Samorodnitsky [51] and Costello, Tao and Vu [10].

We will need some non-standard decoupling inequalities, which allow us to relate quadratic polynomials to *high-dimensional* linear anticoncentration problems. We will state and discuss our decoupling inequalities momentarily (when we start to break down the proofs of Theorems 1.10 and 1.11), but first we introduce a very powerful tool which gives us control over the high-dimensional linear anticoncentration problems that arise from our decoupling inequalities: namely, the so-called "optimal inverse theorem" of Nguyen and Vu [45] (see the discussion in Section 1.1.3 for context). This inverse theorem is stated in terms of generalised arithmetic progressions, defined as follows.

Definition 2.9. Let G be an additive group (i.e., an abelian group, additively written). A symmetric generalised arithmetic progression, or a symmetric GAP for short, is a map of the form

$$\varphi: \{-N_1, -N_1+1, \dots, N_1\} \times \dots \times \{-N_r, -N_r+1, \dots, N_r\} \to G$$

for some $r, N_1, \ldots, N_r \in \mathbb{N}$, such that there exist $v_1, \ldots, v_r \in G$ with $\varphi(a_1, \ldots, a_r) = a_1v_1 + \cdots + a_rv_r$ for all $(a_1, \ldots, a_r) \in \{-N_1, -N_1 + 1, \ldots, N_1\} \times \cdots \times \{-N_r, -N_r + 1, \ldots, N_r\}$ (in other words, such that φ is the restriction of a group homomorphism $\mathbb{Z}^r \to G$).

We say that r is the rank and $\prod_{i=1}^{r} (2N_i + 1)$ is the volume of the symmetric GAP. Furthermore, we say that a subset $G' \subseteq G$ is contained in (or "lies in") a symmetric GAP φ , if G' is contained in the image⁸ of φ .

Note that any symmetric GAP has volume at least 1. Now, the optimal inverse theorem of Nguyen and Vu [45, Theorem 2.5] is as follows.

Theorem 2.10. Fix C > 0. Let v_1, \ldots, v_n be elements of a torsion-free⁹ additive group G and ξ_1, \ldots, ξ_n be *i.i.d.* Rademacher random variables. If

$$\rho := \sup_{z \in G} \mathbb{P}[\xi_1 v_1 + \dots + \xi_n v_n = z] > n^{-C},$$

then for any $\sqrt{n} \leq n' \leq n$, there exists a symmetric GAP with rank $r = O_C(1)$ and volume $O_C(\rho^{-1}(n')^{-r/2})$, that contains all but at most n' of v_1, \ldots, v_n .

Note that the symmetric GAP in the conclusion must automatically have rank $r \leq 4C$ if n is sufficiently large in terms of C (since otherwise the volume $O_C(\rho^{-1}(n')^{-r/2}) \leq O_C(n^C\sqrt{n^{-r/2}})$ would be smaller than 1). In fact, if n' is chosen to be linear in n one obtains an even better bound for the rank (namely, $r \leq 2C$ if $n' \geq \delta n$ for some fixed $\delta > 0$ and n is sufficiently large with respect to C and δ).

If we take G to be a high-dimensional space of the form \mathbb{F}^d (for $\mathbb{F} \in \{\mathbb{R}, \mathbb{C}\}$), then Theorem 2.10 gives us anticoncentration bounds for random variables of the form $A\vec{\xi}$ (where $A \in \mathbb{F}^{d \times n}$ is a matrix, and $\vec{\xi} \in \{-1,1\}^n$ is a vector of i.i.d. Rademacher random variables). The quality of the bound depends on how well the columns of A can be "covered" by a GAP of low rank and small volume. In particular, one can already obtain useful bounds just using information about the "robust rank" of A (without taking into account the volume of the covering GAP); such a bound was previously obtained by Halász [23], and will be used occasionally in this paper (it appears explicitly as Theorem 3.2).

Now we move into the details of the proofs of Theorems 1.10 and 1.11. Both proofs follow a similar proof strategy, but the proof of Theorem 1.11 is simpler, so we start with that.

2.2.1. Bounds in terms of the rank. We now outline the proof of Theorem 1.11, approaching the "generic bound" 1/n for quadratic forms of high rank.

First, our decoupling inequality is as follows. (Recall from Section 1.3 our notation $\bar{\xi}[X]$, A[X, Y] for subvectors and submatrices).

Definition 2.11. We say that a random variable ξ is *shifted Rademacher* if $\xi - \mathbb{E}\xi$ has the Rademacher distribution.

Lemma 2.12. *Fix* $\mathbb{F} \in \{\mathbb{R}, \mathbb{C}\}$ *and a positive integer* k*.*

- Let $f \in \mathbb{F}[x_1, \ldots, x_n]$ be a quadratic polynomial, so there is a unique way to write the quadratic part of $f(\vec{x})$ as $\vec{x}^T A \vec{x}$ for some symmetric matrix $A \in \mathbb{F}^{n \times n}$.
- Let $\vec{\xi} = (\xi_1, \dots, \xi_n)$ be a sequence of independent Rademacher random variables.
- Consider a partition $\{1, \ldots, n\} = X \cup Y$.

⁸It is more common to define GAPs as sets, not maps (i.e., many authors would say that the image of φ is itself a GAP). However, this creates some ambiguity around rank and volume (since the image of φ may not uniquely determine φ), and later in this paper it will be rather important to be precise about these notions.

⁹Recall that an abelian group is *torsion-free* if there is no nonzero element of finite order.

Then,

$$\sup_{z \in \mathbb{F}} \mathbb{P}[f(\vec{\xi}) = z] \le \mathbb{P}\Big[(\vec{\alpha}^{(i)})^T A[X, Y] \vec{\xi}[Y] = \psi(\vec{\alpha}^{(i)}) \text{ for all } i \in \{1, \dots, k\} \Big]^{1/(k+1)},$$
(2.1)

for some function $\psi : \mathbb{F}^X \to \mathbb{F}$, and some *i.i.d.* random vectors $\vec{\alpha}^{(1)}, \ldots, \vec{\alpha}^{(k)} \in \mathbb{F}^X$ with independent shifted Rademacher entries (all independent of $\vec{\xi}[Y]$).

We give the short proof of Lemma 2.12 in Section 8.

Note that, after conditioning on an outcome of $\vec{\alpha}^{(i)}$, the expression $(\vec{\alpha}^{(i)})^T A[X,Y] \vec{\xi}[Y] - \psi(\vec{\alpha}^{(i)})$ becomes a linear function of $\vec{\xi}[Y]$. So, the k = 1 case of Lemma 2.12 provides a direct way to deduce quadratic anticoncentration bounds from linear anticoncentration bounds. This particular case of Lemma 2.12 is well-known (in this form, it seems to have been first observed by Costello and Vu [11], though similar inequalities appeared earlier in [10, 51, 56]).

Although it is near-trivial to generalise the proof of the k = 1 case of Lemma 2.12 to general k, the possibility of "decoupling with multiple copies" (and the advantage of doing so) has been recognised only recently. In particular, a related (though much more delicate) decoupling scheme was recently used by the second and third authors [34] to obtain optimal bounds for the quadratic case of Theorem 1.1. We believe Lemma 2.12 makes it quite transparent why it is helpful to decouple with more than one copy. Indeed, (2.1) involves k simultaneous equations; if we imagine that each equation is satisfied independently with probability at most about 1/n, then Lemma 2.12 provides a bound of about $n^{-k/(k+1)}$, which gets closer and closer to 1/n as we increase k.

Of course, the k simultaneous equations in (2.1) are not independent. The primary challenge in our proof of Theorem 1.11 is to demonstrate that, for the purposes of anticoncentration, these equations are "approximately independent", if **C1** does not hold (i.e., if our quadratic form robustly has high rank). Specifically, given Lemmas 2.7 and 2.12, the following lemma is the main part of the proof of Theorem 1.11. (Recall the definition of a shifted Rademacher random variable from Definition 2.11).

Lemma 2.13. Fix any $\varepsilon, \delta \in (0,1]$ and $k \ge 1$, and let $\mathbb{F} \in \{\mathbb{R}, \mathbb{C}\}$. Let n be sufficiently large (in terms of ε, δ, k), and let $A \in \mathbb{F}^{n \times n}$ be a matrix. Then, at least one of the following holds.

G1 At most a δ -fraction of the $\lceil 2k^2/\varepsilon \rceil \times \lceil 2k^2/\varepsilon \rceil$ submatrices of A are nonsingular, or

G2 letting $\Xi \in \mathbb{F}^{k \times n}$ be a random matrix with independent shifted Rademacher entries, and (independently) letting $\vec{\eta} \in \{-1, 1\}^n$ be a random column vector with independent Rademacher entries, for any (non-random) function $\vec{\varphi} : \mathbb{F}^{k \times n} \to \mathbb{F}^k$ we have

$$\mathbb{P}\big[\Xi A\vec{\eta} = \vec{\varphi}(\Xi)\big] \le n^{-k+\varepsilon}.$$

To prove Lemma 2.13, we first use a counting argument, and some anticoncentration estimates using the rank of A, to show that it is unlikely that most of the columns of ΞA are contained in a generalised arithmetic progression with small rank and volume. Then, we can apply Theorem 2.10 (the Nguyen–Vu optimal linear inverse theorem) with the randomness of $\vec{\eta}$, to prove an upper bound on the probability that $\Xi A \vec{\eta} = \vec{\varphi}(\Xi)$.

Remark 2.14. We suspect that the rank bound in **G1** is far from best possible; indeed, we suspect that the same statement should hold with " $\lceil 2k^2/\varepsilon \rceil$ " replaced by "2k" (if true, this would be best-possible, as one can see by taking A to be an appropriate block-diagonal matrix). Actually, this would constitute a generalisation of the d = 2 case of Theorem 1.9, via Lemma 2.7. We did not fight too hard to optimise the particular rank bound in **G1**, but our methods do seem to be somewhat too crude to prove an optimal bound.

In Section 9, we give the details of the proof of Lemma 2.13, and deduce Theorem 1.11 from Lemmas 2.7, 2.12, and 2.13.

2.2.2. A power-saving improvement in the "robustly irreducible" case. Now, we outline the proof of Theorem 1.10, giving a power-saving improvement for quadratic polynomials which are "robustly irreducible" over \mathbb{C} .

First, the significance of \mathbb{C} is that a quadratic polynomial is reducible over \mathbb{C} if and only if it has rank at most 2. (On the other hand, reducibility over \mathbb{R} cannot be expressed in terms of rank: $x^2 + y^2$ and $x^2 - y^2$ both have rank two, but the former is irreducible and the latter is reducible).

Given this connection between irreducibility and rank, our proof of Theorem 1.10 follows the same general scheme as the proof of Theorem 1.11, though the details are much more delicate. Indeed, for

Theorem 1.11 we could wastefully assume our coefficient matrix A had high rank, but here we can only assume that the rank is at least 3.

Just as for Theorem 1.11, the first ingredient is a decoupling lemma. Instead of Lemma 2.12, we use the following more technical lemma.

Definition 2.15. we say that a random variable α is *lazy Rademacher* if it can be expressed as the difference of two independent Rademacher random variables (explicitly, this means $\mathbb{P}[\alpha = 0] = 1/2$ and $\mathbb{P}[\alpha = -1] = \mathbb{P}[\alpha = 1] = 1/4$).

Lemma 2.16. *Fix* $\mathbb{F} \in \{\mathbb{R}, \mathbb{C}\}$.

- Let $f \in \mathbb{F}[x_1, \ldots, x_n]$ be a quadratic polynomial, so there is a unique way to write the quadratic part of $f(\vec{x})$ as $\vec{x}^T A \vec{x}$ for some symmetric matrix $A \in \mathbb{F}^{n \times n}$.
- Let $\vec{\xi} = (\xi_1, \dots, \xi_n)$ be a sequence of independent Rademacher random variables.
- Consider a partition $\{1, \ldots, n\} = X \cup Y \cup Z$.

Then,

$$\sup_{z\in\mathbb{F}}\mathbb{P}[f(\vec{\xi})=z] \le \mathbb{P}\Big[\vec{\alpha}^{T}A[X,Y]\vec{\beta}=0, \quad \vec{\alpha}^{T}A[X,Z]\vec{\gamma}=\varphi(\vec{\alpha},\vec{\beta}), \quad \vec{\beta}^{T}A[Y,Z]\vec{\gamma}=\psi(\vec{\alpha},\vec{\beta})\Big]^{1/4},$$

for some functions $\varphi, \psi : \mathbb{F}^X \times \mathbb{F}^Y \to \mathbb{F}$, and independent random vectors $\vec{\alpha}, \vec{\beta}, \vec{\gamma}$, where the entries of $\vec{\alpha}, \vec{\beta}$ are *i.i.d.* lazy Rademacher, and the entries of $\vec{\gamma}$ are *i.i.d.* Rademacher.

We prove Lemma 2.16 in Section 8. Then, given Lemmas 2.7 and 2.16, the proof of Theorem 1.10 essentially comes down to the following lemma.

Lemma 2.17. Fix any $\varepsilon \in (0,1]$, and let $\mathbb{F} \in \{\mathbb{R}, \mathbb{C}\}$. Let n be sufficiently large (in terms of ε), and consider matrices $A_1, A_2, A_3 \in \mathbb{F}^{n \times n}$. Then, at least one of the following holds.

- **H1** For one of the matrices A_1, A_2, A_3 , at most an ε -fraction of its 3×3 submatrices are nonsingular, or
- **H2** letting $\vec{\alpha}, \vec{\beta} \in \{-1, 0, 1\}^n$ be vectors of lazy Rademacher random variables, and $\vec{\gamma} \in \{-1, 1\}^n$ be a vector of Rademacher random variables (all independent), for any functions $\varphi, \psi : \mathbb{F}^n \times \mathbb{F}^n \to \mathbb{F}$ we have
 - $\mathbb{P}\Big[\vec{\alpha}^{T}A_{1}\vec{\beta}=0, \quad \vec{\alpha}^{T}A_{2}\vec{\gamma}=\varphi(\vec{\alpha},\vec{\beta}), \quad \vec{\beta}^{T}A_{3}\vec{\gamma}=\psi(\vec{\alpha},\vec{\beta})\Big] \leq n^{-2-1/6+\varepsilon}.$

At a high level, the proof of Lemma 2.17 can be compared to the proof of Lemma 2.13. To study anticoncentration of a bilinear form of random variables, we reveal its two sets of variables in two stages; in the first stage we prepare for the application of Theorem 2.10 (the Nguyen–Vu optimal linear inverse theorem), and in the second stage we actually apply Theorem 2.10. However, the details of Lemma 2.17 are much more complicated than the details of Lemma 2.13, for two main reasons.

First, since Lemma 2.17 concerns three different bilinear forms which depend on each other in a "cyclic" manner, we must be very careful about the way that we reveal the random variables $\vec{\alpha}, \vec{\beta}, \vec{\gamma}$. We actually need to consider three different cases, in which we reveal the random variables $\vec{\alpha}, \vec{\beta}, \vec{\gamma}$ in different ways. (The three cases are defined in terms of the random vectors $\vec{\alpha}^T A_1$ and $\vec{\alpha}^T A_2$; specifically, we need to distinguish whether one of these vectors is almost zero, and whether these vectors are almost collinear).

Second, we have much less room to make crude estimates, and we need to be much more careful about the tradeoffs between the stages of random exposure. In particular, in addition to the simple counting-based estimates featuring in proof of Lemma 2.13, we also need to use some more technical estimates proved by Costello [9] via number-theoretic means.

In Section 10, we give the details of the proof of Lemma 2.17, and deduce Theorem 1.10 from Lemmas 2.7, 2.16, and 2.17.

3. Preliminaries

In this section we collect a few basic tools that will be used throughout the paper. First, the following lemma is related to the "local-to-global" point of view in Lemmas 2.2 and 2.7. (Very similar estimates appear for example in [9, Section 3.4], but we were not able to find an easily citable reference).

Lemma 3.1. Let $p, \delta \in (0, 1)$ and $n \ge r \ge 1$. Suppose I is a random subset of $\{1, \ldots, n\}$ such that $\mathbb{P}[|I| \ge (1-\delta)n] \ge p$. Then, for all but a $2r\delta$ -fraction of the r-element subsets $S \subset \{1, \ldots, n\}$, we have $\mathbb{P}[S \subset I] \ge p/2$.

Proof. For $i \in \{1, \ldots, n\}$, let \mathcal{E}_i be the event that $i \in I$ and let \mathcal{E} be the event that $|I| \ge (1-\delta)n$, so we are assuming $\mathbb{P}[\mathcal{E}] \ge p$. Let \mathcal{Q} be the collection of sets $\{i_1, \ldots, i_r\} \subseteq \{1, \ldots, n\}$ of r distinct indices such that $\mathbb{P}[\mathcal{E}_{i_1} \cap \cdots \cap \mathcal{E}_{i_r}] < p/2 \le \mathbb{P}[\mathcal{E}]/2$. So, our goal is to prove that $|\mathcal{Q}| \le 2r\delta\binom{n}{r}$.

On one hand, we have

$$\mathbb{E}\left[\mathbb{1}_{\mathcal{E}}\binom{|I|}{r}\right] \ge \binom{\lceil (1-\delta)n\rceil}{r} \mathbb{P}[\mathcal{E}],\tag{3.1}$$

but on the other hand we have

$$\mathbb{E}\left[\mathbb{1}_{\mathcal{E}}\binom{|I|}{r}\right] = \sum_{i_1 < \dots < i_r} \mathbb{P}[\mathcal{E} \cap \mathcal{E}_{i_1} \cap \dots \cap \mathcal{E}_{i_r}] \le \sum_{\{i_1,\dots,i_r\} \in \mathcal{Q}} \mathbb{P}[\mathcal{E}_{i_1} \cap \dots \cap \mathcal{E}_{i_r}] + \sum_{\{i_1,\dots,i_r\} \notin \mathcal{Q}} \mathbb{P}[\mathcal{E}] < |\mathcal{Q}| \cdot (\mathbb{P}[\mathcal{E}]/2) + \left(\binom{n}{r} - |\mathcal{Q}|\right) \cdot \mathbb{P}[\mathcal{E}].$$
(3.2)

Combining (3.1) and (3.2) yields the desired inequality

$$|\mathcal{Q}| < 2\left(\binom{n}{r} - \binom{\lceil (1-\delta)n\rceil}{r}\right) \le 2\left(\delta n \cdot \binom{n-1}{r-1}\right) = 2r\delta\binom{n}{r}.$$

(The second inequality has a combinatorial interpretation: given a set of n elements, and an identification of $\lfloor \delta n \rfloor$ of those elements as "special", we are counting the number of ways to choose a subset of r elements, at least one of which is special).

Next, the following high-dimensional Littlewood–Offord theorem is due to Halász [23] (it can also be deduced from Theorem 2.10).

Theorem 3.2. Let $n \ge d \ge 1$ and $\ell \ge 1$ and $\mathbb{F} \in \{\mathbb{R}, \mathbb{C}\}$, and suppose $A \in \mathbb{F}^{d \times n}$ is a matrix containing at least ℓ disjoint nonsingular $d \times d$ submatrices. Let $\vec{\xi} = (\xi_1, \ldots, \xi_n) \in \mathbb{F}^n$ be a vector of independent Rademacher (or shifted Rademacher, or lazy Rademacher¹⁰) random variables. Then,

$$\sup_{\vec{z}\in\mathbb{F}^d}\mathbb{P}\Big[A\vec{\xi}=\vec{z}\Big]\leq O_d(\ell^{-d/2})$$

Halász' theorem has the following consequence for the probability that a random variable falls in an affine-linear subspace (see for example [34, Corollary 3.4]).

Corollary 3.3. Let $n \ge d \ge 1$ and $\ell \ge 1$ and $\mathbb{F} \in \{\mathbb{R}, \mathbb{C}\}$, and suppose a matrix $A \in \mathbb{F}^{d \times n}$ contains at least ℓ disjoint nonsingular $d \times d$ submatrices. Let $\vec{\xi} = (\xi_1, \ldots, \xi_n) \in \mathbb{F}^n$ be a vector of independent Rademacher (or shifted Rademacher, or lazy Rademacher) random variables. Then, for any k-dimensional affine-linear subspace $\mathcal{W} \subseteq \mathbb{F}^d$, we have

$$\mathbb{P}\Big[A\vec{\xi} \in \mathcal{W}\Big] \le O_d(\ell^{-(d-k)/2}).$$

Finally, for Theorems 1.10 and 1.11 we will need some basic observations about matrices. The following fact allows us to translate between the property of having many nonsingular submatrices and the property of having many *disjoint* nonsingular submatrices.

Fact 3.4. Let $n \ge d \ge 1$ and $m \ge 1$ and $\mathbb{F} \in \{\mathbb{R}, \mathbb{C}\}$. If a matrix $A \in \mathbb{F}^{d \times n}$ does not contain more than m disjoint nonsingular $d \times d$ submatrices, then at most an (md^2/n) -fraction of the $d \times d$ submatrices in A are nonsingular.

Proof. Fix a maximal collection of disjoint nonsingular $d \times d$ submatrices in A; these involve at most md different columns. The maximality of the collection implies that any $d \times d$ submatrix that does not use any of these md columns must be singular. Hence, the number of nonsingular $d \times d$ submatrices in A is at most $md\binom{n-1}{d-1} = (md^2/n)\binom{n}{d}$.

The following "monotonicity" fact says that if a matrix has many nonsingular $r \times r$ submatrices, then for any $k \leq r$, there are also many nonsingular $k \times k$ submatrices.

Fact 3.5. Let $\mathbb{F} \in \{\mathbb{R}, \mathbb{C}\}$. Let $A \in \mathbb{F}^{n \times n}$ be a matrix with the property that at least an ε -fraction of its $r \times r$ submatrices are nonsingular. Then, for any $k \leq r$, at least an $(\varepsilon/\binom{r}{k})$ -fraction of the $k \times k$ submatrices of A are nonsingular.

¹⁰It is easy to deduce the lazy and shifted Rademacher cases from the Rademacher case. Indeed, recall that a lazy Rademacher random variable is the difference of two independent Rademacher random variables; after conditioning on one of these it is a shifted Rademacher random variable. By linearity, shifting the entries of $\vec{\xi}$ does not affect anticoncentration.

Proof. Suppose that A has at least $\varepsilon {\binom{n}{r}}^2$ nonsingular $r \times r$ submatrices. For $k \leq r$, each of these nonsingular $r \times r$ submatrices contains at least ${\binom{r}{k}}$ nonsingular $k \times k$ submatrices, and every $k \times k$ submatrix can be extended to exactly ${\binom{n-k}{r-k}}^2$ different $r \times r$ submatrices. So, there are at least $\varepsilon {\binom{n}{r}}^2 {\binom{r}{k}} / {\binom{n-k}{r-k}}^2 = (\varepsilon / {\binom{r}{k}}) {\binom{n}{k}}^2$ nonsingular $k \times k$ submatrices.

The following fact says that if a matrix has many nonsingular $r \times r$ submatrices, then it cannot be close to a matrix of rank less than r. (It can be viewed as the easy "converse direction" of Lemma 2.7).

Fact 3.6. Let $\mathbb{F} \in \{\mathbb{R}, \mathbb{C}\}$. Fix $\varepsilon > 0$ and $r \ge 1$, and let n be sufficiently large (in terms of ε, r). Let $A \in \mathbb{F}^{n \times n}$ be a matrix with the property that at least an $r^2 \varepsilon$ -fraction of its $r \times r$ submatrices are nonsingular. Then there is no matrix $B \in \mathbb{F}^{n \times n}$ with rank less than r such that $||A - B||_0 \le \varepsilon n^2$.

Proof. If $||A - B||_0 \leq \varepsilon n^2$ then there are at most $\varepsilon n^2 {\binom{n-1}{r-1}}^2 = r^2 \varepsilon {\binom{n}{r}}^2$ different $r \times r$ submatrices of A that are not identical to their counterpart in B. So, if B has rank less than r, then at most an $r^2 \varepsilon$ -fraction of the $r \times r$ submatrices in A are nonsingular.

The following lemma says that if a matrix has many nonsingular $r \times r$ submatrices, then we can partition it into q block matrices, each of which individually has many $r \times r$ submatrices.

Lemma 3.7. Let $\mathbb{F} \in \{\mathbb{R}, \mathbb{C}\}$. Fix $\varepsilon > 0$, and $q, r \ge 1$. Let n be sufficiently large (in terms of ε, q, r), and divisible by q. Let $A \in \mathbb{F}^{n \times n}$ be a matrix with the property that at least an ε -fraction of its $r \times r$ submatrices are nonsingular.

Then, there is a partition of $\{1, \ldots, n\} = I_1 \cup \cdots \cup I_q$ with $|I_1| = \cdots = |I_q| = n/q$, such that for any $i, j \in \{1, \ldots, q\}$, more than an $\varepsilon/2$ -fraction of the $r \times r$ submatrices of the matrix $A[I_i, I_j]$ are nonsingular.

Proof. This is a routine application of the probabilistic method. Consider a random partition of $\{1, \ldots, n\}$ into sets I_1, \ldots, I_q of size n/q. Let $X_{i,j}$ be the number of nonsingular $r \times r$ submatrices of $A[I_i, I_j]$.

There are at least $\varepsilon {\binom{n}{r}}^2 - n {\binom{n-1}{r-1}}^2 = (\varepsilon - o(1)) {\binom{n}{r}}^2$ nonsingular $r \times r$ submatrices which do not involve a diagonal entry of A (here and in the rest of this proof, the asymptotic notation is for fixed ε, q, r). For any $i, j \in \{1, \ldots, q\}$, each of these submatrices is contained in $A[I_i, I_j]$ with probability $(1 + o(1))/q^{2r}$. So,

$$\mathbb{E}X_{i,j} \ge \frac{(\varepsilon - o(1))\binom{n}{r}^2}{q^{2r}} = (\varepsilon - o(1))\binom{n/q}{r}^2.$$

Now, moving an element in or out of I_i or I_j can only change $X_{i,j}$ by at most $O(n^{2r-1})$. So, by the Azuma–Hoeffding inequality (see for example [19, Lemma 11] for a statement adapted to our purposes¹¹), for each i, j we have

$$\mathbb{P}\left[X_{i,j} \le (\varepsilon/2) \binom{n/q}{r}^2\right] \le \exp\left(-\Omega\left(\frac{\left(\mathbb{E}X_{i,j}\right)^2}{n \cdot O(n^{2r-1})^2}\right)\right) = \exp(-\Omega(n)) < \frac{1}{q^2}.$$

for sufficiently large *n*. It follows that with positive probability we have $X_{i,j} > (\varepsilon/2) {\binom{n/q}{r}}^2$ for all q^2 choices of *i*, *j*, as desired.

4. A LOCAL-TO-GLOBAL LEMMA FOR REDUCIBILITY OF TENSORS

In this section, we prove Lemma 2.2: if most of the small subtensors of T are reducible, then T itself is close to being reducible. There are different types of reducibility with respect to particular partitions, as follows.

Definition 4.1. For a tensor $T \in \mathbb{F}^{I_1} \otimes \cdots \otimes \mathbb{F}^{I_d}$ and a non-trivial partition $\{1, \ldots, d\} = J_1 \cup J_2$, we say that T is *reducible with respect to* $\{J_1, J_2\}$ if there are tensors $T_1 \in \bigotimes_{j \in J_1} \mathbb{F}^{I_j}$ and $T_2 \in \bigotimes_{j \in J_2} \mathbb{F}^{I_j}$, such that T can be factored as a tensor product $T = T_1 \otimes T_2$.

We note that reducibility of a tensor T with respect to $\{J_1, J_2\}$ can equivalently be interpreted as the property that a certain matrix associated with T has rank at most 1, as follows.

¹¹This statement is for random permutations. But note that a random partition into q parts can easily be defined in terms of a random permutation: simply take the first n/q elements as the first part, the next n/q elements as the second part, and so on.

Fact 4.2. Fix disjoint sets I_1, \ldots, I_d , a field $\mathbb{F} \in \{\mathbb{R}, \mathbb{C}\}$, and a non-trivial partition $\{1, \ldots, d\} = J_1 \cup J_2$. For $s \in \{1, 2\}$, let $r_s = \prod_{j \in J_s} |I_j|$, so we have a vector space isomorphism $\bigotimes_{j \in J_s} \mathbb{F}^{I_j} \cong \mathbb{F}^{r_s}$. Then,

$$\mathbb{F}^{I_1} \otimes \cdots \otimes \mathbb{F}^{I_d} = \left(\bigotimes_{j \in J_1} \mathbb{F}^{I_j}\right) \otimes \left(\bigotimes_{j \in J_2} \mathbb{F}^{I_j}\right) \cong \mathbb{F}^{r_1 \times r_2};$$

that is to say, $\{J_1, J_2\}$ gives rise to a correspondence between $I_1 \times \cdots \times I_d$ tensors and $r_1 \times r_2$ matrices. A tensor in $\mathbb{F}^{I_1} \otimes \cdots \otimes \mathbb{F}^{I_d}$ is reducible with respect to $\{J_1, J_2\}$ if and only if the corresponding $r_1 \times r_2$ matrix has rank at most 1.

We also adopt the following notation for subtensors, similar to that of matrices.

Definition 4.3. Consider a tensor $T : I_1 \times \cdots \times I_d \to \mathbb{F}$ (represented as a *d*-dimensional array). For subsets $S_1 \subseteq I_1, \ldots, S_d \subseteq I_d$, write $T[S_1, \ldots, S_d]$ for the corresponding subtensor (formally, this is a restriction of the function T to the set $S_1 \times \cdots \times S_d$).

We break the proof of Lemma 2.2 into three simpler lemmas. First, by a simple averaging argument, we can reduce our consideration to the subtensors containing a particular nonzero entry.

Lemma 4.4. Fix d, r > 1 and $\delta > 0$ and $c \in [0, 1]$, and any field \mathbb{F} . Consider a d-dimensional tensor $T: I_1 \times \cdots \times I_d \to \mathbb{F}$ (represented as a d-dimensional array). Suppose that at least a δ^c -fraction of the entries of T are nonzero and that all but at most a δ -fraction of the $r \times \cdots \times r$ subtensors are reducible.

Then, there is a nonzero entry $T(i_1^*, i_2^*, \ldots, i_d^*)$ such that all but at most a δ^{1-c} -fraction of the $r \times \cdots \times r$ subtensors which contain the entry $T(i_1^*, i_2^*, \ldots, i_d^*)$ are reducible.

Proof. For each $j \in \{1, \ldots, d\}$, let $i_j^{\star}, i_j^1, \ldots, i_j^{r-1} \in I_j$ be a uniformly random sequence of r distinct indices from I_j . So, by assumption we have

$$\begin{split} \delta &\geq \mathbb{P}\Big[T\big[\{i_1^{\star}, i_1^1, \dots, i_1^{r-1}\}, \dots, \{i_d^{\star}, i_d^1, \dots, i_d^{r-1}\}\big] \text{ is not reducible}\Big] \\ &= \mathbb{E}\Big[\mathbb{P}\Big[T\big[\{i_1^{\star}, i_1^1, \dots, i_1^{r-1}\}, \dots, \{i_d^{\star}, i_d^1, \dots, i_d^{r-1}\}\big] \text{ is not reducible } \Big| i_1^{\star}, \dots, i_d^{\star}\Big]\Big]. \end{split}$$

By Markov's inequality, we deduce that for more than a $(1 - \delta^c)$ -fraction of choices of (i_1^*, \ldots, i_d^*) , we have

$$\mathbb{P}\Big[T\big[\{i_1^{\star}, i_1^{1}, \dots, i_1^{r-1}\}, \dots, \{i_d^{\star}, i_d^{1}, \dots, i_d^{r-1}\}\big] \text{ is not reducible } \Big| i_1^{\star}, \dots, i_d^{\star}\Big] \le \delta^{1-c},$$

meaning that all but at most a δ^{1-c} -fraction of the $r \times \cdots \times r$ subtensors which contain the entry $T(i_1^*, \ldots, i_d^*)$ are reducible. Since at least a δ^c -fraction of the entries of T are nonzero, for at least one of these choices of (i_1^*, \ldots, i_d^*) , the entry $T(i_1^*, \ldots, i_d^*)$ is nonzero.

The second lemma (arguably the most important) allows us to relate reducibility in general to reducibility with respect to a particular partition.

Lemma 4.5. Fix d > 1 and $\delta > 0$, and let $\ell = 2^{d-1} - 1$. Consider a tensor $T : I_1 \times \cdots \times I_d \to \mathbb{F}$ (represented as a d-dimensional array) and suppose $T(i_1^*, \ldots, i_d^*) \neq 0$ for some indices $i_1^* \in I_1, \ldots, i_d^* \in I_d$. Suppose that all but at most a δ -fraction of the $2^{d-1} \times \cdots \times 2^{d-1}$ subtensors of T which contain the entry $T(i_1^*, \ldots, i_d^*)$ are reducible.

Then, there is a non-trivial partition $\{1, \ldots, d\} = J_1 \cup J_2$, such that all but at most a $\delta^{1/\ell}$ -fraction of the $2 \times \cdots \times 2$ subtensors of T which contain the entry $T(i_1^*, \ldots, i_d^*)$ are reducible with respect to $\{J_1, J_2\}$.

Proof. Note that $\ell = 2^{d-1} - 1$ is the number of partitions $\mathcal{P} = \{J_1, J_2\}$ of $\{1, \ldots, d\}$, such that neither J_1 nor J_2 is empty. Let $\mathcal{P}_1, \ldots, \mathcal{P}_\ell$ be an enumeration of all these partitions.

For each $j \in \{1, \ldots, d\}$, let i_j^1, \ldots, i_j^ℓ be uniformly random indices sampled independently from $I_j \setminus \{i_j^\star\}$. Note that there may be fewer than ℓ different indices among i_j^1, \ldots, i_j^ℓ , due to repetitions; let Q_j be a random set of exactly ℓ indices in $I_j \setminus \{i_j^\star\}$, obtained by starting from $\{i_j^1, \ldots, i_j^\ell\}$ and adding the appropriate number of additional random indices. Do this independently for each $j \in \{1, \ldots, d\}$.

Then, we have

$$\delta \geq \mathbb{P}\Big[T\Big[Q_1 \cup \{i_1^{\star}\}, \dots, Q_d \cup \{i_d^{\star}\}\Big] \text{ is not reducible}\Big]$$
$$= \mathbb{P}\Bigg[\bigcap_{t=1}^{\ell} \Big\{T\Big[Q_1 \cup \{i_1^{\star}\}, \dots, Q_d \cup \{i_d^{\star}\}\Big] \text{ is not reducible with respect to } \mathcal{P}_t\Big\}\Bigg]$$

$$\geq \mathbb{P}\left[\bigcap_{t=1}^{\ell} \left\{ T\left[\{i_{1}^{t}, i_{1}^{\star}\}, \dots, \{i_{d}^{t}, i_{d}^{\star}\}\right] \text{ is not reducible with respect to } \mathcal{P}_{t} \right\} \right]$$
$$= \prod_{t=1}^{\ell} \mathbb{P}\left[T\left[\{i_{1}^{t}, i_{1}^{\star}\}, \dots, \{i_{d}^{t}, i_{d}^{\star}\}\right] \text{ is not reducible with respect to } \mathcal{P}_{t} \right].$$

Indeed, the first line is by the assumption on T and the fact that each Q_j is an independent and uniformly random subset of $I_j \setminus \{i_j^*\}$ of size ℓ , the second line is by the definition of reducibility, the third line is due to the fact that $Q_j \cup \{i_j^*\} \supseteq \{i_j^t, i_j^*\}$ for all $j \in \{1, \ldots, d\}$, and the last line is by the independence of (i_1^t, \ldots, i_d^t) between different t.

We deduce that there is some $t \in \{1, \ldots, \ell\}$ such that, with probability at most $\delta^{1/\ell}$, the random subtensor $T[\{i_1^t, i_1^\star\}, \ldots, \{i_d^t, i_d^\star\}]$ is not reducible with respect to \mathcal{P}_t . That is to say, all but at most a $\delta^{1/\ell}$ -fraction of $2 \times \cdots \times 2$ subtensors of T containing the entry $T(i_1^\star, \ldots, i_d^\star)$ are reducible with respect to \mathcal{P}_t , as desired.

In our final lemma, we observe that if T has many $2 \times \cdots \times 2$ subtensors involving a single nonzero entry, which are all reducible with respect to the same partition, then we can construct a reducible tensor that agrees with T on many entries.

Lemma 4.6. Fix $d \ge 1$ and $\delta > 0$. Let n_0 be sufficiently large in terms of d and δ . Consider a tensor $T: I_1 \times \cdots \times I_d \to \mathbb{F}$ (represented as a d-dimensional array) with $|I_1|, \ldots, |I_d| \ge n_0$, and a non-trivial partition $\{1, \ldots, d\} = J_1 \cup J_2$, and suppose $T(i_1^*, \ldots, i_d^*) \ne 0$ for some indices $i_1^* \in I_1, \ldots, i_d^* \in I_d$. Suppose that all but at most a δ -fraction of the $2 \times \cdots \times 2$ subtensors of T which contain the entry $T(i_1^*, \ldots, i_d^*)$ are reducible with respect to $\{J_1, J_2\}$.

Then one can make T reducible by changing up to a 2δ -fraction of its entries.

Proof. Suppose without loss of generality that $J_1 = \{1, \ldots, h\}$ and $J_2 = \{h + 1, \ldots, d\}$. Let $T' : I_1 \times \cdots \times I_d \to \mathbb{F}$ be the tensor defined by

$$T'(i_1,\ldots,i_d) = \frac{T(i_1,\ldots,i_h,i_{h+1}^{\star},\ldots,i_d^{\star}) T(i_1^{\star},\ldots,i_h^{\star},i_{h+1},\ldots,i_d)}{T(i_1^{\star},\ldots,i_d^{\star})}.$$

By definition, T' is reducible. Indeed, we have $T' = T_1 \otimes T_2$ for $T_1 : I_1 \times \cdots \times I_h \to \mathbb{F}$ given by $T_1(i_1, \ldots, i_h) = T(i_1, \ldots, i_h, i_{h+1}^\star, \ldots, i_d^\star)$ and $T_2 : I_{h+1} \times \cdots \times I_d \to \mathbb{F}$ given by $T_1(i_{h+1}, \ldots, i_d) = T(i_1^\star, \ldots, i_h^\star, i_{h+1}, \ldots, i_d)/T(i_1^\star, \ldots, i_d^\star)$. It remains to show that T' differs from T in at most a 2δ -fraction of its entries.

To see this, note that for all $i_1 \in I_1 \setminus \{i_1^*\}, \ldots, i_d \in I_d \setminus \{i_d^*\}$ such that $T[\{i_1, i_1^*\}, \ldots, \{i_d, i_d^*\}]$ is reducible with respect to $\{J_1, J_2\}$, we have

$$T(i_1, \dots, i_d) T(i_1^*, \dots, i_d^*) - T(i_1, \dots, i_h, i_{h+1}^*, \dots, i_d^*) T(i_1^*, \dots, i_h^*, i_{h+1}, \dots, i_d) = 0$$

(this can be interpreted as a determinant of a 2×2 matrix of rank at most 1, recalling Fact 4.2). Rearranging, we obtain $T(i_1, \ldots, i_d) = T'(i_1, \ldots, i_d)$. But, by assumption, the number of such choices of i_1, \ldots, i_d is at least $(1 - \delta) \prod_{i=1}^d (|I_i| - 1) \ge (1 - 2\delta) \prod_{i=1}^d |I_i|$. (Here we are using that I_1, \ldots, I_d are sufficiently large with respect to d and δ).

We now combine Lemmas 4.4 to 4.6 to prove Lemma 2.2.

Proof of Lemma 2.2. As in Lemma 4.5, let $\ell := 2^{d-1} - 1$, and note that then $\delta = (\varepsilon/2)^{2^{d-1}} = (\varepsilon/2)^{\ell+1}$. If at most a $\delta^{1/(\ell+1)}$ -fraction of the entries in T are nonzero, then T can be made zero by changing up to a ε fraction of the entries; we are done in this case. So, assume that T contains more than a $\delta^{1/(\ell+1)}$ -fraction of nonzero entries. We apply Lemma 4.4 with $r = 2^{d-1}$ and $c = 1/(\ell+1)$ to see that T has a nonzero entry $T(i_1^*, i_2^*, \ldots, i_d^*)$ such that all but at most a $\delta^{\ell/(\ell+1)}$ -fraction of the $2^{d-1} \times \cdots \times 2^{d-1}$ subtensors which contain this entry are reducible. By Lemma 4.5, there is a non-trivial partition $\{1, \ldots, d\} = J_1 \cup J_2$ such that all but at most a $\delta^{1/(\ell+1)}$ -fraction of the $2 \times \cdots \times 2$ subtensors of T containing the entry $T(i_1^*, i_2^*, \ldots, i_d^*)$, are reducible with respect to $\{J_1, J_2\}$. Then, Lemma 4.6 guarantees that T can be made reducible by changing a fraction of up to $2\delta^{1/(\ell+1)} = \varepsilon$ of its entries, as desired.

5. Linear subspaces of the variety of reducible tensors

In this section, we prove Lemma 2.6, characterising the linear subspaces of the variety of reducible tensors.

First, note that we can break up the set of reducible tensors according to the partition with respect to which they are reducible.

Fact 5.1. Fix disjoint sets I_1, \ldots, I_d , fix $\mathbb{F} \in \{\mathbb{R}, \mathbb{C}\}$ and let $r = \prod_{j=1}^d |I_j|$. Let $\mathcal{Z} \subseteq \mathbb{F}^{I_1} \otimes \cdots \otimes \mathbb{F}^{I_d}$ be the set of reducible $I_1 \times \cdots \times I_d$ tensors, and for a non-trivial partition $\{1, \ldots, d\} = J_1 \cup J_2$, let \mathcal{Z}_{J_1, J_2} be the set of $I_1 \times \cdots \times I_d$ tensors that are reducible with respect to $\{J_1, J_2\}$. Then

$$\mathcal{Z} = \bigcup_{\{J_1, J_2\}} \mathcal{Z}_{J_1, J_2},$$

where the union is over all $2^{d-1} - 1$ unordered non-trivial partitions $\{1, \ldots, d\} = J_1 \cup J_2$.

The above fact, together with Fact 4.2, makes Lemma 2.6(1) nearly immediate (i.e., the observation that the reducible tensors form a variety).

Proof of Lemma 2.6(1). By Fact 5.1, it suffices to prove that each \mathcal{Z}_{J_1,J_2} is an affine variety, and by Fact 4.2, it suffices to show that the set of $r_1 \times r_2$ matrices with rank at most 1 can be interpreted as an affine variety in $\mathbb{F}^{r_1r_2}$ (i.e., it is the zero locus of a system of polynomial equations). This is well-known to be true: a matrix has rank at most 1 if and only if every 2×2 submatrix has zero determinant.

Next, we prove Lemma 2.6(2), classifying the maximal linear subspaces of the variety of reducible tensors. In this proof, we need the simple algebraic geometry fact that linear subspaces are irreducible as affine varieties. This fact also holds in the case $\mathbb{F} = \mathbb{R}$, even though the underlying field is not algebraically closed (it holds over any infinite field).

Proof of Lemma 2.6(2). Recall that all linear subspaces of \mathbb{F}^r are irreducible as affine varieties. So, recalling Fact 5.1, any linear subspace of $\mathcal{Z} = \bigcup_{\{J_1, J_2\}} \mathcal{Z}_{J_1, J_2}$ must in fact be a linear subspace of some particular \mathcal{Z}_{J_1, J_2} . That is to say, it suffices to classify the maximal linear subspaces of \mathcal{Z}_{J_1, J_2} . By Fact 4.2, this is really the same as classifying the maximal linear subspaces of the variety of $r_1 \times r_2$ matrices with rank at most 1. That is to say, it suffices to prove the lemma in the case d = 2 (recalling that matrices are 2-dimensional tensors).

So, fix disjoint sets I_1, I_2 , and let \mathcal{V} be a linear subspace of the variety \mathcal{Z} of $I_1 \times I_2$ tensors (matrices) which have rank at most 1. Our objective is to prove that there is a "column vector" $\vec{u}^* \in \mathbb{F}^{I_1}$ such that

$$\mathcal{V} \subseteq \{ \vec{u}^{\star} \otimes \vec{x} : \vec{x} \in \mathbb{F}^{I_2} \}$$

or there is a "row vector" $\vec{v}^{\star} \in \mathbb{F}^{I_2}$ such that

$$\mathcal{V} \subseteq \{ \vec{y} \otimes \vec{v}^{\star} : \vec{y} \in \mathbb{F}^{I_1} \}.$$

Indeed, this will show that we have $\mathcal{V} = \{\vec{u}^* \otimes \vec{x} : \vec{x} \in \mathbb{F}^{I_2}\}$ or $\mathcal{V} = \{\vec{y} \otimes \vec{v}^* : \vec{y} \in \mathbb{F}^{I_1}\}$, whenever \mathcal{V} is a maximal linear subspace of \mathcal{Z} .

If \mathcal{V} contains only the zero matrix, the statement trivially holds. So, we can assume that \mathcal{V} contains a nonzero matrix, i.e., a matrix of rank exactly 1. This means that there are nonzero vectors $\vec{u}^* \in \mathbb{F}^{I_1}$ and $\vec{v}^* \in \mathbb{F}^{I_2}$ such that $\vec{u}^* \otimes \vec{v}^* \in \mathcal{V}$.

Claim 5.2. Let $T \in \mathcal{V}$. Then, T is of the form $\vec{u}^* \otimes \vec{x}$ for some $\vec{x} \in \mathbb{F}^{I_2}$, or $\vec{y} \otimes \vec{v}^*$ for some $\vec{y} \in \mathbb{F}^{I_1}$.

Proof of claim. $T \in \mathcal{V} \subseteq \mathcal{Z}$ has rank at most 1, so we can write $T = \vec{u} \otimes \vec{v}$ for some $\vec{u} \in \mathbb{F}^{I_1}$ and $\vec{v} \in \mathbb{F}^{I_2}$. We need to prove that \vec{u} is a multiple of \vec{u}^* or \vec{v} is a multiple of \vec{v}^* .

We can interpret $\vec{u} \otimes \vec{v} + u^* \otimes \vec{v}^*$ as a matrix product PQ, where P is the $|I_1| \times 2$ matrix with columns \vec{u}, \vec{u}^* and Q is the $2 \times |I_2|$ matrix with rows \vec{v}, \vec{v}^* . Since \mathcal{V} is a linear space, we know that $\vec{u} \otimes \vec{v} + u^* \otimes \vec{v}^* = T + u^* \otimes \vec{v}^* \in \mathcal{V} \subseteq \mathcal{Z}$ has rank at most 1, so at least one of P or Q has rank at most 1 (since the product of a $|I_1| \times 2$ matrix of rank 2 with a $2 \times |I_2|$ matrix or rank 2 always has rank 2 as well). That is to say, \vec{u} is a multiple of \vec{u}^* , or \vec{v} is a multiple of \vec{v}^* .

Now, let $\mathcal{V}_1 := \{\vec{u}^* \otimes \vec{x} : \vec{x} \in \mathbb{F}^{I_2}\}$ and $\mathcal{V}_2 := \{\vec{y} \otimes \vec{v}^* : \vec{y} \in \mathbb{F}^{I_1}\}$; they are both affine varieties (in fact, they are both linear subspaces). The above claim shows that $\mathcal{V} \subseteq \mathcal{V}_1 \cup \mathcal{V}_2$. Since \mathcal{V} is a linear space, it is irreducible, so $\mathcal{V} \subseteq \mathcal{V}_1$ or $\mathcal{V} \subseteq \mathcal{V}_2$, as desired.

6. k-multilinear forms

In this section we prove Lemma 2.4 and use it to prove Theorem 1.9.

Proof of Lemma 2.4. For each $i \in I_d = \{1, \ldots, n\}$, let $T_i : I_1 \times I_2 \times \cdots \times I_{d-1} \to \mathbb{F}$ be the tensor defined by $T_i(i_1, \ldots, i_{d-1}) = T(i_1, \ldots, i_{d-1}, i)$, so we can write $T\vec{\xi} = \xi_1 T_1 + \cdots + \xi_n T_n$. Recall that we can interpret an $I_1 \times I_2 \times \cdots \times I_{d-1}$ tensor as a vector in $\mathbb{F}^{|I_1|\cdots|I_{d-1}|}$, and let $\mathcal{Z} \subseteq \mathbb{F}^{|I_1|\cdots|I_{d-1}|}$ be the variety of reducible $I_1 \times I_2 \times \cdots \times I_{d-1}$ tensors. Then, the event that $T\vec{\xi}$ is reducible is precisely the event that $\xi_1 T_1 + \cdots + \xi_n T_n \in \mathcal{Z}$.

Also, note that every maximal affine-linear subspace of Z is in fact a linear subspace. To see this, note that multiplying by a nonzero scalar does not affect reducibility, and therefore we have $\lambda \vec{v} \in Z$ for all $\vec{v} \in Z$ and $\lambda \in \mathbb{F}$. So, for any affine-linear subspace $\mathcal{W} \subseteq Z$, the linear span $\{\lambda \vec{w} : \vec{w} \in \mathcal{W}, \lambda \in \mathbb{F}\}$ of the vectors in \mathcal{W} is a *linear* subspace in Z that contains \mathcal{W} .

So, by Theorem 2.5, $\mathbb{P}[T\vec{\xi} \text{ is reducible}] \leq n^{-1/2+\varepsilon/r} \leq n^{-1/2+\varepsilon}$ (i.e., **D2** holds), or there is a maximal linear subspace $\mathcal{W} \subseteq \mathcal{Z}$ such that all but $(\varepsilon/r)n$ of the T_i lie in \mathcal{W} . We may assume the latter property holds.

By Lemma 2.6, there is a subset $\emptyset \subsetneq J \subsetneq \{1, 2, \dots, d-1\}$ and a tensor $T^* \in \bigotimes_{j \in J} \mathbb{F}^{I_j}$ such that

$$\mathcal{W} = \left\{ T^{\star} \otimes T' : \ T' \in \bigotimes_{j \in \{1, \dots, d-1\} \setminus J} \mathbb{F}^{I_j} \right\}.$$

Suppose without loss of generality that $J = \{1, \ldots, j\}$ for some $j \in \{1, \ldots, d-2\}$. Now, we know that at least a $(1 - \varepsilon/r)$ -fraction of the T_i lie in \mathcal{W} , so for at least a $(1 - \varepsilon)$ -fraction of the r-element subsets $I'_d \subseteq I_d$, we have $T_i \in \mathcal{W}$ for each $i \in I'_d$. To prove **D1**, it suffices to prove that for all such subsets I'_d , the corresponding $r \times \cdots \times r$ subtensor $T[I_1, \ldots, I_{d-1}, I'_d]$ is reducible. In fact, this is nearly immediate: for each $i \in I'_d$ we can write $T_i = T^* \otimes T'_i$, so defining $T' : I_{j+1} \times \cdots \times I_d \to \mathbb{F}$ by $T'(i_{j+1}, \ldots, i_d) = T_{i_d}(i_{j+1}, \ldots, i_{d-1})$, we have $T[I_1, \ldots, I_{d-1}, I'_d] = T^* \otimes T'$.

Now, we turn to the proof of Theorem 1.9. It is a direct consequence of the following slightly more general result.

Theorem 6.1. Fix $d \ge 2$ and $\varepsilon > 0$, and let $\mathbb{F} \in \{\mathbb{R}, \mathbb{C}\}$. Let *n* be sufficiently large (in terms of ε, d), and consider a tensor $T: I_1 \times \cdots \times I_d \to \mathbb{F}$ for some partition $\{1, \ldots, n\} = I_1 \cup \cdots \cup I_d$, and the associated *d*-multilinear form $f \in \mathbb{F}[x_1, \ldots, x_n]$. Then, at least one of the following holds:

F1 T can be made reducible by changing up to εn^d entries, or

F2 letting $\vec{\xi} = (\xi_1, \dots, \xi_n) \in \{-1, 1\}^n$ be a vector of *i.i.d.* Rademacher random variables, we have

$$\sup_{z\in\mathbb{F}}\mathbb{P}[f(\xi_1,\ldots,\xi_n)=z]\leq n^{-1+\varepsilon}$$

Our proof of Theorem 6.1 is by induction on d. The base case d = 2 was essentially proved by Costello [9] (in fact, even with a stronger¹² notion of "close to reducible" than our property **F1**). Specifically, the following theorem is a direct consequence of [9, Theorem 5].

Theorem 6.2. Fix $\gamma > 0$ and $\mathbb{F} \in \{\mathbb{R}, \mathbb{C}\}$, and let r_0 be sufficiently large in terms of γ . For any $m, m' \in \mathbb{N}$, consider a matrix $A \in \mathbb{F}^{m \times m'}$ such that every row has at least $r \geq r_0$ nonzero entries. Let $\vec{\xi} = (\xi_1, \ldots, \xi_m) \in \{-1, 1\}^m$ and $\vec{\xi'} = (\xi'_1, \ldots, \xi'_{m'}) \in \{-1, 1\}^{m'}$ be independent (column) vectors with *i.i.d.* Rademacher entries, and suppose that there is a function $\varphi : \mathbb{F}^{m'} \to \mathbb{F}$ such that

$$\mathbb{P}\left[\vec{\xi}^{T} A \vec{\xi}' = \varphi(\vec{\xi}')\right] \ge r^{-1+\gamma}.$$

Then, A can be turned into a matrix of rank at most 1 by changing up to $\gamma r(m+m')$ entries.

Compared to Theorem 6.1, note that Theorem 6.2 has an additional assumption requiring many nonzero entries in every row of the relevant matrix. We take a moment to state a slight strengthening of the d = 2 case of Theorem 6.1 (this strengthening will be useful later in the paper) and explain how to deduce this statement from Costello's work (using an argument as in [9, Remark 1]).

¹²We remark that Costello's notion of "close to reducible" (which we do not state here) seems to be genuinely stronger than **F1**, in the sense that we do not know how to directly deduce [9, Theorem 5] from the d = 2 case of Theorem 6.1. As discussed in **Q2** in Section 1.1, it may be interesting to investigate further the various different notions of "smallness" or "closeness", and the implications between them.

Theorem 6.3. Fix $\varepsilon > 0$ and $\mathbb{F} \in \{\mathbb{R}, \mathbb{C}\}$. Consider $m, m' \in \mathbb{N}$ such that m + m' is sufficiently large in terms of ε , and consider a matrix $A \in \mathbb{F}^{m \times m'}$. Then, at least one of the following holds:

F1' A can be made into a rank-1 matrix by changing up to $\varepsilon(m+m')^2$ entries, or

F2' letting $\vec{\xi} = (\xi_1, \dots, \xi_m) \in \{-1, 1\}^m$ be a vector of i.i.d. Rademacher random variables, and independently letting $\vec{\xi}' = (\xi'_1, \dots, \xi'_{m'}) \in \{-1, 1\}^{m'}$ be a vector of i.i.d. Rademacher or lazy Rademacher random variables, we have

$$\mathbb{P}[\vec{\xi}^{T}A\vec{\xi}' = \psi(\vec{\xi}')] \le (m+m')^{-1+\varepsilon}$$

for any function $\psi : \mathbb{F}^{m'} \to \mathbb{F}$.

Proof. Let n = m + m' and assume that f does not satisfy **F2'**; that is to say,

$$\mathbb{P}\left[\vec{\xi}^{T}A\vec{\xi'} = \psi(\vec{\xi'})\right] > n^{-1+\varepsilon}.$$

for some function $\psi : \mathbb{F}^{m'} \to \mathbb{F}$. We will show that **F1**' holds. In other words, we will show that we can change at most εn^2 entries of A to obtain a matrix with rank at most 1.

Let I be the set of $i \in \{1, ..., m\}$ such that row i of A has at least εn nonzero entries, and let $I^* = \{1, \ldots, m\} \setminus I$. Let $A' \in \mathbb{F}^{|I| \times m'}$ be the submatrix of A containing the rows indexed by I, and let $A^* \in \mathbb{F}^{|I^*| \times m'}$ be the submatrix of A containing the rows indexed by I^* . Define the random function $\varphi: \{-1,1\}^{m'} \to \mathbb{F}$ by $\varphi(\vec{x}) = \psi(\vec{x}) - \vec{\xi}[I^*]^T A^* \vec{x}$ (this function depends on the random vector $\vec{\xi}[I^*]$). So, we have

$$u^{-1+\varepsilon} < \mathbb{P}\Big[\vec{\xi}^T A \vec{\xi'} = \psi(\vec{\xi'})\Big] = \mathbb{P}\Big[\vec{\xi}[I]^T A' \vec{\xi'} = \varphi(\vec{\xi'})\Big].$$

Consequently, there exists an outcome of $\xi[I^*]$ such that

$$\mathbb{P}\Big[\vec{\xi}[I]^T A' \vec{\xi}' = \varphi(\vec{\xi}') \mid \vec{\xi}[I^*]\Big] > n^{-1+\varepsilon} > (\varepsilon n)^{-1+\varepsilon/2}.$$

Applying Theorem 6.2 (with $r = \varepsilon n$ and $\gamma = \varepsilon/2$) in the conditional probability space where we condition on this particular outcome of $\tilde{\xi}[I^*]$, we see that we can change at most $(\varepsilon/2)(\varepsilon n)(|I|+m') \leq (\varepsilon n)(|I|+m')$ entries of A' to obtain a matrix of rank at most 1. Also, A^* has at most $(\varepsilon n)|I^*|$ nonzero entries by definition, so we can make A^* into the all-zero matrix by changing at most $(\varepsilon n)|I^*| = (\varepsilon n)(m-|I|)$ entries. Recalling the definitions of A' and A^{*}, we see that we can change at most $(\varepsilon n)(|I|+m')+(\varepsilon n)(m-|I|)=$ $(\varepsilon n)(m+m') = \varepsilon n^2$ entries of A to obtain a matrix with rank at most 1, as desired.

We now complete the proof of Theorem 6.1 by presenting the inductive step. Briefly speaking, we first expose $\xi[I_d]$, and obtain a (d-1)-multilinear form in the variables x_i for $i \in I_1 \cup \cdots \cup I_{d-1}$. Equivalently, we "collapse" T, based on $\vec{\xi}[I_d]$, to a (d-1)-dimensional tensor T'. Then, we will apply Lemma 2.4 if T' tends to be close to reducible and apply the inductive hypothesis to T' otherwise.

Proof of Theorem 6.1. We proceed by induction on d; the base case d = 2 is handled by Theorem 6.3. Fix $d \geq 3$ and suppose the statement holds for all (d-1)-multilinear forms. Recall that $f(x_1, \ldots, x_n) =$ $\sum_{i_1 \in I_1} \cdots \sum_{i_d \in I_d} T(i_1, \dots, i_d) x_{i_1} \dots x_{i_d} \text{ for all } \vec{x} \in \mathbb{F}^n. \text{ Our goal is to prove that } \mathbf{F1} \text{ or } \mathbf{F2} \text{ holds.}$ First, recall that we are assuming that n is large with respect to d, ε . We may also assume that

$$|I_1|\cdots|I_d| \ge \varepsilon n^d. \tag{6.1}$$

Indeed, otherwise, T has at most εn^d entries, and F1 trivially holds (as we can change all the entries of T to zero to make T reducible).

Now, let $I' = I_1 \cup \cdots \cup I_{d-1}$, and note that for any outcome of $\vec{\xi}[I_d]$, we can view f as a (d-1)multilinear form $f' \in \mathbb{F}[(x_i)_{i \in I'}]$, in the variables x_i for $i \in I'$. Note that the coefficient tensor of f' is precisely $T\vec{\xi}[I_d]$, in the notation of Lemma 2.4. For convenience we write $T' = T\vec{\xi}[I_d]$. The high-level idea of the proof is to fix an outcome of $\vec{\xi}[I_d]$ (thereby fixing f'), and to apply the induction hypothesis to f'. When f' is far from being reducible, this gives us the desired probability bound, so the main challenge is to carefully study the event that f' is close to reducible. Actually, we need to break down this event further, depending on whether f' is close to the zero polynomial or close to some other (far-from-zero) reducible polynomial. Specifically, let $\delta = (\varepsilon/2)^{2^{d-1}}$ and consider the following two events.

- Let $\varepsilon_1 = \varepsilon^3/32$ and let \mathcal{E}_1 be the event that T' has at most $\varepsilon_1|I'|^{d-1}$ nonzero entries.
- Let $\varepsilon_2 = 2^{-(d-1)^2} \delta \varepsilon / 4$, and let \mathcal{E}_2 be the event that T' can be made reducible by changing at most $\varepsilon_2 |I'|^{d-1}$ entries.

Decomposing the anticoncentration probability. Let $\mathcal{E}_1^c, \mathcal{E}_2^c$ be the complements of \mathcal{E}_1 and \mathcal{E}_2 . Recall that our goal is to prove that **F1** or **F2** holds. Assume that **F2** fails, that is to say, for some $z \in \mathbb{F}$,

$$\mathbb{P}\big[f(\vec{\xi}) = z\big] > n^{-1+\varepsilon}$$

The probability $\mathbb{P}[f(\vec{\xi}) = z]$ can be decomposed as

$$\mathbb{P}\big[f(\vec{\xi}) = z \,\big|\,\mathcal{E}_1\big] \cdot \mathbb{P}\big[\mathcal{E}_1\big] \,+\, \mathbb{P}\big[f(\vec{\xi}) = z \,\big|\,\mathcal{E}_1^{\mathrm{c}} \cap \mathcal{E}_2\big] \cdot \mathbb{P}\big[\mathcal{E}_1^{\mathrm{c}} \cap \mathcal{E}_2\big] \,+\, \mathbb{P}\big[f(\vec{\xi}) = z \,\big|\,\mathcal{E}_1^{\mathrm{c}} \cap \mathcal{E}_2^{\mathrm{c}}\big] \cdot \mathbb{P}\big[\mathcal{E}_1^{\mathrm{c}} \cap \mathcal{E}_2^{\mathrm{c}}\big].$$

and we deduce

$$\mathbb{P}[f(\vec{\xi}) = z] \leq \mathbb{P}[\mathcal{E}_1] + \mathbb{P}[f(\vec{\xi}) = z \mid \mathcal{E}_1^c \cap \mathcal{E}_2] \cdot \mathbb{P}[\mathcal{E}_2] + \mathbb{P}[f(\vec{\xi}) = z \mid \mathcal{E}_1^c \cap \mathcal{E}_2^c].$$
(6.2)

Now, if $\vec{\xi}[I_d]$ satisfies $\mathcal{E}_1^c \cap \mathcal{E}_2^c \subseteq \mathcal{E}_2^c$, then T' cannot be made reducible by changing at most $\varepsilon_2 |I'|^{d-1}$ coefficients. So, if we fix an outcome of $\vec{\xi}[I_d]$ satisfying \mathcal{E}_2^c , then we can condition on this outcome of $\vec{\xi}[I_d]$ and apply the induction hypothesis to f' and T' in the resulting conditional probability space, to obtain

$$\mathbb{P}\big[f(\vec{\xi}) = z \,\big|\,\vec{\xi}[I_d]\big] = \mathbb{P}\big[f'(\vec{\xi}[I']) = z \,\big|\,\vec{\xi}[I_d]\big] \le |I'|^{-1+\varepsilon_2} \le n^{-1+\varepsilon}/3.$$

(Here we used that $|I'| \ge \varepsilon n$, which follows from (6.1)). Averaging over $\xi[I_d]$ satisfying \mathcal{E}_2^c , we deduce that

$$\mathbb{P}\big[f(\vec{\xi}) = z \,\big|\, \mathcal{E}_2^{\mathrm{c}}\big] \le n^{-1+\varepsilon}/3.$$

Similarly, if we fix an outcome of $\vec{\xi}[I_d]$ satisfying $\mathcal{E}_1^c \cap \mathcal{E}_2 \subseteq \mathcal{E}_1^c$, then f' has at least $\varepsilon_1 |I'|^{d-1}$ nonzero coefficients, so Theorem 1.1 yields

$$\mathbb{P}\big[f(\vec{\xi}) = z \,\big|\, \vec{\xi}[I_d]\big] = \mathbb{P}\big[f(\vec{\xi}) = z \,\big|\, \vec{\xi}[I_d]\big] \le |I'|^{-1/2+\varepsilon_1} \le n^{-1/2+\varepsilon/2}.$$

Averaging over all $\vec{\xi}[I_d]$ satisfying $\mathcal{E}_1^c \cap \mathcal{E}_2$, we obtain

$$\mathbb{P}\big[f(\vec{\xi}) = z \,\big|\, \mathcal{E}_1^{\mathrm{c}} \cap \mathcal{E}_2\big] \le |I'|^{-1/2 + \varepsilon_1} \le n^{-1/2 + \varepsilon/2}.$$

Now, from (6.2), it follows that

$$^{-1+\varepsilon} < \mathbb{P}[f(\vec{\xi}) = z] \le \mathbb{P}[\mathcal{E}_1] + n^{-1/2+\varepsilon/2} \cdot \mathbb{P}[\mathcal{E}_2] + n^{-1+\varepsilon}/3.$$

Thus, we must have $\mathbb{P}[\mathcal{E}_1] > n^{-1+\varepsilon}/3$ or $\mathbb{P}[\mathcal{E}_2] > n^{-1/2+\varepsilon/2}/3$. In both cases, we will prove that T can be made reducible by changing $\varepsilon |I_1| \cdots |I_d| \leq \varepsilon n^d$ of its entries; that is, **F1** holds. In the first case we will use Theorem 3.2 (i.e., Halász' inequality), and in the second case we will use Lemma 2.4.

Case 1: f' is likely to be close to the zero polynomial. In this case, we suppose that $\mathbb{P}[\mathcal{E}_1] > n^{-1+\varepsilon}/3$. Let $Q = I_1 \times I_2 \times \cdots \times I_{d-1}$; it follows from (6.1) that $|Q| > \varepsilon n^{d-1}$. Recalling Fact 4.2, let us view T as a $Q \times I_d$ matrix, which we refer to as M.

Translating from tensor to matrix language, in this case we are assuming that with probability exceeding $n^{-1+\varepsilon}/3$, the number of nonzero entries in the matrix-vector product $M\vec{\xi}[I_d]$ is at most $\varepsilon_1|I'|^{d-1} \leq (\varepsilon^2/32)|Q|$. By Lemma 3.1 (with r = 2), applied to the set of zero entries of $M\vec{\xi}[I_d]$, all but at most a $\varepsilon^2/8$ fraction of the pairs of distinct $\vec{i}, \vec{i'} \in Q$ satisfy

$$\mathbb{P}\big[(M\vec{\xi}[I_d])_{\vec{i}} = (M\vec{\xi}[I_d])_{\vec{i}'} = 0\big] = \mathbb{P}\big[M[\{\vec{i},\vec{i}'\},I_d]\vec{\xi}[I_d] = \vec{0}\,\big] > (n^{-1+\varepsilon}/3)/2 > |I_d|^{-1+\varepsilon/2}.$$

(Here we write $M[\{\vec{i}, \vec{i'}\}, I_d]$ for the $2 \times |I_d|$ matrix containing just the rows of M indexed by $\vec{i}, \vec{i'}$ and we use that $|I_d| > \varepsilon n$, which follows from (6.1).) For each such $\vec{i}, \vec{i'}$, by Theorem 3.2, $M[\{\vec{i}, \vec{i'}\}, I_d]$ contains at most $O(|I_d|^{1-\varepsilon/2})$ disjoint nonsingular 2×2 submatrices, so by Fact 3.4 it has at most $O(|I_d|^{1-\varepsilon/2})|I_d| < (\varepsilon^2/8) \binom{|I_d|}{2}$ nonsingular 2×2 submatrices in total. Summing over all these $\vec{i}, \vec{i'}$, the fraction of 2×2 submatrices of M which are nonsingular is less than $\varepsilon^2/8 + \varepsilon^2/8 \le \varepsilon^2/4$. So, by the d = 2 case of Lemma 2.2, we can change at most an ε -fraction of the entries of M to obtain a matrix with rank at most 1. Translating back into tensor language, T can be made reducible by changing at most an ε -fraction of its entries. This proves **F1**, as desired.

Case 2: f' is likely to be close to reducible. In this case, we suppose that $\mathbb{P}[\mathcal{E}_2] > n^{-1/2+\varepsilon/2}/3$. This means, with probability exceeding $n^{-1/2+\varepsilon/2}/3$, there is a reducible tensor $\tilde{T}: I_1 \times \cdots \times I_{d-1} \to \mathbb{F}$ such that T' differs from \tilde{T} in at most $\varepsilon_2 |I'|^{d-1} \leq \varepsilon_2 n^{d-1} \leq 2^{-(d-1)^2} (\delta/4) |I_1| \cdots |I_{d-1}|$ entries (here, we used that $|I_1| \cdots |I_{d-1}| \geq \varepsilon n^{d-1}$, which follows from (6.1)). When this occurs, at least a $(1 - \delta/4)$ fraction of the $2^{d-1} \times \cdots \times 2^{d-1}$ subtensors of T' are the same as the corresponding $2^{d-1} \times \cdots \times 2^{d-1}$ subtensor of \tilde{T} , so are reducible. (Here we used that a $2^{d-1} \times \cdots \times 2^{d-1}$ subtensor has $2^{(d-1)^2}$ entries). By Lemma 3.1 (with r = 1) applied to the set of reducible $2^{d-1} \times \cdots \times 2^{d-1}$ subtensors of T', we see that for at least a $(1 - \delta/2)$ -fraction of choices of 2^{d-1} -element subsets $Q_1 \subseteq I_1, \ldots, Q_{d-1} \subseteq I_{d-1}$, we have

$$\mathbb{P}[T'[Q_1, \dots, Q_{d-1}] \text{ is reducible}] > (n^{-1/2+\varepsilon}/3)/2 \ge |I_d|^{-1/2+\delta/2}$$

(here we are using that $|I_d| \geq \varepsilon n$, which follows from (6.1)). For each such Q_1, \ldots, Q_{d-1} , we apply Lemma 2.4: noting that $T'[Q_1, \ldots, Q_{d-1}] = T[Q_1, \ldots, Q_{d-1}, I_d] \vec{\xi}[I_d]$ in the notation of Lemma 2.4, we see that at least a $(1 - \delta/2)$ -fraction of the $2^{d-1} \times \cdots \times 2^{d-1}$ subtensors of $T[Q_1, \ldots, Q_{d-1}, I_d]$ are reducible. Summing over all these Q_1, \ldots, Q_{d-1} , we see that the fraction of $2^{d-1} \times \cdots \times 2^{d-1}$ subtensors of T that are reducible is at least $(1 - \delta/2)(1 - \delta/2) > 1 - \delta$. In the end, Lemma 2.2 implies that T can be made reducible by changing at most $\varepsilon |I_1| \cdots |I_d| \leq \varepsilon n^d$ of its entries. This proves **F1**, as desired. \Box

7. A local-to-global lemma for low-rank symmetric matrices

In this section we prove Lemma 2.7: if almost all $r \times r$ submatrices of a symmetric matrix A are singular, then A can be made to have rank less than r by changing only a few entries, while maintaining symmetry. Recall that $||A||_0$ denotes the number of nonzero entries in a matrix A.

We split the proof of Lemma 2.7 into the following two lemmas. Roughly speaking, we first approximate A by a matrix of rank less than r, which may not be symmetric. However, this matrix is necessarily *almost* symmetric. In the second step we adjust it to be symmetric while maintaining the rank. Most of the difficulty lies in the second step.

Lemma 7.1. Consider integers $n, m \ge r \ge 1$, let $\alpha \in [0,1]$ and let \mathbb{F} be a field. Let $A \in \mathbb{F}^{n \times m}$ be a matrix such that all but at most an α -fraction of its $r \times r$ submatrices are singular. Then, there exists a matrix $B \in \mathbb{F}^{n \times m}$ of rank less than r such that $||A - B||_0 \le \alpha^{1/r} nm$.

Lemma 7.2. Consider integers $n \ge q \ge 1$, let $\rho \in [0,1]$ and let \mathbb{F} be a field. Let $A \in \mathbb{F}^{n \times n}$ be a matrix of rank at most q such that $||A - A^T||_0 \le \rho n^2$. Then, there exists a symmetric matrix $B \in \mathbb{F}^{n \times n}$ of rank at most q such that $||A - B||_0 \le O(q^4 \rho n^2)$.

Lemma 2.7 follows immediately from Lemma 7.1 and Lemma 7.2, as follows.

Proof of Lemma 2.7. By Lemma 7.1, we first obtain a matrix $B \in \mathbb{F}^{n \times n}$ with rank less than r and $||A - B||_0 \leq \alpha^{1/r} n^2$. Since A is symmetric, it holds that

$$||B - B^T||_0 \le ||B - A^T||_0 + ||A^T - B^T||_0 = ||B - A||_0 + ||A - B||_0 \le 2\alpha^{1/r} n^2.$$

Then, Lemma 7.2 (with q = r - 1) guarantees a symmetric matrix $A' \in \mathbb{F}^{n \times n}$ of rank less than r such that $\|B - A'\|_0 \leq O(r^4 \alpha^{1/r} n^2)$. Also, $\|A - A'\|_0 \leq \|A - B\|_0 + \|B - A'\|_0 \leq O(r^4 \alpha^{1/r} n^2)$, as desired. \Box

The rest of this section is devoted to proving Lemma 7.1 and Lemma 7.2. We first give a short proof of Lemma 7.1. For the rest of this section, we use the notation $[n] = \{1, \ldots, n\}$.

Proof of Lemma 7.1. First, find the minimum $k \ge 1$ such that all but at most an $\alpha^{k/r}$ -fraction of the $k \times k$ submatrices of the matrix A are singular. Clearly, $k \le r$ by assumption. If k = 1, then $||A||_0 \le \alpha^{1/r} nm$, and we are done (by taking B = 0). From now on, we assume that $2 \le k \le r$.

Second, we claim that there exist sets $I \subseteq [n], J \subseteq [m]$, both of size k - 1, such that A[I, J] is nonsingular and $A[I \cup \{i\}, J \cup \{j\}]$ is nonsingular for at most $\alpha^{1/r}(n - k + 1)(m - k + 1)$ choices of $(i, j) \in ([n] \setminus I) \times ([n] \setminus J)$. Indeed, suppose not. Then, whenever A[I, J] is nonsingular, more than $\alpha^{1/r}(n - k + 1)(m - k + 1)$ choices of $(i, j) \in ([n] \setminus I) \times ([n] \setminus J)$ are such that $A[I \cup \{i\}, J \cup \{j\}]$ is nonsingular. In addition, the minimality of k implies there are more than $\alpha^{(k-1)/r} {n \choose k-1} {m \choose k-1}$ choices of I, J such that A[I, J] is nonsingular. In total, we acquire more than

$$\alpha^{(k-1)/r} \binom{n}{k-1} \binom{m}{k-1} \cdot \alpha^{1/r} (n-k+1)(m-k+1) = k^2 \alpha^{k/r} \binom{n}{k} \binom{m}{k}$$

choices of I, J, i, j such that $A[I \cup \{i\}, J \cup \{j\}]$ is a nonsingular $k \times k$ submatrix of A. Since each submatrix is counted at most k^2 times, the number of nonsingular $k \times k$ submatrices of A is larger than $\alpha^{k/r} \binom{n}{k} \binom{m}{k}$, contradicting our definition of k.

We have established that there is some choice of I and J such that A[I, J] is nonsingular, and $A[I \cup \{i\}, J \cup \{j\}]$ is nonsingular for at most $\alpha^{1/r}(n-k+1)(m-k+1)$ choices of (i, j). Fix such I, J, and let $B \in \mathbb{F}^{n \times m}$ be the unique matrix with the same row space as A[I, [m]], such that B[[n], J] = A[[n], J]. (In a bit more detail: since A[I, J] is invertible, for each row A[i, [m]] of A there is a unique vector in the row space of A[I, [m]] which agrees with A[i, [m]] in the coordinates indexed by J; we take that to be

the corresponding row of B.) Note that B[I, [m]] = A[I, [m]] and B[[n], J] = A[[n], J], and furthermore $\operatorname{rank} B = \operatorname{rank}(A[I, [m]]) = |I| = k - 1 < r.$

Now, take any $(i,j) \in ([n] \setminus I) \times ([n] \setminus J)$ such that $A[I \cup \{i\}, J \cup \{j\}]$ is singular; there are at least $(1 - \alpha^{1/r})(n - k + 1)(m - k + 1)$ such (i, j), thanks to the choice of (I, J). Consider the two matrices $A[I \cup \{i\}, J \cup \{j\}]$ and $B[I \cup \{i\}, J \cup \{j\}]$. Their ranks are both less than k, i.e. $\det(A[I \cup \{i\}, J \cup \{j\}]) =$ $\det(B[I \cup \{i\}, J \cup \{j\}]) = 0$. Hence, $\det(A[I \cup \{i\}, J \cup \{j\}]) - \det(B[I \cup \{i\}, J \cup \{j\}]) = 0$. As discussed above, these two matrices are identical except that, possibly, $A[i,j] \neq B[i,j]$. This means that

 $0 = \left| \det(A[I \cup \{i\}, J \cup \{j\}]) - \det(B[I \cup \{i\}, J \cup \{j\}]) \right| = \left| \det(A[I, J])(A[i, j] - B[i, j]) \right|.$

But $\det(A[I,J]) \neq 0$ because A[I,J] is nonsingular. So, A[i,j] = B[i,j]. Recall also that A[i,j] = B[i,j] whenever $i \in I$ or $j \in J$. Therefore, $||A - B||_0 \leq \alpha^{1/r} (n - k + 1)(m - k + 1) < \alpha^{1/r} nm$.

Lemma 7.2 is more complicated. Our strategy is as follows. First, we pass from A to a large submatrix A[I, I], whose rank is "robust" (i.e., the rank is maintained whenever we delete a few rows and columns). Then, we show how to use this robustness to approximate A[I, I] with a symmetric matrix with the same row space. Finally, we show how this approximation can be extended to all of A.

The first of these steps (passing to a large submatrix whose rank is robust) is rather simple, as follows.

Lemma 7.3. Consider integers $n \ge q \ge 1$, let $\gamma > 0$, and let \mathbb{F} be a field. Let $A \in \mathbb{F}^{n \times n}$ be a matrix with rank at most q. Then, there exists $k \in \{0, \ldots, q\}$ and $I \subseteq [n]$ of size $|I| \geq (1 - q\gamma)n$ such that $\operatorname{rank}(A[I', I']) = \operatorname{rank}(A[I, I]) = k \text{ for all } I' \subseteq I \text{ of size } |I'| \ge |I| - \gamma n.$

Proof. Take the minimum $k \in \{0, \ldots, q\}$ such that $\operatorname{rank}(A[I, I]) \leq k$ holds for some $I \subseteq [n]$ of size $|I| \ge (1 - (q - k)\gamma)n$; to see that such k exists, consider k = q. We claim this k and I satisfy the conclusion of the lemma. Clearly, $|I| \ge (1 - q\gamma)n$. Fix an arbitrary $I' \subseteq I$ with $|I'| \ge |I| - \gamma n$. Then, since $|I'| \ge (1 - (q - k + 1)\gamma)n$, the minimality of k implies that rank $(A[I', I']) \ge k$. That is to say, $k \ge \operatorname{rank}(A[I, I]) \ge \operatorname{rank}(A[I', I']) \ge k$, so $\operatorname{rank}(A[I', I']) = \operatorname{rank}(A[I, I]) = k$.

The next lemma shows that when a matrix "has its rank robustly", we can approximate it by a symmetric matrix with the same row space. The proof approach is similar to the approach of the second and third authors in [35, Section 5.7].

Lemma 7.4. Consider integers $n \ge 1$ and $r \ge 0$, let $\rho \in [0,1]$ and let \mathbb{F} be a field. Let $A \in \mathbb{F}^{n \times n}$ be a matrix with $||A - A^T||_0 \leq \rho n^2$, such that $\operatorname{rank}(A[I,I]) = r$ for all $I \subseteq [n]$ of size $|I| \geq (1 - r\sqrt{\rho})n$. Then, there exists a symmetric matrix $B \in \mathbb{F}^{n \times n}$ such that $||A - B||_0 \leq (r^2 + 1)\rho n^2$, and such that the row spaces of A and B are the same.

Proof. Note that in the case r = 0 we must have A = 0, and so we can take B = 0. So let us from now on assume $r \geq 1$.

For $i, j \in [n]$, say that the pair (i, j) is bad if $a_{ij} \neq a_{ji}$. For each $1 \leq i \leq n$, write $N(i) := \{1 \leq j \leq i \leq n\}$ $n: a_{ij} \neq a_{ji}$. Also, define I_{bad} to be the set of indices i such that $|N(i)| \geq \sqrt{\rho}n$; hence,

$$|I_{\text{bad}}| \le ||A - A^T||_0 / \sqrt{\rho}n \le \sqrt{\rho}n.$$

First, we claim that there exists $V \subseteq [n] \setminus I_{\text{bad}}$ of size r such that A[V, V] is symmetric and rank(A[[n], V]) =r. To this end, we iteratively find, for each $\ell \in [r]$, a sequence $v_1, \ldots, v_\ell \in [n] \setminus I_{\text{bad}}$ such that $\operatorname{rank}(A[[n], \{v_1, \ldots, v_\ell\}]) = \ell$ and such that no (v_i, v_j) is a bad pair. Suppose $\ell \in \{0, 1, \ldots, r-1\}$ and we have found v_1, \ldots, v_ℓ with these properties. For all $i \in \{1, \ldots, \ell\}$, we have $|N(v_i)| \leq \sqrt{\rho}n$ because $v_i \notin I_{\text{bad}}$. Thus, $\tilde{I} := [n] \setminus (I_{\text{bad}} \cup \bigcup_{i=1}^{\ell} N(v_i))$ has size $|\tilde{I}| \ge n - \sqrt{\rho}n - \ell\sqrt{\rho}n \ge (1 - r\sqrt{\rho})n$. By our assumption on A, $\operatorname{rank}(A[[n], \tilde{I}]) \ge \operatorname{rank}(A[\tilde{I}, \tilde{I}]) = r > \ell$, so there exists $v_{\ell+1} \in \tilde{I}$ such that $\operatorname{rank}(A[[n], \{v_1, v_2, \dots, v_{\ell+1}\}]) = \ell + 1$. In addition, for each $i \in [\ell]$, the pair $(v_i, v_{\ell+1})$ is not bad because $v_{\ell+1} \in I \subseteq [n] \setminus N(v_i)$. This completes the iteration step. Having found v_1, \ldots, v_r at the end of the iterative procedure, we set $V = \{v_1, \ldots, v_r\}$, and note that A[V, V] is symmetric since none of the pairs (v_i, v_j) is bad.

Fix V as above. We next claim that $\operatorname{rank}(A[V,V]) = r$, i.e., A[V,V] is invertible. To see this, first note that since $\operatorname{rank}(A[[n], V]) = r = \operatorname{rank}(A)$, the column vectors of A[[n], V] span those of A = A[[n], [n]]. Let $I := [n] \setminus \bigcup_{v \in V} N(v)$; it holds that

$$|I| \ge n - \sum_{i=1}^{r} |N(v_i)| \ge (1 - r\sqrt{\rho})n.$$
(7.1)

The above discussion indicates that the column vectors of A[I, V] span those of A[I, I], so rank $(A[I, V]) \ge$ $\operatorname{rank}(A[I, I])$. Similarly, we have $\operatorname{rank}(A[V, V]) \geq \operatorname{rank}(A[V, I])$. In addition, the definition of I guarantees that $a_{vi} = a_{iv}$ for all $v \in V$ and all $i \in I$. In short, $A[V, I] = A[I, V]^T$. So,

 $\operatorname{rank}(A[V,V]) \ge \operatorname{rank}(A[V,I]) = \operatorname{rank}(A[I,V]) \ge \operatorname{rank}(A[I,I]) = r,$

where the last inequality follows as $|I| \ge (1 - r\sqrt{\rho})n$. That is to say, $\operatorname{rank}(A[V,V]) = r$, i.e., A[V,V] is invertible.

Now, take $B := A[V, [n]]^T A[V, V]^{-1} A[V, [n]]$. Note that B is symmetric, and rank $(B) \leq |V| = r$. By the preceding discussion,

$$B[V,V] = A[V,V]^T A[V,V]^{-1} A[V,V] = A[V,V]$$

has rank r, so rank(B) = r. Also, the row vectors of B are spanned by those of A. This, plus the fact that $\operatorname{rank}(B) = r = \operatorname{rank}(A)$, demonstrates that the row space of B is the same as that of A.

To complete the proof, it remains to show that $||A - B||_0 \leq (r^2 + 1)\rho n^2$. Since $r = \operatorname{rank}(A) \geq r$ $\operatorname{rank}(A[V, [n]]) \geq \operatorname{rank}(A[V, V]) = r$, we see that A[V, [n]] has the same row space as A, i.e., for each $i \in [n]$, we have $A[i,[n]] = \vec{x}_i A[V,[n]]$ for some $\vec{x}_i \in \mathbb{F}^V$. In particular, $A[i,V] = \vec{x}_i A[V,V]$, i.e. $\vec{x}_i = A[i,V]A[V,V]^{-1}$. Hence, $A[i,[n]] = A[i,V]A[V,V]^{-1}A[V,[n]]$. By considering all $i \in I$, we deduce

$$A[I, [n]] = A[I, V]A[V, V]^{-1}A[V, [n]] = A[V, I]^T A[V, V]^{-1}A[V, [n]] = B[I, [n]].$$
(7.2)

Writing $I^{c} = [n] \setminus I$, note that for any $i \in I^{c}$ and $j \in I$, the above equality implies that $b_{ij} = b_{ji} = a_{ji}$. So, we have $b_{ij} \neq a_{ij}$ if and only if $a_{ij} \neq a_{ji}$. We deduce that

$$||A[I^{c}, I] - B[I^{c}, I]||_{0} \le ||A - A^{T}||_{0} \le \rho n^{2}.$$
(7.3)

Considering separately the subsets of entries indexed by $I \times [n]$, $I^c \times I$ and $I^c \times I^c$, using (7.1) to (7.3), we conclude

$$\|A - B\|_{0} \le \|A[I, [n]] - B[I, [n]]\|_{0} + \|A[I^{c}, I] - B[I^{c}, I]\|_{0} + |I^{c}|^{2} \le (r^{2} + 1)\rho n^{2}.$$

Before completing the proof of Lemma 7.2, we record a necessary inequality on matrix ranks.

Lemma 7.5. Consider a matrix $A \in \mathbb{F}^{n \times n}$ over any field \mathbb{F} (for any integer $n \geq 1$). Then, for any $I, J \subseteq [n]$, we have

$$\operatorname{rank}(A) \ge \operatorname{rank}(A[I, [n]]) + \operatorname{rank}(A[[n], J]) - \operatorname{rank}(A[I, J]).$$

Proof. Write $\Delta := \operatorname{rank}(A[[n], J]) - \operatorname{rank}(A[I, J]) \ge 0$. There exist distinct $i_1, i_2, \ldots, i_\Delta \in [n] \setminus I$ such that for each $k \in \{1, \ldots, \Delta\}$, the row vector $A[i_k, J]$ is not spanned by the row vectors of A[I, J] along with $A[i_1, J], \ldots, A[i_{k-1}, J]$. Hence, for each $k \in \{1, \ldots, \Delta\}$, the row vector $A[i_k, [n]]$ is not spanned by the row vectors of A[I, [n]] along with $A[i_1, [n]], \ldots, A[i_{k-1}, [n]]$. In other words, rank $(A[I \cup \{i_1, \ldots, i_{\Delta}\}, [n]]) \ge 1$ $\operatorname{rank}(A[I, [n]]) + \Delta$. This shows $\operatorname{rank}(A) \ge \operatorname{rank}(A[I, [n]]) + \Delta$, as desired.

We now complete the proof of Lemma 7.2.

Proof of Lemma 7.2. We first apply Lemma 7.3 with $\gamma = q\sqrt{\rho}$ to obtain $k \in \{0, \ldots, q\}$ and $I \subseteq [n]$ of size $|I| \ge (1 - q^2 \sqrt{\rho})n$ such that rank $(A[I', I']) = \operatorname{rank}(A[I, I]) = k$ for all $I' \subseteq I$ of size $|I'| \ge |I| - q\sqrt{\rho}n$. Write rank $(A[[n], I]) = k + \Delta_1$ and rank $(A[I, [n]]) = k + \Delta_2$; clearly, $\Delta_1, \Delta_2 \ge 0$. According to Lemma 7.5, $q \geq \operatorname{rank}(A) \geq k + \Delta_1 + k + \Delta_2 - k = k + \Delta_1 + \Delta_2$. Without loss of generality, we assume $\Delta_1 \leq \Delta_2$ (otherwise replace A with A^T). Hence, $k + 2\Delta_1 \leq k + \Delta_1 + \Delta_2 \leq q$.

Now, Lemma 7.4 applied to A[I, I] (with r = k) yields a symmetric matrix $C \in \mathbb{F}^{I \times I}$ such that

$$||A[I,I] - C||_0 \le (q^2 + 1)\rho n^2 \tag{7.4}$$

and such that the row space of C is the same as that of A[I, I]. The latter property implies that $A[I, I] = P_1 C$ for some $P_1 \in \mathbb{F}^{I \times I}$.

We have shown how to approximate A[I, I] with the symmetric matrix C; we now need to deal with the other parts of the matrix A. Recall that $\operatorname{rank}(A[[n], I]) = \operatorname{rank}(A[I, I]) + \Delta_1$. Letting $I^c = [n] \setminus I$, this means we can write $A[I^c, I] = P_2 A[I, I] + Q$, for some $P_2, Q \in \mathbb{F}^{I^c \times I}$ and $\operatorname{rank}(Q) = \Delta_1$. Let $P := P_2 P_1 \in \mathbb{F}^{I^c \times I}, \text{ so } A[I^c, I] = PC + Q.$

Now, let $B \in \mathbb{F}^{n \times n}$ be the symmetric matrix defined by

$$B[I, I] = C, \quad B[I^{c}, I] = B[I, I^{c}]^{T} = PC + Q, \quad B[I^{c}, I^{c}] = PCP^{T} + PQ^{T} + QP^{T}$$

By construction, and our assumption on A, we have

$$||B[I^{c}, I] - A[I^{c}, I]||_{0} = 0,$$
(7.5)

$$||B[I, I^{c}] - A[I, I^{c}]||_{0} = ||B[I^{c}, I]^{T} - A[I, I^{c}]||_{0} = ||A[I^{c}, I]^{T} - A[I, I^{c}]||_{0} \le \rho n^{2}.$$
(7.6)

Combining (7.5) and (7.6) with (7.4) yields

$$\begin{split} \|A - B\|_{0} &\leq \|C - A[I, I]\|_{0} + \|B[I^{c}, I] - A[I^{c}, I]\|_{0} + \|B[I, I^{c}] - A[I, I^{c}]\|_{0} + |I^{c}|^{2} \\ &\leq (q^{2} + 1)\rho n^{2} + \rho n^{2} + q^{4}\rho n^{2} = O(q^{4}\rho n^{2}). \end{split}$$

It remains to show that $rank(B) \leq q$. This follows from a short sequence of elementary row and column operations:

$$\operatorname{rank}(B) = \operatorname{rank}\left[\frac{C}{PC+Q} \left| \frac{CP^{T}+Q^{T}}{PC+Q} \right| = \operatorname{rank}\left[\frac{C}{PC+Q} \left| \frac{Q^{T}}{PQ} \right| = \operatorname{rank}\left[\frac{C}{Q} \left| \frac{Q^{T}}{Q} \right| \right] = \operatorname{rank}\left[\frac{C}{Q} \left| \frac{Q^{T}}{Q} \right| \right] \\ \leq \operatorname{rank}(C) + 2\operatorname{rank}(Q) = k + 2\Delta_{1} \leq q.$$

8. Decoupling inequalities

In this section we prove Lemmas 2.12 and 2.16, which are the decoupling inequalities that will be used in the proofs of Theorems 1.10 and 1.11 respectively. Both lemmas use the following fact, which is an easy application of Jensen's inequality.

Lemma 8.1 (Decoupling with multiple copies). Let $\mathcal{E}(E, F)$ be an event depending on independent random objects E, F. Then, for any $k \in \mathbb{N}$, we have

$$\mathbb{P}[\mathcal{E}(E,F)] \leq \mathbb{P}\big[\mathcal{E}(E_0,F) \cap \cdots \cap \mathcal{E}(E_k,F)\big]^{1/(k+1)},$$

where E_0, \ldots, E_k are independent copies of E.

Proof. By the law of total probability and the independence of E_0, \ldots, E_k , we have

$$\mathbb{P}\big[\mathcal{E}(E_0,F)\cap\cdots\cap\mathcal{E}(E_k,F)\big] = \mathbb{E}\big[\mathbb{P}\big[\mathcal{E}(E_0,F)\cap\cdots\cap\mathcal{E}(E_k,F)\,\big|F\,\big]\big] = \mathbb{E}\big[\mathbb{P}\big[\mathcal{E}(E,F)\,\big|F\,\big]^{k+1}\big].$$

By Jensen's inequality, this is at least

$$\mathbb{E}\Big[\mathbb{P}\big[\mathcal{E}(E,F)\,\big|\,F\,\big]\Big]^{k+1} = \mathbb{P}\big[\mathcal{E}(E,F)\big]^{k+1}$$

Rearranging yields the desired result.

We now prove Lemma 2.12 and Lemma 2.16. The proof approach is the same for both of these statements, but Lemma 2.16 involves somewhat more complicated manipulations of equations.

Proof of Lemma 2.12. Fix any $z \in \mathbb{F}$, and let $\vec{\xi}^{(0)}, \ldots, \vec{\xi}^{(k)}$ be independent copies of $\vec{\xi}$. Applying Lemma 8.1 with $E = \vec{\xi}[X]$ and $F = \vec{\xi}[Y]$, and with $\mathcal{E}(E,F)$ being the event that $f(\vec{\xi}[X], \vec{\xi}[Y]) = z$, we see

$$\mathbb{P}[f(\vec{\xi}) = z] \le \mathbb{P}[f(\vec{\xi}^{(i)}[X], \vec{\xi}[Y]) = z \text{ for all } i \in \{0, \dots, k\}]^{1/(k+1)}$$

Subtracting the equation for i = 0 from each of the other k equations, and using that $f(\vec{x}) = \vec{x}^T A \vec{x} + \vec{b}^T \vec{x} + c$ (for some $\vec{b} \in \mathbb{F}^n$ and $c \in \mathbb{F}$), we deduce

$$\begin{split} \sup_{z \in \mathbb{F}} \mathbb{P}[f(\vec{\xi}) &= z] \\ &\leq \mathbb{P}\Big[f(\vec{\xi}^{(i)}[X], \vec{\xi}[Y]) - f(\vec{\xi}^{(0)}[X], \vec{\xi}[Y]) = 0 \text{ for all } i \in \{1, \dots, k\}\Big]^{1/(k+1)} \\ &= \mathbb{P}\Big[(\vec{\xi}^{(i)}[X] - \vec{\xi}^{(0)}[X])^T A[X, Y] \vec{\xi}[Y] + \psi_0(\vec{\xi}^{(0)}[X], \vec{\xi}^{(i)}[X]) = 0 \text{ for all } i \in \{1, \dots, k\}\Big]^{1/(k+1)}, \end{split}$$

where

$$\begin{split} \psi_0(\vec{\xi}^{(0)}[X], \vec{\xi}^{(i)}[X]) &= \frac{1}{2} \Big(\vec{\xi}^{(i)}[X]^T A[X, X] \vec{\xi}^{(i)}[X] - \vec{\xi}^{(0)}[X]^T A[X, X] \vec{\xi}^{(0)}[X] + \vec{b}[X]^T (\vec{\xi}^{(i)}[X] - \vec{\xi}^{(0)}[X]) \Big). \\ \text{Let } \vec{x} \in \{-1, 1\}^X \text{ maximise the expression} \end{split}$$

$$\mathbb{P}\Big[(\vec{\xi}^{(i)}[X] - \vec{x})^T A[X, Y] \vec{\xi}[Y] + \psi_0(\vec{x}, \vec{\xi}^{(i)}[X]) = 0 \text{ for all } i \in \{1, \dots, k\}\Big].$$

Then, we conclude

$$\sup_{z \in \mathbb{F}} \mathbb{P}[f(\vec{\xi}) = z] \le \mathbb{P}\Big[(\vec{\alpha}^{(i)})^T A[X, Y] \, \vec{\xi}[Y] = \psi(\vec{\alpha}^{(i)}) \text{ for all } i \in \{1, \dots, k\} \Big]^{1/(k+1)}$$

where $\vec{\alpha}^{(i)} = \vec{\xi}^{(i)}[X] - \vec{x}$ is a shifted Rademacher random variable and $\psi(\vec{\alpha}^{(i)}) = -\psi_0(\vec{x}, \vec{\alpha}^{(i)} + \vec{x})$. \Box

Proof of Lemma 2.16. Fix any $z \in \mathbb{F}$ and let $\vec{\xi'}$ be an independent copy of $\vec{\xi}$. We apply the k = 1 case of Lemma 8.1 twice, first with

$$E = \vec{\xi}[X], \quad F = \vec{\xi}[Y \cup Z], \quad \mathcal{E}(E, F) = \{f(\vec{\xi}[X], \vec{\xi}[Y], \vec{\xi}[Z]) = z\},$$

and then with

$$E = \vec{\xi}[Y], \quad F = (\vec{\xi}[X], \vec{\xi}'[X], \vec{\xi}[Z]), \quad \mathcal{E}(E, F) = \begin{cases} f(\vec{\xi}[X], \vec{\xi}[Y], \vec{\xi}[Z]) = z, \\ f(\vec{\xi}'[X], \vec{\xi}[Y], \vec{\xi}[Z]) = z \end{cases}$$

This yields

$$\mathbb{P}[f(\vec{\xi}) = z] \le \mathbb{P} \begin{bmatrix} f(\vec{\xi}[X], \vec{\xi}[Y], \vec{\xi}[Z]) = z, & (8.1) \\ f(\vec{\xi}'[X], \vec{\xi}[Y], \vec{\xi}[Z]) = z, & (8.2) \\ f(\vec{\xi}[X], \vec{\xi}'[Y], \vec{\xi}[Z]) = z, & (8.3) \\ f(\vec{\xi}'[X], \vec{\xi}'[Y], \vec{\xi}[Z]) = z & (8.4) \end{bmatrix}^{1/4}$$

Now, we write $f(\vec{x}) = \vec{x}^T A \vec{x} + \vec{b}^T \vec{x} + c$ for some $\vec{b} \in \mathbb{F}^n$ and $c \in \mathbb{F}$. If we subtract (8.2) from (8.1), and (8.4) from (8.3), and (8.3) from (8.1), we obtain

$$2(\vec{\xi}[X] - \vec{\xi}'[X])^T A[X, Y \cup Z] \vec{\xi}[Y \cup Z] + \varphi_0(\vec{\xi}[X], \vec{\xi}'[X]) = 0$$
(8.5)

$$2(\vec{\xi}[X] - \vec{\xi}'[X])^T A[X, Y \cup Z](\vec{\xi}'[Y], \vec{\xi}[Z]) + \varphi_0(\vec{\xi}[X], \vec{\xi}'[X]) = 0,$$
(8.6)

$$2(\vec{\xi}[Y] - \vec{\xi}'[Y])^T A[Y, X \cup Z]\vec{\xi}[X \cup Z] + \psi_0(\vec{\xi}[Y], \vec{\xi}'[Y]) = 0.$$
(8.7)

where

$$\begin{aligned} \varphi_0(\vec{\xi}[X], \vec{\xi}'[X]) &= \vec{\xi}[X]^T A[X, X] \vec{\xi}[X] - \vec{\xi}'[X] A[X, X] \vec{\xi}'[X] + \vec{b}[X]^T (\vec{\xi}[X] - \vec{\xi}'[X]), \\ \psi_0(\vec{\xi}[Y], \vec{\xi}'[Y]) &= \vec{\xi}[Y]^T A[Y, Y] \vec{\xi}[Y] - \vec{\xi}'[Y] A[Y, Y] \vec{\xi}'[Y] + \vec{b}[Y]^T (\vec{\xi}[Y] - \vec{\xi}'[Y]). \end{aligned}$$

Now, let

$$\vec{\alpha} = \vec{\xi}[X] - \vec{\xi}'[X], \qquad \vec{\beta} = \vec{\xi}[Y] - \vec{\xi}'[Y], \qquad \vec{\gamma} = \vec{\xi}[Z]$$
) from (8.5) yields

Subtracting (8.6) from (8.5) yields

$$\vec{\alpha}^{\,T} A[X,Y]\vec{\beta} = 0.$$

(8.6) and (8.7) can be written as

$$\vec{\alpha}^{T}A[X,Z]\vec{\gamma} = \varphi_{1}(\vec{\xi}[X],\vec{\xi}'[X],\vec{\xi}[Y],\vec{\xi}'[Y]) \quad \text{and} \quad \vec{\beta}^{T}A[Y,Z]\vec{\gamma} = \psi_{1}(\vec{\xi}[X],\vec{\xi}'[X],\vec{\xi}[Y],\vec{\xi}'[Y])$$
respectively, where

$$\varphi_1(\vec{\xi}[X], \vec{\xi}'[X], \vec{\xi}[Y], \vec{\xi}'[Y]) = -(\vec{\xi}[X] - \vec{\xi}'[X])^T A[X, Y] \vec{\xi}'[Y] - \varphi_0(\vec{\xi}[X], \vec{\xi}'[X])/2$$

$$\psi_1(\vec{\xi}[X], \vec{\xi}'[X], \vec{\xi}[Y], \vec{\xi}'[Y]) = -(\vec{\xi}[Y] - \vec{\xi}'[Y])^T A[Y, X] \vec{\xi}[X] - \psi_0(\vec{\xi}[Y], \vec{\xi}'[Y])/2.$$

Hence,

$$\mathbb{P}[f(\vec{\xi}) = z] \le \mathbb{P} \begin{bmatrix} \vec{\alpha}^T A[X, Y] \vec{\beta} = 0, \\ \vec{\alpha}^T A[X, Z] \vec{\gamma} = \varphi_1(\vec{\xi}[X], \vec{\xi}'[X], \vec{\xi}[Y], \vec{\xi}'[Y]), & (8.8) \\ \vec{\beta}^T A[Y, Z] \vec{\gamma} = \psi_1(\vec{\xi}[X], \vec{\xi}'[X], \vec{\xi}[Y], \vec{\xi}'[Y]) & (8.9) \end{bmatrix}^{1/4}$$

This is almost the desired conclusion, except that φ_1 and ψ_1 do not only depend on $\vec{\alpha}, \vec{\beta}$. To address this issue, for all possible outcomes of $\vec{\alpha} \in \{-1, 0, 1\}^X$ and $\vec{\beta} \in \{-1, 0, 1\}^Y$, choose $\vec{x}_{\vec{\alpha}, \vec{\beta}} \in \{-1, 1\}^X$ and $\vec{y}_{\vec{\alpha}, \vec{\beta}} \in \{-1, 1\}^Y$ to maximise the conditional probability

$$\mathbb{P}\Big[(8.8) \text{ and } (8.9) \text{ hold } \Big| \vec{\xi}'[X] = \vec{x}_{\vec{\alpha},\vec{\beta}}, \ \vec{\xi}'[Y] = \vec{y}_{\vec{\alpha},\vec{\beta}}, \ \vec{\alpha},\vec{\beta}\Big]$$

(among all choices of $\vec{x}_{\vec{\alpha},\vec{\beta}}$ and $\vec{y}_{\vec{\alpha},\vec{\beta}}$ such that this conditional probability is well-defined). In addition, let $\varphi(\vec{\alpha},\vec{\beta}) = \varphi_1(\vec{x}_{\vec{\alpha},\vec{\beta}} + \vec{\alpha}, \vec{x}_{\vec{\alpha},\vec{\beta}}, \vec{y}_{\vec{\alpha},\vec{\beta}} + \vec{\beta}, \vec{y}_{\vec{\alpha},\vec{\beta}})$ and $\psi(\vec{\alpha},\vec{\beta}) = \psi_1(\vec{x}_{\vec{\alpha},\vec{\beta}} + \vec{\alpha}, \vec{x}_{\vec{\alpha},\vec{\beta}}, \vec{y}_{\vec{\alpha},\vec{\beta}} + \vec{\beta}, \vec{y}_{\vec{\alpha},\vec{\beta}})$. Then we have

$$\sup_{z \in \mathbb{F}} \mathbb{P}[f(\vec{\xi}) = z] \le \mathbb{P}\Big[\vec{\alpha}^T A[X, Y]\vec{\beta} = 0, \quad \vec{\alpha}^T A[X, Z]\vec{\gamma} = \varphi(\vec{\alpha}, \vec{\beta}), \quad \vec{\beta}^T A[Y, Z]\vec{\gamma} = \psi(\vec{\alpha}, \vec{\beta})\Big]^{1/4},$$

as desired.

9. Bounds in terms of rank, for quadratic polynomials

In this section we prove Lemma 2.13, and use it (together with Lemmas 2.7 and 2.12) to deduce Theorem 1.11. First, we show how to deduce Theorem 1.11 from Lemma 2.13.

Proof of Theorem 1.11 assuming Lemma 2.13. In what follows, we assume n is even¹³ and sufficiently large. Write the quadratic part of $f(\vec{x})$ as $\vec{x}^T A \vec{x}$.

Suppose C1 does not hold, i.e., for any symmetric matrix $B \in \mathbb{F}^{n \times n}$ with rank less than $2k^2$, $f(\vec{x}) - \vec{x}^T B \vec{x}$ must have more than εn^2 nonzero coefficients. Our goal is to show that C2 holds, i.e., $\sup_{z \in \mathbb{F}} \mathbb{P}[f(\xi_1, \ldots, \xi_n) = z] \leq n^{-1+2/k}$.

Since f has at most n+1 non-quadratic terms, our assumption that **C1** does not hold implies that for any symmetric matrix $B \in \mathbb{F}^{n \times n}$ with rank less than $2k^2$, it must be that $||A - B||_0 > \varepsilon n^2 - n - 1$. By Lemma 2.7, this means that, for some $\delta \in (0, 1)$ depending on ε and k, at least a δ -fraction of $2k^2 \times 2k^2$ submatrices of A are nonsingular.

By Lemma 3.7, there is a partition of $\{1, \ldots, n\} = X \cup Y$ into parts of size |X| = |Y| = n/2, such that more than a $(\delta/2)$ -fraction of $2k^2 \times 2k^2$ submatrices of A[X, Y] are nonsingular. By Lemma 2.12 we then have

$$\sup_{z \in \mathbb{F}} \mathbb{P}[f(\vec{\xi}) = z] \le \mathbb{P}\Big[(\vec{\alpha}^{(i)})^T A[X, Y] \, \vec{\xi}[Y] = \psi(\vec{\alpha}^{(i)}) \text{ for all } i \in \{1, \dots, k\} \Big]^{1/(k+1)}$$

for some function $\psi : \mathbb{F}^X \to \mathbb{F}$, and some i.i.d. random vectors $\vec{\alpha}^{(1)}, \ldots, \vec{\alpha}^{(k)} \in \mathbb{F}^X$ with independent shifted Rademacher entries (all independent of $\vec{\xi}[Y]$). By Lemma 2.13 (with Ξ as the matrix with rows $\vec{\alpha}^{(1)}, \ldots, \vec{\alpha}^{(k)}$, and $\vec{\eta} = \vec{\xi}[Y]$, and $\vec{\varphi}(\Xi) = (\psi(\vec{\alpha}^{(i)}))_{i=1}^k$, and $\varepsilon = 1$), we have

$$\mathbb{P}\Big[(\vec{\alpha}^{(i)})^T A[X,Y]\vec{\xi}[Y] = \psi(\vec{\alpha}^{(i)}) \text{ for all } i \in \{1,\dots,k\}\Big] \le (n/2)^{-k+1} \le n^{-(k+1)(1-2/k)}.$$

The desired bound follows.

In the rest of this section, we prove Lemma 2.13. Recalling the strategy outlined in Section 2.2.1, via the Nguyen–Vu inverse theorem (Theorem 2.10), the proof boils down to bounding the probability that a given tuple of columns of ΞA lies in a GAP of "small" rank and of "small" volume (recall Definition 2.9). To be precise, we need the following two lemmas.

Lemma 9.1. Fix $k \leq m$ and $\mathbb{F} \in \{\mathbb{R}, \mathbb{C}\}$. For some $\ell \leq n$, consider a matrix $A \in \mathbb{F}^{n \times m}$ which has at least ℓ disjoint nonsingular $m \times m$ submatrices, and let $\Xi \in \mathbb{F}^{k \times n}$ be a matrix with independent shifted Rademacher entries. Then,

$$\mathbb{P}[\operatorname{rank}(\Xi A) < k] = O_m(\ell^{-(m-k+1)/2}).$$

Proof. This follows from Corollary 3.3. Indeed, write $\vec{\xi}^{(i)} \in \mathbb{F}^n$ for the *i*-th row of Ξ . If rank $(\Xi A) < k$, then there is an index $i \in \{1, \ldots, k\}$ such that $\vec{\xi}^{(i)}A$ lies in span $(\vec{\xi}^{(j)}A : j < i)$ (which is a linear subspace with dimension at most k - 1). Thus, applying Corollary 3.3 (after transposing, with the matrix A^T), we conclude

$$\mathbb{P}[\operatorname{rank}(\Xi A) < k] \le \sum_{i=1}^{k} \mathbb{P}[\vec{\xi}^{(i)}A \in \operatorname{span}(\vec{\xi}^{(j)}A : j < i)] = k \cdot O_m(\ell^{-(m-k+1)/2}) = O_m(\ell^{-(m-k+1)/2}). \quad \Box$$

Lemma 9.2. Fix $k \leq r \leq m$ and $\mathbb{F} \in \{\mathbb{R}, \mathbb{C}\}$. For some $\ell \leq n$, consider a matrix $A \in \mathbb{F}^{n \times m}$ which has at least ℓ disjoint nonsingular $m \times m$ submatrices, and let $\Xi \in \mathbb{F}^{k \times n}$ be a matrix with independent shifted Rademacher (or lazy Rademacher) entries. Then for any $V \geq 1$,

 $\mathbb{P}[\text{the columns of } \Xi A \text{ lie in a symmetric } GAP \text{ of rank at most } r \text{ and volume at most } V]$

$$\leq O_m (V^m (1 + \log V)^{r-1} \cdot \ell^{-k(m-r)/2}).$$

Proof. Let $\mathcal{Z}(N_1, \ldots, N_r)$ be the set of integer matrices $Z \in \mathbb{Z}^{r \times m}$ such that, for each $i \in \{1, \ldots, r\}$, the absolute value of the entries in the *i*th row is at most N_i . Then, let $\mathcal{Z}(V)$ be the union of all the sets $\mathcal{Z}(N_1, \ldots, N_r)$, among N_1, \ldots, N_r satisfying $(2N_i + 1)(2N_2 + 1) \ldots (2N_r + 1) \leq V$.

By definition, if the column vectors of ΞA lie in a single symmetric GAP of rank at most r and volume at most V, then $\Xi A = UZ$ for some $U \in \mathbb{F}^{k \times r}$ and $Z \in \mathcal{Z}(V)$. We will proceed by studying the probability that $\Xi A = UZ$ for some particular Z, and summing over all $Z \in \mathcal{Z}(V)$.

Claim 9.3. For every fixed $Z \in \mathcal{Z}(V)$, we have $\mathbb{P}[\Xi A = UZ \text{ for some } U \in \mathbb{F}^{k \times r}] \leq \ell^{-k(m-r)/2}$.

 $^{^{13}}$ For the odd-*n* case, we can add a single dummy variable to reduce to the even-*n* case.

Proof of claim. This is a consequence of Corollary 3.3, noting that we can interpret ΞA as a (matrix-valued) linear function of the kn entries of Ξ , and we can interpret $\{UZ : U \in \mathbb{F}^{k \times r}\}$ as a kr-dimensional linear subspace of the vector space of $k \times m$ matrices.

In more detail: let $\vec{\xi}^* \in \mathbb{F}^{kn}$ be the vector obtained by concatenating the k rows of Ξ (so, this is a row vector of kn independent shifted Rademacher random variables or kn independent lazy Rademacher random variables), and let $A^* \in \mathbb{F}^{nk \times mk}$ be the block-diagonal matrix obtained by concatenating k copies of A along the diagonal. Then, $\vec{\xi}^*A^*$ is the concatenation of the k rows of ΞA . If we let $\mathcal{W} \in \mathbb{F}^{km}$ be the set of all vectors obtained by concatenating the k rows of a matrix UZ, for some $U \in \mathbb{F}^{k \times r}$, then the event that $\Xi A = UZ$ for some $U \in \mathbb{F}^{k \times r}$ precisely corresponds to the event that $\vec{\xi}^*A^* \in \mathcal{W}$. Note that A^* has at least ℓ disjoint nonsingular $m \times m$ matrices in A), and \mathcal{W} is a kr-dimensional linear subspace of \mathbb{F}^{km} , so the desired bound follows from Corollary 3.3 (after transposing, with the matrix $(A^*)^T$).

Claim 9.4.
$$|\mathcal{Z}(V)| = O_m(V^m(1 + \log V))^{r-1}).$$

Proof of claim. Let d be the minimal positive integer such that $V \leq 2^d$, so $d = \log_2 V + O(1)$.

Consider $N_1, \ldots, N_r \in \mathbb{N}$ such that $(2N_i + 1)(2N_2 + 1) \ldots (2N_r + 1) \leq V$. For each $i \in \{1, \ldots, r\}$, let $d_i \in \mathbb{N}$ be the minimal integer such that $N_i \leq 2^{d_i}$, so $2N_i + 1 \geq 2^{d_i}$. We have $2^{d_1 + \cdots + d_r} \leq V \leq 2^d$, i.e., $d_1 + \cdots + d_r \leq d$. It follows that

$$\mathcal{Z}(V) \subseteq \bigcup_{0 \le d_1 + \dots + d_r \le d} \mathcal{Z}(2^{d_1}, \dots, 2^{d_r}),$$

and hence

$$|\mathcal{Z}(V)| \leq \sum_{0 \leq d_1 + \dots + d_r \leq d} |\mathcal{Z}(2^{d_1}, \dots, 2^{d_r})|.$$

Then, note that $|\mathcal{Z}(2^{d_1}, \ldots, 2^{d_r})| = \prod_{i=1}^r (2^{d_i+1}+1)^m = O_m(2^{m(d_1+\cdots+d_r)})$. We deduce that $|\mathcal{Z}(V)|$ is at most (recall $r \leq m$)

$$\sum_{\substack{d_1,\dots,d_r\in\mathbb{N}\\d_1+\dots+d_r\leq d}} O_m(2^{m(d_1+\dots+d_r)}) \leq \sum_{D=0}^d \binom{D+r-1}{r-1} O_m(2^{mD}) = O_m(d^{r-1}2^{md}) = O_m(V^m(1+\log V)^{r-1}). \blacksquare$$

The desired bound immediately follows from the above two claims.

Remark 9.5. Both Lemma 9.1 and Lemma 9.2 seem far from the truth. In the setting of Lemma 9.1 we anticipate a bound of $\ell^{-m/2+o(1)}$ (see [29, 30] for related work in the analytic number theory literature, which can be viewed as special cases of this anticipated bound). If one could prove this, as well as an improved bound of the form $V^{m-r+o(1)} \cdot \ell^{-k(m-r)/2+o(1)}$ in Lemma 9.2, then it would be possible to deduce an optimal version of Lemma 2.13 (see Remark 2.14), using some of the ideas we will see shortly in the proof of Theorem 1.10 (namely, consideration of "expected inverse volume" quantities as in Lemma 10.6).

Equipped with Lemmas 9.1 and 9.2, we now present the proof of Lemma 2.13.

Proof of Lemma 2.13. Suppose **G2** does not hold, i.e. $\mathbb{P}[\Xi A \vec{\eta} = \vec{\varphi}(\Xi)] > n^{-k+\varepsilon}$ for some function $\varphi : \mathbb{F}^{k \times n} \to \mathbb{F}^k$ and some $\varepsilon \in (0, 1]$. Taking $m := \lceil 2k^2/\varepsilon \rceil$, the goal is to prove that **G1** holds, namely that at most a δ -fraction of the $m \times m$ submatrices of A are nonsingular (for any $\delta > 0$, assuming that n is sufficiently large in terms of m, δ).

Let \mathcal{E} be the event that Ξ satisfies $\mathbb{P}[\Xi A \vec{\eta} = \varphi(\Xi) | \Xi] > n^{-k+\varepsilon}/2$. We have

$$n^{-k+\varepsilon} < \mathbb{P}\big[\Xi A \vec{\eta} = \varphi(\Xi)\big] \le \mathbb{P}[\mathcal{E}] + \mathbb{P}\big[\Xi A \vec{\eta} = \varphi(\Xi)\big|\mathcal{E}^{c}\big] \le \mathbb{P}[\mathcal{E}] + n^{-k+\varepsilon}/2,$$

so we must have $\mathbb{P}[\mathcal{E}] > n^{-k+\varepsilon}/2$ (and so in particular $\mathbb{P}[\mathcal{E}] > n^{-k}$).

Now, note that $\Xi A \vec{\eta}$ can be viewed as a linear function of $\vec{\eta}$ (taking values in \mathbb{F}^k), where the coefficients of this linear function are the columns of ΞA . So, if Ξ satisfies \mathcal{E} , then Theorem 2.10 tells us that most of the columns of ΞA are contained in a symmetric GAP of small rank and volume. More precisely, for Ξ satisfying \mathcal{E} , applying Theorem 2.10 with $n' = (\delta/(4m))n$, we conclude there exists a symmetric GAP φ in \mathbb{F}^k with rank $r = O_k(1)$ and volume $V = O_{k,\delta,\varepsilon}(n^{-r/2+k-\varepsilon}) \leq n^{-r/2+k-7\varepsilon/8}$ such that all but at most n' of the columns of ΞA lie in φ . Observe that r < 2k, since otherwise V = o(1), which is impossible according to the definition of V. We define a suitable GAP to be a symmetric GAP in \mathbb{F}^k of rank r < 2k and volume $V \leq n^{-r/2+k-7\varepsilon/8}$. For an arbitrary outcome of Ξ , define $I = I(\Xi) \subseteq \{1, \ldots, n\}$ to be a maximum-size set of (indices of) columns of ΞA that are contained in a suitable GAP. Whenever Ξ satisfies \mathcal{E} , we have $|I| \geq n - n' \geq (1 - \delta/(4m))n$. By Lemma 3.1 (with m as the "r" in Lemma 3.1), it follows that for all but at most a $\delta/2$ -fraction of m-sets $J \subseteq \{1, \ldots, n\}$, we have $\mathbb{P}[J \subseteq I] \geq \mathbb{P}[\mathcal{E}]/2 \geq n^{-k+\varepsilon}/4$. Say such subsets J are good.

If J is good, then the corresponding columns of ΞA are quite likely to lie in a suitable GAP. We will use Lemmas 9.1 and 9.2 to show that this implies that A has many disjoint nonsingular $m \times m$ submatrices in the columns indexed by J; summing over good J will then yield the desired result.

Claim 9.6. If J is good, then at most a $\delta/2$ -fraction of the $m \times m$ submatrices of A in the columns indexed by J are nonsingular.

Proof of claim. Let \mathcal{F}_r be the event that the columns of ΞA indexed by J lie in a suitable GAP with rank r, so the assumption that J is good implies that $\mathbb{P}[\mathcal{F}_0 \cup \cdots \cup \mathcal{F}_{2k-1}] \ge n^{-k+\varepsilon}/4$.

Suppose for the purpose of contradiction that the columns indexed by J contain at least $(\delta/2)\binom{n}{m}$ nonsingular $m \times m$ submatrices. By Fact 3.4, then these columns in fact contain $\ell = \Omega_{m,\delta}(n)$ disjoint $m \times m$ submatrices.

Recall that $m = \lceil 2k^2/\varepsilon \rceil$. First, Lemma 9.1 implies that

$$\mathbb{P}[\mathcal{F}_0 \cup \cdots \cup \mathcal{F}_{k-1}] \le O_m(\ell^{-(m-k+1)/2}) = o_{m,\delta}(n^{-k+\varepsilon}),$$

with plenty of room to spare. Second, if $r \ge k$, Lemma 9.2 implies that, for $V = n^{k-7\varepsilon/8-r/2}$,

$$\mathbb{P}[\mathcal{F}_r] \le O_m(V^m(\log V)^{r-1} \cdot \ell^{-k(m-r)/2}) = O_{m,\delta}(n^{m(k-7\varepsilon/8-r/2)+o(1)} \cdot n^{-k(m-r)/2})$$
$$= O_{m,\delta}(n^{km/2-7\varepsilon m/8-r(m/2-k/2)+o(1)})$$
$$\le O_{m,\delta}(n^{km/2-7\varepsilon m/8-k(m/2-k/2)+o(1)})$$
$$\le O_{m,\delta}(n^{-7k^2/4+k^2/2+o(1)}) = o_{m,\delta}(n^{-k+\varepsilon}).$$

Summing the above bounds yields $\mathbb{P}[\mathcal{F}_0 \cup \cdots \cup \mathcal{F}_{2k-1}] = o_{m,\delta}(n^{-k+\varepsilon})$, which is the desired contradiction.

Now, using the above claim, the number of nonsingular $m \times m$ submatrices in good sets of columns is at most $(\delta/2) \binom{n}{m}^2$. Recalling that all but a $(\delta/2)$ -fraction of *m*-sets of columns are good, the number of nonsingular $m \times m$ submatrices which do *not* lie in a good set of columns is also at most $(\delta/2) \binom{n}{m}^2$. That is to say, at most a δ -fraction of $m \times m$ submatrices of *A* are nonsingular, which proves **G1**.

10. A power-saving improvement for robustly irreducible quadratic polynomials

In this section we prove Lemma 2.17, and use it (together with Lemmas 2.7 and 2.16) to deduce Theorem 1.10. First, we show how to deduce Theorem 1.10 from Lemma 2.17.

Proof of Theorem 1.10 assuming Lemma 2.16 and Lemma 2.17. In what follows, we assume n is divisible¹⁴ by 3, and sufficiently large. Write the quadratic part of $f(\vec{x})$ as $\vec{x}^T A \vec{x}$ for some symmetric matrix $A \in \mathbb{C}^{n \times n}$.

Suppose **B1** does not hold, i.e., for every reducible polynomial $g \in \mathbb{C}[x_1, \ldots, x_n]$ with degree at most 2, the polynomial f - g must have more than εn^2 nonzero coefficients. Our objective is to prove that **B2**' holds, i.e., $\sup_{z \in \mathbb{C}} \mathbb{P}[f(\vec{\xi}) = z] \leq n^{-1/2 - 1/24 + \varepsilon}$.

Note that a quadratic form is reducible over \mathbb{C} if and only if its rank is less than 3. Thus, for every symmetric matrix $B \in \mathbb{C}^{n \times n}$ of rank less than 3, the polynomial $g(\vec{x}) = \vec{x}^T B \vec{x}$ is reducible, and so f - g must have more than εn^2 nonzero coefficients. Since f - g has at most n + 1 non-quadratic terms, this implies that we must have $||A - B||_0 > \varepsilon n^2 - n - 1$ for every symmetric matrix $B \in \mathbb{C}^{n \times n}$ with rank less than 3. By Lemma 2.7, this means that, for some $\delta > 0$ depending on ε , at least a δ -fraction of the 3×3 submatrices of A are nonsingular.

By Lemma 3.7, there is a partition of $\{1, \ldots, n\} = X \cup Y \cup Z$ with |X| = |Y| = |Z| = n/3 such that for each of the matrices A[X, Y], A[X, Z], A[Y, Z], more than a $(\delta/2)$ -fraction of their 3×3 submatrices are nonsingular. By Lemma 2.16 we then have

$$\sup_{z\in\mathbb{C}}\mathbb{P}[f(\vec{\xi})=z] \leq \mathbb{P}\Big[\vec{\alpha}^{\,T}A[X,Y]\vec{\beta}=0, \quad \vec{\alpha}^{\,T}A[X,Z]\vec{\gamma}=\varphi(\vec{\alpha},\vec{\beta}), \quad \vec{\beta}^{\,T}A[Y,Z]\vec{\gamma}=\psi(\vec{\alpha},\vec{\beta})\Big]^{1/4},$$

¹⁴in general, we can add one or two dummy variables to reduce to the case where n is divisible by 3.

for some functions $\varphi, \psi : \mathbb{C}^X \times \mathbb{C}^Y \to \mathbb{C}$, and independent random vectors $\vec{\alpha}, \vec{\beta}, \vec{\gamma}$, where the entries of $\vec{\alpha}, \vec{\beta}$ are i.i.d. lazy Rademacher, and the entries of $\vec{\gamma}$ are i.i.d. Rademacher. By Lemma 2.17 (taking the " ε " in Lemma 2.17 to be min $(\delta/2, \varepsilon)$, we have

$$\mathbb{P}\Big[\vec{\alpha}^{T}A[X,Y]\vec{\beta} = 0, \quad \vec{\alpha}^{T}A[X,Z]\vec{\gamma} = \varphi(\vec{\alpha},\vec{\beta}), \quad \vec{\beta}^{T}A[Y,Z]\vec{\gamma} = \psi(\vec{\alpha},\vec{\beta})\Big] \leq (n/3)^{-2-1/6+\varepsilon} \leq n^{-4(1/2+1/24-\varepsilon)},$$

he desired bound follows.
$$\Box$$

and the desired bound follows.

In the rest of this section, we prove Lemma 2.17. As in the proof of Theorem 1.11, we will use the Nguyen–Vu inverse theorem (Theorem 2.10), but the estimates are much more delicate. It will be convenient to use the following reformulation of Theorem 2.10 (which we will apply with q = 3).

Definition 10.1. Let G be a torsion-free additive group. For $r \in \mathbb{N}$ and $\delta > 0$, and a sequence $\vec{v} = (v_1, \ldots, v_n) \in G^n$ of elements of G, let $h_r^{\delta}(\vec{v})$ be the minimum volume of a symmetric GAP of rank at most r that contains all but at most $n^{1-\delta}$ entries of \vec{v} (let $h_r^{\delta}(\vec{v}) = \infty$ if no such GAP exists).

We remark that the only possible rank-0 symmetric GAPs is the singleton $\{0\}$. So, $h_0^{\delta}(\vec{v}) = 1$ if at most $n^{1-\delta}$ elements of \vec{v} are nonzero; otherwise $h_0^{\delta}(\vec{v}) = \infty$.

Theorem 10.2. Fix $q \in \mathbb{N}$ and $0 < \delta < 1/2$, and let G be a torsion-free additive group. Let n be sufficiently large (in terms of δ, q) and consider some $\vec{v} = (v_1, \ldots, v_n) \in G^n$. Let ξ_1, \ldots, ξ_n be i.i.d. Rademacher (or lazy Rademacher) random variables. We have

$$\sup_{z \in G} \mathbb{P}[\xi_1 v_1 + \dots + \xi_n v_n = z] \le n^{-q/2 + q\delta} + \sum_{r=0}^{q-1} n^{-r/2 + q\delta} \cdot \frac{1}{h_r^{\delta}(\vec{v})}$$

Proof. Let $\rho = \sup_{z \in G} \mathbb{P}[\xi_1 v_1 + \dots + \xi_n v_n = z]$ and suppose for the purpose of contradiction that the desired bound does not hold. Then, we have $\rho \ge n^{-q/2+q\delta}$ and by¹⁵ Theorem 2.10 there is a symmetric GAP that contains all but $n^{1-\delta}$ of v_1, \ldots, v_n , such that the rank r and the volume V of this GAP satisfy

$$V = O_q(\rho^{-1}n^{-(1-\delta)(r/2)}).$$
(10.1)

We are assuming that $\rho > n^{-r/2+q\delta}/h_r^{\delta}(\vec{v})$ for all r < q, so (10.1) implies that the volume of our GAP satisfies $V = O_q(n^{-q\delta+\delta r/2}h_r^{\delta}(\vec{v})) < h_r^{\delta}(\vec{v})$ for r < q. By the definition of $h_r^{\delta}(\vec{v})$, it follows that we cannot have r < q. But if $r \ge q$, then (10.1) also implies that the volume of our GAP is $V = O_q(n^{-q\delta+q/2-(1-\delta)r/2}) = O_q(n^{-q\delta+q/2-(1-\delta)q/2}) = O_q(n^{-q\delta/2}) < 1$ (recalling that $\rho \ge n^{-q/2+q\delta}$), which is impossible.

10.1. Reduction to technical estimates. Recall that we are working towards a proof of Lemma 2.17; namely, we are attempting to upper bound the probability of the event that $\vec{\alpha}^T A_1 \vec{\beta} = 0$, and $\vec{\alpha}^T A_2 \vec{\gamma} =$ $\varphi(\vec{\alpha},\vec{\beta})$, and $\vec{\beta}^T A_3 \vec{\gamma} = \psi(\vec{\alpha},\vec{\beta})$. We will want to break this event into sub-events, depending on the structure of the vectors $\vec{\alpha}^T A_1$, $\vec{\alpha}^T A_2$ and $\vec{\beta}^T A_3$; in each case we will reveal our random vectors in slightly different orders, and estimate probabilities in slightly different ways.

For example, it is quite unlikely that almost all the entries of $\vec{\alpha}^T A_1$ are zero. But if $\vec{\alpha}^T A_1$ has plenty of nonzero entries, then we can get quite a strong bound on the probability that $\vec{\alpha}^T A_1 \vec{\beta} = 0$ using the randomness of $\vec{\beta}$. Similarly, it is quite unlikely that $\vec{\alpha}^T A_2$ and $\vec{\beta}^T A_3$ are nearly collinear. If they are not nearly collinear, then we can get quite a strong bound on the joint probability that $\vec{\alpha}^T A_2 \vec{\gamma} = \varphi(\vec{\alpha}, \vec{\beta})$, and $\vec{\beta}^T A_3 \vec{\gamma} = \psi(\vec{\alpha}, \vec{\beta})$, using the randomness of $\vec{\gamma}$.

In this subsection, we collect various technical probabilistic estimates of this kind (some of whose proofs will be deferred until the next subsection), and show how to piece together these estimates to prove Lemma 2.17. To state all these estimates, it will also be convenient to introduce some notation for the property that a matrix robustly has rank at least 3.

Definition 10.3. We say a matrix $A \in \mathbb{C}^{n \times n}$ is said to be ε -robust if more than an ε -fraction of its 3×3 submatrices are nonsingular.

 $^{^{15}}$ Theorem 2.10 is only stated for Rademacher random variables, but it is easy to deduce the same statement in the lazy Rademacher case. Indeed, a sequence of independent lazy Rademacher random variables can be expressed as a sum $\bar{\xi}^{(1)} + \bar{\xi}^{(2)}$ of two independent sequences of independent Rademacher random variables. We can simply apply Theorem 2.10 after conditioning on an outcome of $\vec{\xi}^{(1)}$.

First, the following lemma, which is a simple corollary of Halász' inequality (Theorem 3.2), will be used to bound the probability that almost all of the entries of $\vec{\alpha}^T A_1$ or $\vec{\alpha}^T A_2$ are zero.

Lemma 10.4. Fix $\delta > 0$ and let $n \in \mathbb{N}$ be sufficiently large (in terms of δ). Consider a δ -robust matrix $A \in \mathbb{C}^{n \times n}$. Let $\vec{\alpha} \in \{-1, 0, 1\}^n$ be a vector of *i.i.d.* lazy Rademacher random variables. Then

$$\mathbb{P}\left[\|\vec{\alpha}^T A\|_0 \le n^{1-\delta}\right] < n^{-3/2+\delta}.$$

In the proof of Lemma 10.4, and some other lemmas later in this section, it is convenient to interpret a $m \times n$ matrix $A \in \mathbb{F}^{m \times n}$ as a sequence of n column vectors in \mathbb{F}^m , so in particular, for a subset $J \subseteq \{1, \ldots, n\}$ we write A[J] to denote the $m \times |J|$ submatrix of A consisting of just the columns indexed by J.

Proof of Lemma 10.4. Suppose for the purpose of contradiction that $\mathbb{P}[\|\vec{\alpha}^T A\|_0 \leq n^{1-\delta}] \geq n^{-3/2+\delta}$. By Lemma 3.1, for all but a $6n^{-\delta}$ -fraction of the 3-element subsets J of $\{1, \ldots, n\}$, we have $\mathbb{P}[\vec{\alpha}^T A[J] = \vec{0}] \geq \frac{1}{2}n^{-3/2+\delta}$. Say such a J is good.

By Theorem 3.2 (after transposing), if J is good, then A[J] contains at most $O(n^{1-2\delta/3})$ disjoint nonsingular 3×3 submatrices. By Fact 3.4, it follows that at most a $O(n^{-2\delta/3})$ -fraction of the 3×3 submatrices of A[J] are nonsingular.

Summing over both good and non-good triples of columns, the total number of nonsingular 3×3 submatrices in A is at most $6n^{-\delta}\binom{n}{3}^2 + O(n^{-2\delta/3}) \cdot \binom{n}{3}^2 < \delta\binom{n}{3}^2$, contradicting δ -robustness of A. \Box

Next, given an outcome of $\vec{\alpha}^T A_2$ with many nonzero entries, the following lemma will be used to bound the probability that $\vec{\alpha}^T A_2$ and $\vec{\beta}^T A_3$ are nearly collinear, while $\vec{\alpha}^T A_1 \vec{\beta} = 0$ (when applying the lemma, we will take \vec{u} and \vec{v} to be $\vec{\alpha}^T A_1$ and $\vec{\alpha}^T A_2$, respectively). It is proved in a similar way to Lemma 10.4, using Theorem 10.2 instead of Halász' inequality. (Recall the notation h_r^{δ} from Definition 10.1).

Lemma 10.5. Fix $\varepsilon, \delta > 0$ and let $n \in \mathbb{N}$ be sufficiently large (in terms of ε, δ). Consider (row) vectors \vec{u}, \vec{v} with $\|\vec{v}\|_0 > n^{1-\delta}$, and consider an ε -robust matrix $A \in \mathbb{C}^{n \times n}$. Let $\vec{\beta} \in \{-1, 0, 1\}^n$ be a (column) vector of *i.i.d.* lazy Rademacher random variables. Then

$$\mathbb{P}\Big[\vec{u}\cdot\vec{\beta}=0, \text{ and } \|\vec{a}\vec{v}+b\vec{\beta}^TA\|_0 \le n^{1-\delta} \text{ for some } a, b \in \mathbb{C}, \text{ not both zero}\Big] \le 2n^{-1+3\delta} \cdot \frac{1}{h_2^{\delta}(\vec{u})} + 2n^{-3/2+3\delta}.$$

Proof. First note that if b = 0 and $a \neq 0$ then $||a\vec{v} + b\vec{\beta}^T A||_0 = ||\vec{v}||_0 > n^{1-\delta}$, so it suffices to consider $b \neq 0$.

Suppose for the purpose of contradiction that the desired inequality does not hold. Then Lemma 3.1 tells us that for all but a $6n^{-\delta}$ -fraction of the triples of columns J, we have

$$\mathbb{P}\left[\vec{u}\cdot\vec{\beta}=0, \text{ and } a\vec{v}[J]+b\vec{\beta}^{T}A[J]=\vec{0} \text{ for some } a\in\mathbb{C} \text{ and } b\in\mathbb{C}\setminus\{0\}\right] > n^{-1+3\delta}\frac{1}{h_{2}^{\delta}(\vec{u})} + n^{-3/2+3\delta}.$$

Say that such J are good.

For each good triple of columns J, let $P_J \in \mathbb{C}^{3\times 2}$ be a rank-2 matrix whose columns are both in the orthogonal complement of span $(\vec{v}[J])$, and let $((A[J]P_J)^T; \vec{u})$ be the $3 \times n$ matrix obtained by attaching the row \vec{u} to the $2 \times n$ matrix $(A[J]P_J)^T$. So, if $\vec{u} \cdot \vec{\beta} = 0$ and $a\vec{v}[J] + b\vec{\beta}^T A[J] = \vec{0}$ for some $a \in \mathbb{C}$ and $b \in \mathbb{C} \setminus \{0\}$, then we have $((A[J]P_J)^T; \vec{u})\vec{\beta} = \vec{0}$.

Associating $((A[J]P_J)^T; \vec{u})$ with its sequence of *n* column vectors, we have $h_2^{\delta}((A[J]P_J)^T; \vec{u}) \ge h_2^{\delta}(\vec{u})$, so we deduce

$$\mathbb{P}[((A[J]P_J)^T; \vec{u})\vec{\beta} = \vec{0}] > n^{-1+3\delta} \frac{1}{h_2^{\delta}((A[J]P_J)^T; \vec{u})} + n^{-3/2+3\delta}$$

By Theorem 10.2, it follows that $h_r^{\delta}((A[J]P_J)^T; \vec{u}) < \infty$ for some $r \leq 1$. This means that all but at most $n^{1-\delta}$ of the rows of the $n \times 2$ matrix $A[J]P_J$ lie in a single symmetric GAP of rank at most 1. So, every 3×2 submatrix of $A[J]P_J$ which has rank 2 must involve at least one of the rows outside this GAP, meaning that A[J] has at most $n^{1-\delta} \cdot \binom{n}{2}$ nonsingular 3×3 submatrices.

Summing over both good and non-good triples of columns, the total number of nonsingular 3×3 submatrices in A is at most $6n^{-\delta}\binom{n}{3}^2 + n^{1-\delta}\binom{n}{2} \cdot \binom{n}{3} < \varepsilon\binom{n}{3}^2$, contradicting the ε -robustness of A. \Box

We also need some estimates on the quantities h_r^{δ} from Definition 10.1.

Lemma 10.6. Fix any $\delta > 0$, and let n be sufficiently large (in terms of δ). Let $\vec{\alpha} \in \mathbb{C}^n$ be a vector of independent lazy Rademacher random variables.

(1) For a δ -robust matrix $A \in \mathbb{C}^{n \times n}$, we have

$$\mathbb{E}\left[\frac{1}{h_2^{\delta}(\vec{\alpha}^T A)}\right] \le n^{-1/6+\delta}.$$

(2) For δ -robust matrices $A_1, A_2 \in \mathbb{C}^{n \times n}$, we have

$$\mathbb{E}\left[\frac{1}{h_1^{\delta}(\vec{\alpha}^T A_1)h_2^{\delta}(\vec{\alpha}^T A_2)}\right] \le n^{-2/3+\delta}.$$

Remark 10.7. The bounds in Lemma 10.6 are likely far from being optimal; we made no attempt to optimise them.

We defer the proof of Lemma 10.6 to the next subsection. First, we show how to deduce Lemma 2.17. *Proof of Lemma 2.17.* Let

$$\mathcal{E} := \left\{ \vec{\alpha}^T A_1 \vec{\beta} = 0 \right\} \cap \left\{ \vec{\alpha}^T A_2 \vec{\gamma} = \varphi(\vec{\alpha}, \vec{\beta}) \right\} \cap \left\{ \vec{\beta}^T A_3 \vec{\gamma} = \psi(\vec{\alpha}, \vec{\beta}) \right\}.$$

We will prove that $\mathbb{P}[\mathcal{E}] \leq n^{-2-1/6+\varepsilon}$ (i.e., **H2**), assuming that A_1, A_2, A_3 are ε -robust (i.e., assuming that **H1** does not hold). Throughout this proof, we assume *n* is sufficiently large whenever necessary. Let $\delta = \varepsilon/100$, and define two auxiliary events:

• Let \mathcal{E}_0 be the event that $\|\vec{\alpha}^T A_1\|_0 \leq n^{1-\delta}$ or $\|\vec{\alpha}^T A_2\|_0 \leq n^{1-\delta}$. Hence, \mathcal{E}_0 fails to hold if and only if $h_0^{\delta}(\vec{\alpha}^T A_1) = h_0^{\delta}(\vec{\alpha}^T A_2) = \infty$.

• Let \mathcal{E}_1 be the event that there are $a, b \in \mathbb{C}$, not both zero, such that $\|a\vec{\alpha}^T A_2 + b\vec{\beta}^T A_3\|_0 \leq n^{1-\delta}$. Note that if the event \mathcal{E}_1 does not hold, then there is no symmetric GAP of rank at most 1 containing all but $n^{1-\delta}$ of the columns in the $2 \times n$ matrix whose rows are $\vec{\alpha}^T A_2$ and $\vec{\beta}^T A_3$.

Now, $\mathbb{P}[\mathcal{E}]$ is at most

$$\mathbb{P}[\mathcal{E} \cap \mathcal{E}_0] + \mathbb{P}[\mathcal{E} \cap \mathcal{E}_0^c \cap \mathcal{E}_1] + \mathbb{P}[\mathcal{E} \cap \mathcal{E}_0^c \cap \mathcal{E}_1^c].$$
(10.2)

In the rest of the proof, we will show that each of the three terms above is at most $n^{-2-1/6+\varepsilon}/3$. Throughout the proof, we will put subscripts on probabilistic notation to remind the reader whether we are using the randomness of $\vec{\alpha}$, $\vec{\beta}$ or $\vec{\gamma}$. Also, we simply write h_0, h_1, h_2 instead of $h_0^{\delta}, h_1^{\delta}, h_2^{\delta}$.

The first term. By Lemma 10.4, we have $\mathbb{P}_{\vec{\alpha}}[\mathcal{E}_0] \leq 2n^{-3/2+\delta}$. By Fact 3.6, A_3 differs in more than $(\varepsilon/9)n^2$ entries from every matrix of rank at most 2 (in particular, this holds for every matrix of rank at most 1). So by Theorem 6.3, for any outcome of $\vec{\alpha}$, we have

$$\mathbb{P}_{\vec{\beta},\vec{\gamma}}[\mathcal{E} \mid \vec{\alpha}] \leq \mathbb{P}_{\vec{\beta},\vec{\gamma}}\left[\vec{\beta}^{\,T}A_{3}\vec{\gamma} = \psi(\vec{\alpha},\vec{\beta}) \mid \vec{\alpha}\right] \leq n^{-1+\varepsilon/9}$$

We deduce

$$\mathbb{P}[\mathcal{E} \cap \mathcal{E}_0] \le 2n^{-3/2+\delta} \cdot n^{-1+\varepsilon/9} < n^{-2-1/6+\varepsilon}/3$$

(with plenty of room to spare in the final inequality).

The second term. For any outcome of $\vec{\alpha}$ for which \mathcal{E}_0 does not hold (in particular, $\|\vec{\alpha}^T A_2\|_0 > n^{1-\delta}$), Lemma 10.5 tells us that

$$\mathbb{P}_{\vec{\beta}}[\vec{\alpha}^{T}A_{1}\vec{\beta}=0, \text{ and } \mathcal{E}_{1} \text{ occurs } | \vec{\alpha}] \leq 2 \cdot n^{-1+3\delta} \cdot \frac{1}{h_{2}(\vec{\alpha}^{T}A_{1})} + 2 \cdot n^{-3/2+3\delta} \leq n^{-1+4\delta} \cdot \frac{1}{h_{2}(\vec{\alpha}^{T}A_{1})} + n^{-3/2+4\delta} \cdot \frac{1$$

Then, for any outcomes of $\vec{\alpha}, \vec{\beta}$ for which \mathcal{E}_0 does not hold (in particular, $\|\vec{\alpha}^T A_2\|_0 > n^{1-\delta}$), we can apply Theorem 10.2 (with q = 2, using that $h_0(\vec{\alpha}^T A_2) = \infty$) to see that

$$\mathbb{P}_{\vec{\gamma}}[\vec{\alpha}^{T}A_{2}\vec{\gamma} = \varphi(\vec{\alpha},\vec{\beta}) \,|\, \vec{\alpha},\vec{\beta}\,] \le n^{-1/2+2\delta} \frac{1}{h_{1}(\vec{\alpha}^{T}A_{2})} + n^{-1+2\delta}.$$

So, we have

$$\begin{split} \mathbb{P}[\mathcal{E} \cap \mathcal{E}_{0}^{c} \cap \mathcal{E}_{1}] &\leq \mathbb{E}_{\vec{\alpha}} \left[\left(n^{-1+4\delta} \cdot \frac{1}{h_{2}(\vec{\alpha}^{T}A_{1})} + n^{-3/2+4\delta} \right) \cdot \left(n^{-1/2+2\delta} \frac{1}{h_{1}(\vec{\alpha}^{T}A_{2})} + n^{-1+2\delta} \right) \right] \\ &\leq n^{-3/2+6\delta} \mathbb{E}_{\vec{\alpha}} \left[\frac{1}{h_{1}(\vec{\alpha}^{T}A_{2})h_{2}(\vec{\alpha}^{T}A_{1})} \right] + n^{-2+6\delta} \mathbb{E}_{\vec{\alpha}} \left[\frac{1}{h_{2}(\vec{\alpha}^{T}A_{1})} \right] \\ &+ n^{-2+6\delta} \mathbb{E}_{\vec{\alpha}} \left[\frac{1}{h_{2}(\vec{\alpha}^{T}A_{2})} \right] + n^{-5/2+6\delta} \end{split}$$

(Here we used that $h_1(\vec{\alpha}^T A_2) \ge h_2(\vec{\alpha}^T A_2)$). Using Lemma 10.6, we deduce

$$\mathbb{P}[\mathcal{E} \cap \mathcal{E}_0^c \cap \mathcal{E}_1] \le n^{-2-1/6+7\delta} + 2n^{-2-1/6+7\delta} + n^{-5/2+6\delta} \le n^{-2-1/6+\varepsilon}/3$$

The third term. For any outcome of $\vec{\alpha}$ for which \mathcal{E}_0 does not hold (in particular, $\|\vec{\alpha}^T A_1\|_0 > n^{1-\delta}$), Theorem 10.2 (with q = 2, using that $h_0(\vec{\alpha}^T A_1) = \infty$) tells us that

$$\mathbb{P}_{\vec{\beta}}[\vec{\alpha}^{T}A_{1}\vec{\beta}=0\,|\,\vec{\alpha}\,] \leq \frac{1}{h_{1}(\vec{\alpha}^{T}A_{1})}n^{-1/2+2\delta} + n^{-1+2\delta}.$$

Then, write M for the $2 \times n$ matrix whose rows are $\vec{\alpha}^T A_2$ and $\vec{\beta}^T A_3$ (which we may view as a sequence of n vectors in \mathbb{F}^2). For any outcomes of $\vec{\alpha}, \vec{\beta}$ for which \mathcal{E}_1 does not hold, another application of Theorem 10.2 (with q = 3, using that $h_0(M) = h_1(M) = \infty$, and $h_2(M) \ge h_2(\vec{\alpha}^T A_2)$) tells us that

$$\mathbb{P}_{\vec{\gamma}}[\vec{\alpha}^{T}A_{2}\vec{\gamma} = \varphi(\vec{\alpha},\vec{\beta}) \text{ and } \vec{\beta}^{T}A_{3}\vec{\gamma} = \psi(\vec{\alpha},\vec{\beta}) \,|\,\vec{\alpha},\vec{\beta}\,] \le \frac{1}{h_{2}(\vec{\alpha}^{T}A_{2})}n^{-1+3\delta} + n^{-3/2+3\delta}.$$

So, we have

$$\begin{split} \mathbb{P}[\mathcal{E} \cap \mathcal{E}_{0}^{c} \cap \mathcal{E}_{1}^{c}] &\leq \mathbb{E}_{\vec{\alpha}} \left[\left(\frac{1}{h_{1}(\vec{\alpha}^{T}A_{1})} n^{-1/2+2\delta} + n^{-1+2\delta} \right) \cdot \left(\frac{1}{h_{2}(\vec{\alpha}^{T}A_{2})} n^{-1+3\delta} + n^{-3/2+3\delta} \right) \right] \\ &\leq n^{-3/2+5\delta} \mathbb{E}_{\vec{\alpha}} \left[\frac{1}{h_{1}(\vec{\alpha}^{T}A_{1})h_{2}(\vec{\alpha}^{T}A_{2})} \right] + n^{-2+5\delta} \mathbb{E}_{\vec{\alpha}} \left[\frac{1}{h_{2}(\vec{\alpha}^{T}A_{1})} \right] \\ &+ n^{-2+5\delta} \mathbb{E}_{\vec{\alpha}} \left[\frac{1}{h_{2}(\vec{\alpha}^{T}A_{2})} \right] + n^{-5/2+5\delta} \end{split}$$

(Here we used that $h_1(\vec{\alpha}^T A_1) \ge h_2(\vec{\alpha}^T A_1)$). Proceeding in the same way as for the second term, it follows from Lemma 10.6 that $\mathbb{P}[\mathcal{E} \cap \mathcal{E}_0^c \cap \mathcal{E}_1^c] \le n^{-2-1/6+\varepsilon}/3$, as desired. \Box

10.2. **Proofs of technical estimates.** In this subsection we prove Lemma 10.6, which is the remaining ingredient for the proof of Lemma 2.17. It will be convenient to introduce some additional notation.

Definition 10.8. For a matrix $A \in \mathbb{C}^{n \times n}$ and a subset $J \subseteq \{1, \ldots, n\}$, recall that A[J] denotes the submatrix of A containing just the columns indexed by J. For $r \in \mathbb{N}$ and a vector $\vec{u} \in \mathbb{C}^m$, let $h_r^*(\vec{u})$ be the minimum volume of a symmetric GAP of rank at most r that contains all the entries of \vec{u} .

We start by proving Lemma 10.6(1), which is a fairly simple consequence of Lemma 9.2.

Proof of Lemma 10.6(1). Since $1/h_2^{\delta}(\vec{\alpha}^T A)$ only takes values in the range [0, 1], we have

$$\mathbb{E}\left[\frac{1}{h_2^{\delta}(\vec{\alpha}^T A)}\right] = \int_0^1 \mathbb{P}\left[\frac{1}{h_2^{\delta}(\vec{\alpha}^T A)} > u\right] du = \int_1^\infty V^{-2} \mathbb{P}[h_2^{\delta}(\vec{\alpha}^T A) < V] dV.$$
(10.3)

To apply this identity, we need to estimate probabilities of the form $\mathbb{P}[h_2^{\delta}(\vec{\alpha}^T A) < V]$, as follows.

Claim 10.9. For any $V \ge 1$, we have $\mathbb{P}[h_2^{\delta}(\vec{\alpha}^T A) < V] \le O_{\delta}(V^3(1 + \log V)n^{-1/2}).$

Proof of claim. We will apply Lemma 3.1 to reduce studying $h_2^{\delta}(\vec{\alpha}^T A)$ to studying $h_2^{\delta}(\vec{\alpha}^T A[J])$ for 3element subsets $J \subseteq \{1, \ldots, n\}$. Similar arguments appeared in the proof of Lemma 2.13, so we will be quite brief with details. Suppose that $\mathbb{P}[h_2^{\delta}(\vec{\alpha}^T A) < V] > C_{\delta}V^3(1 + \log V)n^{-1/2}$. We will show that this is a contradiction for large C_{δ} .

First, Lemma 3.1 tells us that for all but an $6n^{-\delta}$ -fraction of 3-element subsets $J \subseteq \{1, \ldots, n\}$ we have $\mathbb{P}[h_2^*(\vec{\alpha}^T A[J]) < V] > C_{\delta} V^3(1 + \log V)n^{-1/2}/2$. Say that such J are good.

For each good J, Lemma 9.2 (with m = 3 and k = 1 and r = 2) tells us that A[J] has at most $O(n/C_{\delta}^2)$ disjoint nonsingular 3×3 submatrices, so by Fact 3.4 at most a $O(C_{\delta}^{-2})$ -fraction of all the 3×3 submatrices of A[J] are nonsingular.

Recalling that all but a $6n^{-\delta}$ -fraction of J are good, we see that the total number of nonsingular 3×3 submatrices in A is at most $6n^{-\delta} {\binom{n}{3}}^2 + O(C_{\delta}^{-2}) {\binom{n}{3}}^2$, which contradicts the δ -robustness of A for sufficiently large C_{δ} .

Now, given (10.3) and Claim 10.9, we compute

$$\mathbb{E}\left[\frac{1}{h_2^{\delta}(\vec{\alpha}^T A)}\right] \le \int_1^{n^{1/6}} V^{-2} \cdot O_{\delta}(V^3(1+\log V)n^{-1/2}) \, dV + \int_{n^{1/6}}^{\infty} V^{-2} \, dV = n^{-1/6+o(1)},$$

as desired.

The proof of Lemma 10.6(2) is similar, but Lemma 9.2 is not quite strong enough to get a non-trivial bound. Instead, we need the following more refined estimate of Costello [9, Lemma 8]¹⁶. He proved this via (elementary) number-theoretic considerations.

Lemma 10.10. Consider any vectors $\vec{u}, \vec{v} \in \mathbb{C}^m$, such that for each $i \in \{1, \ldots, m\}$, at least one of u_i and v_i is nonzero. Let $\vec{\alpha} \in \mathbb{C}^m$ be a vector of *i.i.d.* lazy Rademacher random variables. For any real number $q \geq 1$, let $\mathcal{E}(q)$ be the event that there are integers $x, y \in \mathbb{Z}$ such that

- $|x|, |y| \le q$,
- $x\vec{u} y\vec{v}$ has at least m/10 nonzero entries,
- $x(\vec{u} \cdot \vec{\alpha}) = y(\vec{v} \cdot \vec{\alpha}).$

Then for any $1 \leq q \leq \sqrt{m}$ we have

$$\mathbb{P}[\mathcal{E}(q)] \le qm^{-1/2 + o(1)}.$$

Lemma 10.10 has the following corollary in terms of rank-1 GAPs.

Corollary 10.11. Suppose $A \in \mathbb{C}^{n \times 2}$ contains at least ℓ disjoint nonsingular 2×2 submatrices. Then, for every $1 \le q \le \sqrt{2\ell}$,

$$\mathbb{P}[h_1^*(\vec{\alpha}^T A) < q] \le q \cdot n^{-1/2 + o(1)}$$

Proof. Write \vec{u} and \vec{v} for the two columns of A. Let $I \subseteq \{1, \ldots, n\}$ be a set of 2ℓ row indices corresponding to ℓ disjoint nonsingular 2×2 submatrices, so for each $i \in I$ at least one of u_i and v_i is nonzero. Note that for any $x, y \in \mathbb{C}$ which are not both zero, the vector $x\vec{u}[I] - y\vec{v}[I]$ has at least ℓ nonzero entries.

Now, if the two entries $\vec{u} \cdot \vec{\alpha}$ and $\vec{v} \cdot \vec{\alpha}$ of $\vec{\alpha}^T A$ lie in a symmetric rank-1 GAP with volume 2q + 1, then there are integers $x, y \in \mathbb{Z}$ with $|x|, |y| \leq q$ and $x(\vec{u} \cdot \vec{\alpha}) = y(\vec{v} \cdot \vec{\alpha})$, such that x and y are not both zero. We also know that $x(\vec{u} \cdot \vec{\alpha}) = x(\vec{u}[I] \cdot \vec{\alpha}[I])$ and $y(\vec{v} \cdot \vec{\beta}) = y(\vec{v}[I] \cdot \vec{\beta}[I])$. We then apply Lemma 10.10 (with $m = 2\ell$) to $\vec{u}[I]$ and $\vec{v}[I]$.

Now we prove Lemma 10.6(2).

Proof of Lemma 10.6(2). Here, our starting point is the inequality

$$\mathbb{E}\left[\frac{1}{h_{1}^{\delta}(\vec{\alpha}^{T}A_{1})h_{2}^{\delta}(\vec{\alpha}^{T}A_{2})}\right] = \int_{1}^{\infty} t^{-2} \mathbb{P}[h_{1}^{\delta}(\vec{\alpha}^{T}A_{1})h_{2}^{\delta}(\vec{\alpha}^{T}A_{2}) < t] dt$$

$$\leq \int_{1}^{n^{2/3}} t^{-2} \mathbb{P}[h_{1}^{\delta}(\vec{\alpha}^{T}A_{1}) < t^{3/4} \text{ or } h_{2}^{\delta}(\vec{\alpha}^{T}A_{2}) < t^{1/4}] dt + \int_{n^{2/3}}^{\infty} t^{-2} dt$$

$$\leq \int_{1}^{n^{2/3}} t^{-2} \left(\mathbb{P}[h_{1}^{\delta}(\vec{\alpha}^{T}A_{1}) < t^{3/4}] + \mathbb{P}[h_{2}^{\delta}(\vec{\alpha}^{T}A_{2}) < t^{1/4}]\right) dt + n^{-2/3}. \quad (10.4)$$

To apply (10.4), we need Claim 10.9, together with the following additional estimate.

Claim 10.12. For any $1 \le V \le \sqrt{2\delta n}$, we have $\mathbb{P}[h_1^{\delta}(\vec{\alpha}^T A_1) < V] \le V \cdot n^{-1/2 + o(1)}$.

Claim 10.12 can be proved using Corollary 10.11, in a very similar way to the proof of Claim 10.9 using the r = 2 case of Lemma 9.2, but now studying $h_1^{\delta}(\vec{\alpha}^T A[J])$ for 2-element subsets $J \subseteq \{1, \ldots, n\}$ and counting nonsingular 2×2 submatrices (recalling Fact 3.5 to obtain the final contradiction). Now, substituting Claims 10.9 and 10.12 into (10.4), we obtain

$$\mathbb{E}\left[\frac{1}{h_1^{\delta}(\vec{\alpha}^T A_1)h_2^{\delta}(\vec{\alpha}^T A_2)}\right] \leq \int_1^{n^{2/3}} t^{-2} \left(t^{3/4}n^{-1/2+o(1)} + t^{3/4} \cdot n^{-1/2+o(1)}\right) dt + n^{-2/3}$$
$$\leq n^{-1/2+o(1)} \cdot \int_1^{n^{2/3}} t^{-5/4} dt + n^{-2/3}$$
$$= n^{-1/2+o(1)} \cdot (n^{2/3})^{-1/4} + n^{-2/3} = n^{-2/3+o(1)}.$$

¹⁶[9, Lemma 8] is stated for Rademacher random variables, not lazy Rademacher random variables. This does not have any impact on the proof (also, one can deduce the lazy Rademacher case from the Rademacher case by first revealing which of ξ_1, \ldots, ξ_n are zero; conditionally, the nonzero ξ_i are Rademacher).

References

- N. Alon and R. Beigel, Lower bounds for approximations by low degree polynomials over z/sub m/, Proceedings 16th Annual IEEE Conference on Computational Complexity, CCC-01, IEEE Comput. Soc, 2001, p. 184–187. 36
- [2] M.-F. Balcan, Y. Li, D. P. Woodruff, and H. Zhang, *Testing matrix rank, optimally*, Proceedings of the Thirtieth Annual ACM-SIAM Symposium on Discrete Algorithms, SIAM, Philadelphia, PA, 2019, pp. 727–746.
- [3] A. Bhattacharyya and Y. Yoshida, Property testing—problems and techniques, Springer, Singapore, 2022. 6
- [4] T. D. Browning and D. R. Heath-Brown, Counting rational points on hypersurfaces, J. Reine Angew. Math. 584 (2005), 83–115. 5
- [5] T. D. Browning, D. R. Heath-Brown, and P. Salberger, Counting rational points on algebraic varieties, Duke Math. J. 132 (2006), no. 3, 545–578. 5
- [6] T. D. Browning, Quantitative arithmetic of projective varieties, Progress in Mathematics, vol. 277, Birkhäuser Verlag, Basel, 2009. 2
- [7] A. Cohen and G. Moshkovitz, Partition and analytic rank are equivalent over large fields, Duke Math. J. 172 (2023), no. 12, 2433–2470. 5
- [8] S. D. Cohen, The distribution of Galois groups and Hilbert's irreducibility theorem, Proc. London Math. Soc. (3) 43 (1981), no. 2, 227–250. 5
- K. P. Costello, Bilinear and quadratic variants on the Littlewood-Offord problem, Israel J. Math. 194 (2013), no. 1, 359–394. 2, 3, 6, 12, 18, 33
- [10] K. P. Costello, T. Tao, and V. Vu, Random symmetric matrices are almost surely nonsingular, Duke Math. J. 135 (2006), no. 2, 395–413. 2, 9, 11
- [11] K. P. Costello and V. H. Vu, The rank of random graphs, Random Structures Algorithms 33 (2008), no. 3, 269–285.
 11
- [12] V. De la Pena and E. Giné, Decoupling: from dependence to independence, Springer Science & Business Media, 1999.
 9
- [13] F. de Zeeuw, A survey of Elekes-Rónyai-type problems, New trends in intuitive geometry, Bolyai Soc. Math. Stud., vol. 27, János Bolyai Math. Soc., Budapest, 2018, pp. 95–124. 2
- [14] W. Doeblin, Sur les sommes d'un grand nombre de variables aléatoires indépendantes, Bull. Sci. Math 63 (1939), no. 2, 23–32. 2
- [15] P. Erdős, On a lemma of Littlewood and Offord, Bull. Amer. Math. Soc. 51 (1945), 898–902. 2
- [16] A. Ferber, V. Jain, and Y. Zhao, On the number of Hadamard matrices via anti-concentration, Combin. Probab. Comput. 31 (2022), no. 3, 455–477. 6
- [17] K. Ford, The distribution of integers with a divisor in a given interval, Ann. of Math. (2) 168 (2008), no. 2, 367–433.
 5
- [18] J. Fox, M. Kwan, and H. Spink, Geometric and o-minimal Littlewood-Offord problems, Ann. Probab. 51 (2023), no. 1, 101–126. 7, 8, 9
- [19] A. Frieze and B. Pittel, Perfect matchings in random graphs with prescribed minimal degree, Mathematics and computer science. III, Trends Math., Birkhäuser, Basel, 2004, pp. 95–132. 14
- [20] O. Goldreich, Introduction to property testing, Cambridge University Press, Cambridge, 2017. 6
- [21] W. T. Gowers and T. Karam, Equidistribution of high-rank polynomials with variables restricted to subsets of \mathbb{F}_p , arXiv preprint arXiv:2209.04932 (2022). 4
- [22] B. Green and T. Tao, The distribution of polynomials over finite fields, with applications to the Gowers norms, Contrib. Discrete Math. 4 (2009), no. 2, 1–36. 4
- [23] G. Halász, Estimates for the concentration function of combinatorial number theory and probability, Period. Math. Hungar. 8 (1977), no. 3-4, 197–211. 6, 10, 13
- [24] D. R. Heath-Brown, The density of rational points on curves and surfaces, Ann. of Math. (2) 155 (2002), no. 2, 553–595. 5
- [25] R. Howard, Estimates on the concentration function of sets in R^d: Notes on lectures of Oskolkov, https://people.math.sc.edu/howard/Notes/concentration.pdf, 2000. 6
- [26] O. Janzer, Polynomial bound for the partition rank vs the analytic rank of tensors, Discrete Anal. (2020), Paper No. 7, 18. 4, 5
- [27] D. M. Kane, The correct exponent for the Gotsman-Linial conjecture, Comput. Complexity 23 (2014), no. 2, 151–175.
- [28] T. Karam, High-rank subtensors of high-rank tensors, arXiv preprint arXiv:2207.08030 (2022). 6
- [29] Y. R. Katznelson, Singular matrices and a uniform bound for congruence groups of $SL_n(\mathbf{Z})$, Duke Math. J. **69** (1993), no. 1, 121–136. 27
- [30] Y. R. Katznelson, Integral matrices of fixed rank, Proc. Amer. Math. Soc. 120 (1994), no. 3, 667–675. 27
- [31] D. J. Kleitman, On a lemma of Littlewood and Offord on the distributions of linear combinations of vectors, Advances in Math. 5 (1970), 155–157 (1970). 2
- [32] A. Kolmogorov, Sur les propriétés des fonctions de concentrations de M. P. Lévy, Ann. Inst. H. Poincaré 16 (1958), 27–34. 2
- [33] R. Krauthgamer and O. Sasson, Property testing of data dimensionality, Proceedings of the Fourteenth Annual ACM-SIAM Symposium on Discrete Algorithms (Baltimore, MD, 2003), ACM, New York, 2003, pp. 18–27. 6
- [34] M. Kwan and L. Sauermann, Resolution of the quadratic Littlewood-Offord problem, arXiv preprint arXiv:2312.13826. 2, 5, 11, 13
- [35] M. Kwan and L. Sauermann, An algebraic inverse theorem for the quadratic Littlewood-Offord problem, and an application to Ramsey graphs, Discrete Anal. (2020), Paper No. 12, 34. 4, 22
- [36] A. Lampert and T. Ziegler, Relative rank and regularization, Forum Math. Sigma 12 (2024), Paper No. e29, 26. 5

- [37] Y. Li, Z. Wang, and D. P. Woodruff, Improved testing of low rank matrices, Proceedings of the 20th ACM SIGKDD international conference on Knowledge discovery and data mining, KDD '14, ACM, August 2014, p. 691–700. 6
- [38] J. E. Littlewood and A. C. Offord, On the number of real roots of a random algebraic equation. III, Rec. Math. [Mat. Sbornik] N.S. 12(54) (1943), 277–286. 2
- [39] S. Lovett, The analytic rank of tensors and its applications, Discrete Anal. (2019), Paper No. 7, 10. 5
- [40] R. Meka, O. Nguyen, and V. Vu, Anti-concentration for polynomials of independent random variables, Theory Comput. 12 (2016), Paper No. 11, 16. 1, 2
- [41] L. Milićević, Polynomial bound for partition rank in terms of analytic rank, Geom. Funct. Anal. 29 (2019), no. 5, 1503–1530. 4, 5
- [42] G. Moshkovitz and D. G. Zhu, Quasi-linear relation between partition and analytic rank, arXiv preprint arXiv:2211.05780 (2022). 5
- [43] E. Naslund, Exponential bounds for the Erdős-Ginzburg-Ziv constant, J. Combin. Theory Ser. A 174 (2020), 105185, 19. 5
- [44] E. Naslund, The partition rank of a tensor and k-right corners in \mathbb{F}_q^n , J. Combin. Theory Ser. A **174** (2020), 105190, 25.5
- [45] H. Nguyen and V. Vu, Optimal inverse Littlewood-Offord theorems, Adv. Math. 226 (2011), no. 6, 5298–5319. 6, 10
- [46] H. H. Nguyen, Inverse Littlewood-Offord problems and the singularity of random symmetric matrices, Duke Math. J. 161 (2012), no. 4, 545–586.
- [47] H. H. Nguyen and V. H. Vu, Small ball probability, inverse theorems, and applications, Erdős centennial, Bolyai Soc. Math. Stud., vol. 25, János Bolyai Math. Soc., Budapest, 2013, pp. 409–463. 2
- [48] J. Pila, Density of integral and rational points on varieties, no. 228, 1995, Columbia University Number Theory Seminar (New York, 1992), pp. 4, 183–187. 5
- [49] A. Razborov and E. Viola, Real advantage, ACM Trans. Comput. Theory 5 (2013), no. 4, Art. 17, 8. 2
- [50] D. Ron, Property testing: A learning theory perspective, now Publishers Inc, 2007. 6
- [51] J. Rosiński and G. Samorodnitsky, Symmetrization and concentration inequalities for multilinear forms with applications to zero-one laws for Lévy chaos, Ann. Probab. 24 (1996), no. 1, 422–437. 2, 9, 11
- [52] M. Rudelson and R. Vershynin, The Littlewood-Offord problem and invertibility of random matrices, Adv. Math. 218 (2008), no. 2, 600–633. 6
- [53] P. Salberger, Counting rational points on projective varieties, Proc. Lond. Math. Soc. (3) 126 (2023), no. 4, 1092–1133.
- [54] W. Sawin and T. Tao, Notes on the "slice rank" of tensors, https://terrytao.wordpress.com/2016/08/24/ notes-on-the-slice-rank-of-tensors/., 2011. 8
- [55] W. M. Schmidt, The density of integer points on homogeneous varieties, Acta Math. 154 (1985), no. 3-4, 243-296. 5
- [56] A. Sidorenko, A correlation inequality for bipartite graphs, Graphs Combin. 9 (1993), no. 2, 201–204. 11
- [57] T. Tao, A symmetric formulation of the croot-lev-pach-ellenberg-gijswijt capset bound, https://terrytao.wordpress. com/2016/05/18/a-symmetric-formulation-of-the-croot-lev-pach-ellenberg-gijswijt-capset-bound/., 2011. 8
- [58] T. Tao and V. Vu, A sharp inverse Littlewood-Offord theorem, Random Structures Algorithms 37 (2010), no. 4, 525–539.6
- [59] T. Tao and V. H. Vu, Inverse Littlewood-Offord theorems and the condition number of random discrete matrices, Ann. of Math. (2) 169 (2009), no. 2, 595–632. 6
- [60] F. Vermeulen, Dimension growth for affine varieties, Int. Math. Res. Not. IMRN (2024), no. 15, 11464–11483. 5
- [61] V. Vu, Anti-concentration inequalities for polynomials, A journey through discrete mathematics, Springer, Cham, 2017, pp. 801–810. 2
- [62] M. N. Walsh, Bounded rational points on curves, Int. Math. Res. Not. IMRN (2015), no. 14, 5644-5658. 5

Appendix A. A Littlewood–Offord theorem for varieties

In this appendix we prove Theorem 2.5, giving a bound of the form $\mathbb{P}[\xi_1 \vec{a}_1 + \cdots + \xi_n \vec{a}_n \in \mathbb{Z}] \leq n^{-1/2+o(1)}$ unless almost all the vectors $\vec{a}_1, \ldots, \vec{a}_n$ lie in a linear subspace \mathcal{W} such that some translate $\vec{w} + \mathcal{W}$ of \mathcal{W} is contained in \mathbb{Z} .

Proof of Theorem 2.5. Assume E1 does not hold; we will prove the bound in E2. We may also assume that $\mathcal{Z} \neq \emptyset$ (otherwise E2 holds trivially). Let P_1, \ldots, P_z be nonzero polynomials with $\mathcal{Z} = \{\vec{y} \in \mathbb{F}^d : P_1(\vec{y}) = \cdots = P_z(\vec{y}) = 0\}$, and let $q^* = \max(\deg P_1, \ldots, \deg P_z)$. We first find a subspace $\mathcal{W} \subseteq \mathbb{F}^d$, for which one can find "many" disjoint basis among the vectors

We first find a subspace $\mathcal{W} \subseteq \mathbb{F}^d$, for which one can find "many" disjoint basis among the vectors $\vec{a}_1, \ldots, \vec{a}_n$. Writing $n_b = n - \varepsilon n(d-b)/d$, let $b \in \{0, \ldots, d\}$ be the minimum integer such that at least n_b of the vectors \vec{a}_i lie in a common linear subspace $\mathcal{W} \subseteq \mathbb{F}^d$ of dimension b (certainly such a b exists, considering b = d). We cannot have b = 0, as otherwise **E1** would hold with $\mathcal{W} = \{\vec{0}\}$ and with \vec{w} as an arbitrary element of \mathcal{Z} . In addition, there is no proper subspace $\mathcal{W}' \subsetneq \mathcal{W}$ containing at least n_{b-1} of the vectors \vec{a}_i , meaning that among the vectors \vec{a}_i we can find at least $(n_b - n_{b-1})/b \ge \varepsilon n/d^2$ disjoint bases of \mathcal{W} .

Upon relabelling the indices, without loss of generality we may assume that $\vec{a}_{jb+1}, \vec{a}_{jb+2}, \ldots, \vec{a}_{(j+1)b}$ form a basis of \mathcal{W} for all integers $0 \leq j < \lfloor \varepsilon n/d^2 \rfloor$. Let $m = \lfloor \varepsilon n/d^2 \rfloor \cdot b$, then we have $\vec{a}_1, \ldots, \vec{a}_m \in \mathcal{W}$. Now, condition on any outcome of $(\xi_i)_{i>m}$; we will prove the desired bound conditional on this outcome (all probabilistic notation for the rest of the proof is implicitly with respect to the corresponding conditional probability space).

Let $\vec{w} = \sum_{i>m} \xi_i \vec{a}_i$ (which we no longer view as a random variable). Fix a linear isomorphism $\varphi : \mathcal{W} \to \mathbb{F}^b$ and let $\mathcal{Z}' = \varphi(\mathcal{W} \cap (\mathcal{Z} - \vec{w})) \subseteq \mathbb{F}^b$. Recall that at least $n_b \ge (1 - \varepsilon)n$ of the vectors \vec{a}_i lie in \mathcal{W} , and we are assuming **E1** does not hold, so $\vec{w} + \mathcal{W}$ is not fully contained in \mathcal{Z} , meaning that $\mathcal{Z}' \subseteq \mathbb{F}^b$.

Let $P \in \mathbb{F}[x_1, \ldots, x_b]$ be a nonzero polynomial of degree $q := \deg P \leq q^*$ vanishing on \mathcal{Z}' . To see that such a polynomial exists, note that $\mathcal{W} \cap (\mathcal{Z} - \vec{w}) = \{\vec{y} \in \mathcal{W} : P_1(\vec{y} + \vec{w}) = \cdots = P_z(\vec{y} + \vec{w}) = 0\}$, and hence $\mathcal{Z}' \subset \mathbb{F}^b$ can be described as the vanishing set of up to z polynomials with degrees at most deg $P_1, \ldots, \deg P_z$, respectively (at least one of these polynomials is nonzero, since $\mathcal{Z}' \neq \mathbb{F}^b$).

Now, for i = 1, ..., m, let $\vec{a}'_i = \varphi(\vec{a}_i)$. Then for each $j = 0, ..., \lfloor \varepsilon n/d^2 \rfloor -1$, the vectors $\vec{a}'_{jb+1}, ..., \vec{a}'_{(j+1)b}$ form a basis of \mathbb{F}^b . Furthermore, we have

$$\mathbb{P}[\xi_{1}\vec{a}_{1} + \dots + \xi_{n}\vec{a}_{n} \in \mathcal{Z}] = \mathbb{P}[\xi_{1}\vec{a}_{1} + \dots + \xi_{m}\vec{a}_{m} \in \mathcal{Z} - \vec{w}] = \mathbb{P}[\xi_{1}\vec{a}_{1} + \dots + \xi_{m}\vec{a}_{m} \in \mathcal{W} \cap (\mathcal{Z} - \vec{w})] \\ = \mathbb{P}[\xi_{1}\vec{a}_{1}' + \dots + \xi_{m}\vec{a}_{m}' \in \mathcal{Z}'] \le \mathbb{P}[P(\xi_{1}\vec{a}_{1}' + \dots + \xi_{m}\vec{a}_{m}') = 0].$$

Note that we can interpret $P(\xi_1 \vec{a}'_1 + \dots + \xi_m \vec{a}'_m)$ as a polynomial of degree (at most) q in the variables ξ_1, \dots, ξ_m . It suffices to show that this polynomial has at least $\varepsilon' m^q$ nonzero coefficients for $\varepsilon' = \min(\varepsilon/2, 1/(2d)^{q^*})$. Indeed, then the Meka–Nguyen–Vu bound for the polynomial Littlewood–Offord problem (see Theorem 1.1) implies

$$\mathbb{P}[\xi_1 \vec{a}_1 + \dots + \xi_n \vec{a}_n \in \mathcal{Z}] \le \mathbb{P}[P(\xi_1 \vec{a}_1' + \dots + \xi_m \vec{a}_m') = 0] \le m^{-1/2 + \varepsilon'} \le n^{-1/2 + \varepsilon},$$

establishing **E2**.

For any distinct indices $i(1), \ldots, i(q) \in \{1, \ldots, m\}$, the coefficient of $\xi_{i(1)} \cdots \xi_{i(q)}$ in the polynomial $P(\xi_{l}\vec{a}'_{1} + \cdots + \xi_{m}\vec{a}'_{m})$ is the same as the coefficient of $\xi_{i(1)} \cdots \xi_{i(q)}$ in $P(\xi_{i(1)}\vec{a}'_{i(1)} + \cdots + \xi_{i(q)}\vec{a}'_{i(q)})$ and hence the same as the coefficient of $t_1 \cdots t_q$ in $P(t_1\vec{a}'_{i(1)} + \cdots + t_q\vec{a}'_{i(q)})$. But note that for any vectors $\vec{v}_1, \ldots, \vec{v}_q \in \mathbb{F}^b$ the coefficient of $t_1 \cdots t_q$ in $P(t_1\vec{v}_1 + \cdots + t_q\vec{v}_q)$ is a multilinear function of $\vec{v}_1, \ldots, \vec{v}_q$ (more formally, the function assigning a q-tuple $(\vec{v}_1, \ldots, \vec{v}_q)$ the value of this coefficient is a multilinear function $\mathbb{F}^b \times \cdots \times \mathbb{F}^b \to \mathbb{F}$). Since P has degree q, this multilinear function is nonzero (indeed, the homogeneous degree-q part P_q of the polynomial P is non-zero, so we can find a vector $\vec{w} \in \mathbb{F}^b$ with $P_q(\vec{w}) \neq 0$, and observe that the coefficient of $t_1 \cdots t_q$ in $P(t_1\vec{w} + \cdots + t_q\vec{w})$ is the same as in $P_q(t_1\vec{w} + \cdots + t_q\vec{w}) = (t_1 + \cdots + t_q)^q P_q(\vec{w})$ and therefore equal to $q!P_q(\vec{w}) \neq 0$). This means that, given any bases B_1, \ldots, B_q of \mathbb{F}^b , we can choose vectors $\vec{v}_1 \in B_1, \ldots, \vec{v}_q \in B_q$ such that the coefficient of $t_1 \cdots t_q$ in $P(t_1\vec{v}_{11} + \cdots + t_q\vec{v}_q)$ is nonzero. In particular, for any distinct $j(1), \ldots, j(q) \in \{0, \ldots, \lfloor \varepsilon n/d^2 \rfloor - 1\}$, we can find $i(1) \in \{j(1)b + 1, \ldots, (j(1) + 1)b\}, \ldots, i(q) \in \{j(q)b + 1, \ldots, (j(q) + 1)b\}$ such that the coefficient of $t_1 \cdots t_q$ in $P(\xi_{i(1)}\vec{a}'_{i(1)} + \cdots + \xi_{i(q)}\vec{a}'_{i(q)})$ is nonzero. This means that we can find at least $\binom{\lfloor \varepsilon n/d^2 \rfloor}{q} = \binom{m/b}{q} > (m/(2d))^q \geq m^q/(2d)^{q^*} \geq \varepsilon' m^q$ distinct q-element sets $\{i(1), \ldots, i(q)\} \subset \{1, \ldots, m\}$, such that the coefficient of $\xi_{i(1)} \cdots \xi_{i(q)}$ in the polynomial $P(\xi_1\vec{a}'_1 + \cdots + \xi_m\vec{a}'_m)$ is nonzero. Thus, the polynomial $P(\xi_1\vec{a}'_1 + \cdots + \xi_m\vec{a}'_m)$ has indeed at least $\varepsilon' m^q$ nonzero coefficients. This completes the proof.

Appendix B. A counterexample to some conjectures of Costello (By Matthew Kwan, Ashwin Sah and Mehtaab Sawhney)

Here we prove Proposition 1.4. First, the second bullet point is straightforward: note that

$$\mathbb{P}[f(\xi_1, \dots, \xi_n) = 0] \ge \mathbb{P}[L_1(\xi_1, \dots, \xi_n) = 0 \text{ and } L_{2d}(\xi_1, \dots, \xi_n) = 0] = \Omega((1/\sqrt{n})^2) \ge \varepsilon/n$$

for sufficiently small ε .

In the rest of this appendix we prove that **B1** does not hold for this polynomial f, showing that there are no non-constant polynomials g_1, g_2 for which $f - g_1g_2$ has fewer than εn^d nonzero terms. The idea is that if $f - g_1g_2$ had fewer than εn^d nonzero terms (i.e., if f had an "approximate factorisation"), this would lead to an *exact* factorisation of the irreducible polynomial $y_1 \cdots y_d + y_{d+1} \cdots y_{2d}$, which is impossible.

To execute this idea, we start with a random sampling argument to strengthen our "approximate factorisation", showing that the only obstruction comes from non-multilinear terms. We then use a Ramsey-theoretic argument (related to ideas of Alon and Beigel [1]) to "clean up" g_1 and g_2 . We start with some preparations.

Definition B.1. Fix a vector $\vec{t} = (t_1, \ldots, t_r) \in \mathbb{N}^r$, with entries summing to $k = \|\vec{t}\|_1$. Let I_1, \ldots, I_r be disjoint sets of size m, and let $K_m(\vec{t})$ be the k-uniform hypergraph on the vertex set $I_1 \cup \cdots \cup I_r$, containing every possible edge e for which $|e_1 \cap I_1| = t_1, \ldots, |e_r \cap I_r| = t_r$.

For example, $K_m((k))$ is the complete k-uniform hypergraph on m vertices, and $K_m((1,1))$ is the complete bipartite graph on m + m vertices. We will need the following Ramsey-type theorem.

Lemma B.2. For any $r, s, b \ge 1$ and $\vec{t} \in \mathbb{N}^r$, there is $M \in \mathbb{N}$ such that the following holds. Let $m \ge M$ and consider any colouring of the hyperedges of $K_m(\vec{t})$ with b different colours. Then there are subsets $I'_1 \subseteq I_1, \ldots, I'_r \subseteq I'_r$ of size at least s such that all the edges of $K_m(\vec{t})$ inside $I'_1 \cup \cdots \cup I'_r$ have the same colour.

Proof. Put $k := \|\vec{t}\|_1$ and let G be the complete k-uniform hypergraph on the vertex set $\{1, \ldots, m\}$. We will use the edge-colouring of $K_m(\vec{t})$ to define an edge-colouring of G, to which we will apply Ramsey's theorem.

Let j(0) = 0, and for $\ell \in \{1, \ldots, k\}$, define $j(\ell)$ to satisfy $t_1 + \cdots + t_{j(\ell)-1} < \ell \leq t_1 + \cdots + t_{j(\ell)}$. Writing $I_j = \{i_1^j, \ldots, i_m^j\}$ for each $j \in \{1, \ldots, r\}$, we define a mapping φ from edges of G to edges of $K_m(\vec{t})$ as follows. For an edge $e = \{q(1), \ldots, q(k)\}$ of G, where $q(1) < q(2) < \cdots < q(k)$, let $\varphi(e) = \{i_{q(1)}^{j(1)}, \ldots, i_{q(k)}^{j(k)}\}$.

Via this mapping, our edge-colouring of $K_m(\vec{t})$ induces an edge-colouring of G. So, by Ramsey's theorem, assuming m is sufficiently large there is a subset Q of sr vertices of G such that all edges of G inside Q have the same colour. Order the elements of Q as

$$q(1,1) < \dots < q(1,s) < \dots < q(r,1) < \dots < q(r,s),$$

} let $I'_{i} = \{i^{j}_{\sigma(i,1)}, \dots, i^{j}_{\sigma(i,c)}\}.$

and for each $j \in \{1, ..., r\}$ let $I'_j = \{i^j_{q(j,1)}, ..., i^j_{q(j,s)}\}.$

We also need a simple combinatorial fact about of sums of vectors.

Definition B.3. For a vector $\vec{v} \in \mathbb{N}^r$ let \vec{v}_{\downarrow} be the *decreasing rearrangement* of \vec{v} (obtained by sorting the entries of \vec{v} in decreasing order). Recall that in the *lexicographic order* \preceq on \mathbb{N}^r , we write $\vec{p} \preceq \vec{q}$ if $\vec{p} = \vec{q}$ or if $p_i < q_i$ for the first i where p_i and q_i differ.

Lemma B.4. Let $\vec{s_1}, \vec{s_2}, \vec{t_1}, \vec{t_2} \in \mathbb{N}^r$ be vectors with $\vec{s_1} + \vec{s_2} = \vec{t_1} + \vec{t_2} = \vec{v}$. If $(\vec{s_1}, \vec{s_2}) \neq (\vec{t_1}, \vec{t_2})$ then $(\vec{s_1} + \vec{t_2})_{\downarrow} \succ \vec{v_{\downarrow}}$ or $(\vec{t_1} + \vec{s_2})_{\downarrow} \succ \vec{v_{\downarrow}}$.

Proof. Without loss of generality, suppose that $\vec{v}_{\downarrow} = \vec{v}$. Note that the condition $(\vec{s}_1, \vec{s}_2) \neq (\vec{t}_1, \vec{t}_2)$ is equivalent to the condition $\vec{s}_1 \neq \vec{t}_1$ (since $\vec{s}_1 + \vec{s}_2 = \vec{t}_1 + \vec{t}_2$). Let $i \in \{1, \ldots, r\}$ be minimal such that $s_{1,i} \neq t_{1,i}$. If $s_{1,i} > t_{1,i}$ then $s_{2,i} < t_{2,i}$ and $(\vec{s}_1 + \vec{t}_2)_{\downarrow} \succeq \vec{s}_1 + \vec{t}_2 \succ \vec{v}$. On the other hand, if $s_{1,i} < t_{1,i}$ then $s_{2,i} > t_{2,i}$ and $(\vec{t}_1 + \vec{s}_2)_{\downarrow} \succeq \vec{v}$.

Now we prove Proposition 1.4.

Proof of Proposition 1.4. Fix $d \ge 2$. As discussed at the start of this appendix, we need to show that for our particular polynomial f, **B1** does not hold. Suppose for the purpose of contradiction that

$$f = g_1 g_2 + p, \tag{B.1}$$

for some polynomials $g_1g_2, p \in \mathbb{C}[x_1, \ldots, x_n]$ such that $\deg(g_1), \deg(g_2) \ge 1$ and $\deg(g_1g_2), \deg(p) \le d$, and p has at most εn^d nonzero coefficients. Let $d_1 = \deg(g_1)$ and $d_2 = d - d_1 \ge \deg(g_2)$.

Let N be a large integer (we will need it to be sufficiently large in terms of d, at a later point in the proof).

Claim B.5. If $\varepsilon > 0$ is sufficiently small (in terms of N, d), then there are subsets $I'_1 \subseteq I_1, \ldots, I'_{2d} \subseteq I_{2d}$ of size N, such that for every set $\{i_1, \ldots, i_d\} \in I'_1 \cup \cdots \cup I'_{2d}$ of d distinct indices, the coefficient of $x_{i_1} \ldots x_{i_d}$ in p is zero.

Proof of claim. Independently for each $j \in \{1, \ldots, 2d\}$, let $I'_j \subseteq I_j$ be a uniformly random subset of size N. We will show that the sets I'_j satisfy the desired property with positive probability.

Indeed, note that for any particular set $Q \subseteq I_j$, we have $\mathbb{P}[Q \subseteq I'_j] \leq (N/|I_j|)^{|Q|}$. Thus, for each monomial $x_{i_1} \ldots x_{i_d}$ of p with nonzero coefficient and distinct indices i_1, \ldots, i_d (there are at most εn^d of them), the probability of having $\{i_1, \ldots, i_d\} \subseteq I'_1 \cup \cdots \cup I'_{2d}$ is at most $(N/(2\lfloor n/(4d) \rfloor))^d < 1/(\varepsilon n^d)$ for sufficiently small ε . The claim follows from a union bound over all such monomials.

Claim B.5 tells us that if we restrict our polynomials f, g_1, g_2, p to the variables in $I'_1 \cup \cdots \cup I'_{2d}$, then we have a slightly "cleaner" version of (B.1) in which we have eliminated all the degree-d multilinear terms in p. Unfortunately, the non-multilinear degree-d terms in p can still cause problems, so we need to refine the situation further.

For $S \subset I'_1 \cup \cdots \cup I'_{2d}$, let type $(S) = (|S \cap I'_1|, \ldots, |S \cap I'_{2d}|)$ be a vector encoding the number of elements of S that lie in each I'_j . Slightly abusively, we conflate the set $S \subset I'_1 \cup \cdots \cup I'_{2d}$ with the multilinear monomial $\prod_{i \in S} x_i$ (so we can also talk about the type of a multilinear monomial, or the type of a multilinear term of a polynomial). Say a *k*-type is a vector $\vec{t} \in \mathbb{N}^{2d}$ with entries summing to *k*. So, the degree- d_1 multilinear terms of g_1 can be categorised by d_1 -type, the degree- d_2 multilinear terms of g_2 can be categorised by d_2 -type, and the terms of f (all of which are degree-d, and multilinear) can be categorised by d-type.

For $z \in \mathbb{C}$, with argument $\arg(z) \in [0, 2\pi)$, define the *direction*

$$\operatorname{dir}(z) = \begin{cases} * & \text{if } z = 0 \\ \pi/2 & \text{if } \arg(z) \in [\pi/4, 3\pi/4) \\ \pi & \text{if } \arg(z) \in [3\pi/4, 5\pi/4) \\ 3\pi/2 & \text{if } \arg(z) \in [5\pi/4, 7\pi/4) \\ 0 & \text{otherwise.} \end{cases}$$

In other words, if $z \in \mathbb{C}$ is nonzero, then dir(z) is chosen such that $\arg(z)$ differs from dir(z) by an angle of at most $\pi/4$. Slightly abusively, we conflate each term of a polynomial with its coefficient (so we can talk about the directions of terms of g_1 and g_2). We will need the following (easy) fact about directions.

Fact B.6. For any $x_1, \ldots, x_k, y_1, \ldots, y_k \in \mathbb{C} \setminus \{0\}$, if

$$\operatorname{dir}(x_1) = \cdots = \operatorname{dir}(x_k), \quad \operatorname{dir}(y_1) = \cdots = \operatorname{dir}(y_k),$$

then $x_1y_1 + \cdots + x_ky_k \neq 0$.

Now, by iteratively applying Lemma B.2 (once for each d_1 -type and once for each d_2 -type), assuming N is sufficiently large (in terms of d), we can find $I''_1 \subseteq I'_1, \ldots, I''_{2d} \subseteq I'_{2d}$, each of size d, such that, among multilinear degree- d_1 terms of g_1 (respectively, degree- d_2 terms of g_2) containing variables indexed by $I''_1 \cup \cdots \cup I''_{2d}$, the direction of a term depends only on its type.

Now, let $h_1, h_2, q, R_1, \ldots, R_{2d} \in \mathbb{C}[x_i : i \in I''_1 \cup \cdots \cup I''_{2d}]$ be the $2d^2$ -variable polynomials obtained from $g_1, g_2, q, L_1, \ldots, L_{2d}$ by setting all variables not indexed by $I''_1 \cup \cdots \cup I''_{2d}$ to zero. (So, in particular, we have $R_j = \sum_{i \in I''_i} x_i$ for $j = 1, \ldots, 2d$.) Then, (B.1) gives rise to the "cleaner" identity

$$R_1 \cdots R_d + R_{d+1} \cdots R_{2d} = h_1 h_2 + q, \tag{B.2}$$

where q has no multilinear degree-d terms (by Claim B.5), and in h_1 (respectively, h_2), the direction of a degree- d_1 (respectively, degree d_2) multilinear term only depends on its type. So, we can talk about the direction of a d_1 -type in h_1 or the direction of a d_2 -type in h_2 .

Now, say that a type is *simple* if all of its entries are at most 1.

Claim B.7. Let $\vec{t_1}$ be a d_1 -type and let $\vec{t_2}$ be a d_2 -type, such that $\vec{t_1} + \vec{t_2}$ is not simple. Then $\vec{t_1}$ has direction * in h_1 or $\vec{t_2}$ has direction * in h_2 .

Proof of claim. Suppose for the purpose of contradiction that the statement of this claim is false, so there are $\vec{t_1}, \vec{t_2}$ such that $\vec{t_1} + \vec{t_2}$ is not simple but $\vec{t_1}$ does not have direction * in h_1 and $\vec{t_2}$ does not have direction * in h_2 . Choose such a pair $\vec{t_1}, \vec{t_2}$ such that $(\vec{t_1} + \vec{t_2})_{\downarrow}$ is lexicographically maximal.

Now, let $\{i_1, \ldots, i_d\} \subseteq I''_1 \cup \cdots \cup I''_{2d}$ be a set of distinct indices with type $\vec{t}_1 + \vec{t}_2$ (such a set exists, since each I''_j has size d). Since $\vec{t}_1 + \vec{t}_2$ is not simple, the coefficient of $x_{i_1} \cdots x_{i_d}$ in $R_1 \cdots R_d + R_{d+1} \ldots R_{2d}$ is zero. By (B.2), this coefficient can also be written as

$$0 = \sum_{A,B} \alpha_A \beta_B, \tag{B.3}$$

where we write α_A for the coefficient of $\prod_{i \in A} x_i$ in h_1 , and we write β_B for the coefficient of $\prod_{i \in B} x_i$ in h_2 , and the sum is over all partitions of $\{i_1, \ldots, i_d\}$ into a set of A of size d_1 and a set B of size d_2 . Note that for all such A, B we always have type(A) + type $(B) = \vec{t_1} + \vec{t_2}$.

By Fact B.6, and the choice of \vec{t}_1, \vec{t}_2 (and the meaning of the direction *), there is a nonzero contribution to (B.3) from pairs A, B with type $(A) = \vec{t}_1$ and type $(B) = \vec{t}_2$. So, there must be an additional d_1 -type \vec{s}_1 and an additional d_2 -type \vec{s}_2 , such that $(\vec{s}_1, \vec{s}_2) \neq (\vec{t}_1, \vec{t}_2)$ and such that $\vec{s}_1 + \vec{s}_2 = \vec{t}_1 + \vec{t}_2$, but \vec{s}_1 does not have direction * in h_1 and \vec{s}_2 does not have direction * in h_2 . But by Lemma B.4, at least one of $(\vec{t}_1 + \vec{s}_2)_{\downarrow}$ or $(\vec{s}_1 + \vec{t}_2)_{\downarrow}$ is lexicographically greater than $(\vec{t}_1 + \vec{t}_2)_{\downarrow}$ (and therefore not simple), which contradicts the choice of \vec{t}_1, \vec{t}_2 .

Now, say a term of h_1 (respectively, of h_2) is good if it has degree d_1 (respectively, degree d_2) and is multilinear. If two terms with types \vec{t}_1, \vec{t}_2 share a variable, then $\vec{t}_1 + \vec{t}_2$ is not simple. So, Claim B.7 tells us that no variable appears in nonzero good terms of both h_1 and h_2 . Now, arbitrarily choose $i_1 \in I_1'', \ldots, i_{2d} \in I_{2d}''$. Let $h_1^*, h_2^* \in \mathbb{C}[x_{i_1}, \ldots, x_{i_{2d}}]$ be the 2*d*-variable polynomials obtained from h_1, h_2 by setting all variables other than $x_{i_1}, \ldots, x_{i_{2d}}$ to zero, and deleting all terms which are not good. By the above discussion, h_1^* and h_2^* involve disjoint sets of variables (so each term of $h_1^*h_2^*$ is multilinear, so there can be no cancellation between $h_1^*h_2^*$ and q). Also, there can be no cancellation between $h_1^*h_2^*$ and terms arising from the non-good terms omitted in h_1^* and h_2^* (since the terms arising this way cannot be degree-d multilinear). So, (B.2) gives rise to the identity

$$x_{i_1} \cdots x_{i_d} + x_{i_{d+1}} \cdots x_{i_{2d}} = h_1^* h_2^*.$$

(Specifically, on both sides of (B.2), we have set all variables other than $x_{i_1}, \ldots, x_{i_{2d}}$ to zero, and restricted to degree-*d* multilinear terms). But it is easy to see that this factorisation is impossible.

Department of Mathematics, ETH, Zürich, Switzerland $\mathit{Email}\ address:\ \texttt{zhihan.jin@math.ethz.ch}$

Institute of Science and Technology Austria (ISTA). Am Campus 1, 3400 Klosterneuburg, Austria. *Email address:* matthew.kwan@ist.ac.at

Institute for Applied Mathematics, University of Bonn. Endenicher Allee 60, 53115 Bonn, Germany. *Email address:* sauermann@iam.uni-bonn.de

Institute of Science and Technology Austria (ISTA). Am Campus 1, 3400 Klosterneuburg, Austria. *Email address:* yiting.wang@ist.ac.at