Inducing Graphs, Hypergraphs, and Tournaments

A Thesis

submitted to
Indian Institute of Science Education and Research Pune
in partial fulfilment of the requirements for the
BS-MS Dual Degree Programme

by

Arjun Ranganathan

Indian Institute of Science Education and Research Pune
Dr. Homi Bhabha Road,
Pashan, Pune 411008, INDIA.

May, 2024

Supervisors: Dr Matthew Kwan, Dr Jonathan Noel

Certificate

This is to certify that this dissertation entitled "Inducing Graphs, Hypergraphs, and Tournaments" towards the partial fulfilment of the BS-MS dual degree programme at the Indian Institute of Science Education and Research, Pune represents study/work carried out by Arjun Ranganathan at the Institute of Science and Technology Austria and the Indian Institute of Science Education and Research, Pune under the joint supervision of Dr Matthew Kwan, Assistant Professor, Institute of Science and Technology Austria and Dr Jonathan Noel, Assistant Professor, University of Victoria during the academic year 2023-2024.

Jon Noel

Dr Matthew Kwan

Dr Jonathan Noel

Committee:

Dr Matthew Kwan

Dr Jonathan Noel

Prof. Soumen Maity

Declaration

I hereby declare that the matter embodied in the report entitled "Inducing Graphs, Hypergraphs, and Tournaments" are the results of the work carried out by me at the Institute of Science and Technology Austria and the Indian Institute of Science Education and Research, Pune, under the joint supervision of Dr Matthew Kwan, Institute of Science and Technology Austria and Dr Jonathan Noel, University of Victoria, and the same has not been submitted elsewhere for any other degree. Wherever others contribute, every effort is made to indicate this clearly, with due reference to the literature and acknowledgement of collaborative research and discussions.

Arjun Ranganathan

Registration Number: 20181043

Acknowledgements

First and foremost, I would like to thank my supervisors, Matthew Kwan and Jonathan Noel, for their exceptional mentoring. They always went the extra mile and advised me on much more than just this thesis. Their guidance has lent me perspective on how to approach mathematics, teaching, and pursuing a career in this field, and I hope to carry forward all I have learnt from them – both mathematical and otherwise.

My journey in mathematics began because of the wonderful faculty at IISER, and I am forever grateful for their stimulating courses and discussions. In particular, I wish to thank Prof. Soumen Maity for teaching me graph theory, guiding my first reading project in combinatorics, allowing me to participate in his group meetings, and providing me with my first teaching opportunity. All these experiences significantly inspired my enthusiasm for combinatorics. I also wish to thank Dr Moumanti Podder for taking the time to discuss all kinds of difficult problems and guide me through advanced concepts.

My time at IISER has been fraught with ups and downs, and I am ever thankful to have had my friends and my family by my side to celebrate the joys, help me through the hard times, and everything in between. Dropping a year with COVID-19 was a tremendously difficult chapter in my life, but it led to me forming more relationships than I could have imagined. I am sure many of these people would object to me tarnishing their names by mentioning them in my humble thesis (a convenient reason to excuse myself from thanking a list of far too many friends), but I hope that all of those who are important to me know that I am ever appreciative for their companionship.

Abstract

In this thesis, we consider two problems in extremal and probabilistic combinatorics.

The first question relates to the so-called *Edge-Statistics conjecture* of Alon, Hefetz, Krivelevich and Tyomkyn, stated as follows: given integers $k \geq 1$ and $0 < \ell < \binom{k}{2}$, when a uniformly random k-vertex subset is sampled from a large graph, the probability of it inducing precisely ℓ edges is at most $1/e + o_k(1)$. Independent recent studies have proven this conjecture in different parametrised regimes of k and ℓ, yielding a complete proof overall. We begin by briefly discussing some of the methods used in these papers and then shift our focus to the generalisation of the conjecture to uniform hypergraphs, a direction of research suggested by Alon *et al.*. We detail our attempts to extend the techniques used in the aforementioned papers to the hypergraph framework and conclude with a discussion on the limitations of our approaches and potential ways to address them.

The second problem we tackle addresses the notion of quasirandom tournaments. A sequence of tournaments $\{T_n\}_{n\geq 1}$ is said to be *quasirandom* if and only if every tournament on k vertices appears as a subtournament in T_n with a density that is $1+o(1)$ times its expected density in a random tournament. A tournament H is said to be *quasirandom-forcing* if it has the following property: if the density of H in T_n is $1 + o(1)$ times the expected density of H in a random tournament, then this is true for every subtournament of T_n, i.e., $\{T_n\}_{n\geq 1}$ is quasirandom. Recent papers have proven that the only quasirandom-forcing tournaments are transitive tournaments on at least four vertices and one particular tournament on five vertices. We generalise this notion of forcing quasirandomness to pairs of tournaments and make partial progress towards characterising all non-transitive quasirandom-forcing pairs. We then state the methods we intend to use to handle the remaining pairs and finish with some open problems regarding pairs containing a transitive tournament.

Contributions

All the contents of this thesis were primarily completed by the author under the joint supervision of Dr Matthew Kwan and Dr Jonathan Noel. Additionally, the work contained in Chapters 4 and 5 involved research discussions with Lina Maria Simbaqueba Marin, University of Victoria.

Contents

Acknowledgements **vii**

Abstract **ix**

Contributions **xi**

1 Edge-Statistics in Uniform Hypergraphs: Introduction **1**

1.1 Introduction 1

1.2 Preliminaries 3

2 The Superlinear Regime **5**

2.1 Main Results 5

2.2 Notation 6

2.3 Broad Outline of the Proofs 6

2.4 Probabilistic Techniques and Results 7

2.5 Proof of Theorem 2.1.1 10

2.6 Proof of Theorem 2.1.2 12

2.7 Attempts to Generalise to r-graphs 14

3 The Sparse Regime **17**

3.1 Main Results 17

3.2 Notation 18

3.3 Proof of Theorem 3.1.2 18

3.4 Probabilistic Results 21

3.5 Proof of Theorem 3.1.1 21

3.6 Extension to r-graphs and Further Discussion 31

4 Quasirandom-Forcing Pairs of Tournaments: Introduction and Preliminaries 35

4.1 Introduction 35

4.2 Preliminaries 38

5 Classifying Pairs of Tournaments 41

5.1 Tournament Limits 41

5.2 Finding Tournamentons 44

5.3 Further Directions of Research: Characterising the Remaining Pairs 47

5.4 Open Problems 50

Chapter 1

Edge-Statistics in Uniform Hypergraphs: Introduction

1.1 Introduction

The notion of *graph inducibility* was introduced by Pippenger and Golumbic in 1975 [40]. Roughly speaking, the inducibility of a graph H measures the maximum number of induced copies of H a large graph can have. Determining the inducibility of different graphs is a longstanding question that has enjoyed a recent surge in popularity (see [3, 28, 49, 25]).

Motivated by its ties to graph inducibility, Alon, Hefetz, Krivelevich and Tyomkyn [1] initiated the study of the following question: given integers $k \geq 1$ and $0 \leq \ell \leq \binom{k}{2}$, what is the maximum probability that a uniformly random k-vertex subset of a very large graph induces exactly ℓ edges? In order to make this precise, we will first recall some of their notation. Given a graph G and any subset of its vertex set $U \subseteq V(G)$, let $G[U]$ denote the subgraph induced by U and $e(G[U])$ be the number of edges it contains. Let A be a uniformly random k-vertex subset of $V(G)$ and set $X_{G,k} := e\left(G[A]\right)$ to be the random variable corresponding to the number of edges induced by the random set A. Then define

$$I(n,k,\ell) := \max\left\{\mathbb{P}\left[X_{G,k} = \ell\right] : |V(G)| = n\right\},$$

the maximum probability of $X_{G,k} = \ell$ across all n-vertex graphs G. Note that $I(n,k,\ell)$ may also be interpreted as the maximum density of k-vertex ℓ-edge induced subgraphs among all n-vertex graphs. Standard averaging arguments show that $I(n,k,\ell)$ monotonically decreases in n (see, for instance, [40, Proposition 1]). Hence, one can define the *edge-inducibility of* k *and* ℓ as the limit

$$\mathrm{ind}(k,\ell) := \lim_{n\to\infty} I(n,k,\ell).$$

Note that by taking complements, we have

$$\mathrm{ind}(k,\ell) = \mathrm{ind}\left(k, \binom{k}{2} - \ell\right)$$

By considering an empty or a complete graph respectively, it is easy to see that $\mathrm{ind}(k,0) = \mathrm{ind}\left(k,\binom{k}{2}\right) = 1$. However, Ramsey's theorem states that if n is large enough, then any graph on n vertices must have either a clique or an independent set on k vertices, meaning that $I(n,k,\ell) < 1$ for $\ell \notin \left\{0, \binom{k}{2}\right\}$, and hence $\mathrm{ind}(k,\ell) < 1$ as $I(n,k,\ell)$ monotonically decreases in n. Consequently, Alon *et al.* [1] put forth the *Edge-Statistics conjecture*, which states that $\mathrm{ind}(k,\ell) \le 1/e + o_k(1)$ whenever $0 < \ell < \binom{k}{2}$.

They provide two examples to motivate this conjecture and show its tightness. Firstly, by considering the random graph $G\left(n, 1/\binom{k}{2}\right)$ and computing the expected density of k-vertex subsets that induce only one edge, one can easily see that $\mathrm{ind}(k,1) \ge 1/e + o_k(1)$. Additionally, by looking at the complete bipartite graph on n vertices with parts of size n/k and $(k-1)n/k$ and computing the asymptotic density of copies of $K_{1,k-1}$ inside it, one obtains $\mathrm{ind}(k,k-1) \ge 1/e + o_k(1)$.

As noted previously, since the edge-inducibility is invariant under complements, it suffices to focus on the setting where $1 \le \ell \le \frac{1}{2}\binom{k}{2}$. Kwan, Sudakov and Tran [30] proved the Edge-Statistics conjecture when $c \cdot k \le \ell \le \binom{k}{2} - c \cdot k$ for some sufficiently large constant c. Subsequently, Fox and Sauermann [16] proved it when $1 \le \ell \le c \cdot k$ for any constant $c > 0$, and Martinsson, Mousset, Noever and Trujić [34] proved it in the regime $1 \le \ell \le o_k(k^{6/5})$. The first result along with either of the latter two yields a complete proof of the conjecture, and hence, we may formally state it as a theorem as below:

Theorem 1.1.1 (Edge-Statistics for Graphs). *For all $k, \ell \in \mathbb{N}$ satisfying $0 < \ell < \binom{k}{2}$, we have*

$$\mathrm{ind}(k,\ell) \le 1/e + o_k(1).$$

Furthermore, Kwan *et al.* [30] initiated the study of the edge-inducibility in the setting of uniform hypergraphs, a line of investigation originally proposed by Alon *et al.* [1]. Given integers $r \geq 2$, $k \geq 1$ and $0 \leq \ell \leq \binom{k}{r}$, for any r-uniform hypergraph G on at least k vertices, let $I_r(G, k, \ell)$ be the probability that a uniformly random k-vertex subset of G induces precisely ℓ edges. Analogous to the graph scenario, we set $I_r(n, k, \ell)$ to be the maximum of $I_r(G, k, \ell)$ over all r-uniform hypergraphs G on n vertices. Again, as $I_r(n, k, \ell)$ monotonically decreases with n, the limit $\mathrm{ind}_r(k, \ell) := \lim_{n\to\infty} I_r(n, k, \ell)$ is well defined. With this notation, we see that $\mathrm{ind}(k, \ell) = \mathrm{ind}_2(k, \ell)$. As before, the edge-inducibility for $\ell \in \{0, \binom{k}{r}\}$ is exactly equal to 1 and is strictly smaller than 1 otherwise.

Thus, Alon *et al.* [1] suggested the following natural extension of the Edge-Statistics conjecture:

Conjecture 1.1.2 (Edge-Statistics for Uniform Hypergraphs). *For all $r \geq 2$, $k \geq 1$ and $0 < \ell < \binom{k}{r}$, we have*

$$\mathrm{ind}_r(k, \ell) \leq 1/e + o_k(1).$$

This hypergraph generalisation of the Edge-Statistics conjecture has been settled only in very few cases (note that there is now a much wider regime for ℓ in terms of k). For $r = 3$, Kwan *et al.* [30] resolved it for very dense hypergraphs, where $\ell = \Omega_k(k^3)$. In the sublinear setting, Fox *et al.* [16] proved the conjecture for all $r \geq 2$ and $1 \leq \ell \leq o_k(k)$.

Our aim in this project is to study this extension of the Edge-Statistics conjecture to the edge-inducibility of uniform hypergraphs. Primarily, we have attempted to generalise the approaches of [30] and [34]. After introducing some preliminary notation and definitions in the next section, we then proceed to Chapters 2 and 3 where we discuss the proofs used by these two papers in the graph case for specific regimes of ℓ and k, mention our progress to date in extending these ideas, and the further improvements and alternative approaches that could overcome the issues we faced and complete the proofs.

1.2 Preliminaries

As discussed in Chapter 1, we have mainly attempted to extend the techniques of [30] and [34] to the r-uniform hypergraph generalisation of the Edge-Statistics conjecture 1.1.2.

The approaches of these two papers are vastly different, so specialised notation is introduced individually in Chapters 2 and 3. However, we shall still introduce some basic notation that shall be common throughout.

Our (hyper)graph theoretic notation is standard. A *hypergraph* $G = (V(G), E(G))$ consists of a pair of sets, where $E(G)$ is a family of subsets of $V(G)$. The elements of $V(G)$ are called the *vertices* of G, and those of $E(G)$ are called its *egdes* or *hyperedges*. The number of edges is denoted $e(G) = |E(G)|$. If $\{v_1, v_2, \ldots, v_r\} \in E(G)$ is an edge, we sometimes write $v_1 v_2 \ldots v_r \in E(G)$. If every edge consists of exactly r vertices, we say that G is an *r-uniform hypergraph* or an *r-graph* for short. A *graph* is simply a 2-uniform hypergraph. In a graph, two vertices u and v are said to be *neighbours* or are *adjacent* if they constitute an edge, and we write $u \in N(v)$ and $v \in N(u)$.

For any vertex $v \in V(G)$, the *degree* of v in G, denoted $d_G(v)$ is the number of edges containing v. More generally, given a set of i vertices $\{v_1, v_2, \ldots, v_i\} \in V(G)$, we define the *co-degree* of $v_1, v_2, \ldots, v_i$, denoted $d_G(v_1, v_2, \ldots, v_i)$, to be the number of edges containing the set $\{v_1, v_2, \ldots, v_i\}$. Often, in case there is no risk of ambiguity, we may drop the subscript G. For any subset $U \subseteq V(G)$ of the vertex set, we let $G[U]$ denote the hypergraph induced by U. For any vertex $u \in U$, we may write $d_U(u)$ in place of $d_{G[U]}(u)$. For short, we may also write $e(U)$ for $e(G[U])$.

For a positive integer n, we write $[n]$ to denote the set $\{1, 2, \ldots, n\}$. For a set S, we use $\binom{S}{k}$ to denote all k-vertex subsets of S, and $\binom{S}{\leq k}$ to denote all subsets of S of size at most k.

We use standard asymptotic notation. Given any two functions $f(n)$ and $g(n)$, we say $f = O(g)$ if there is some positive constant C such that $|f(n)| \leq C|g(n)|$ for all sufficiently large n, we say $f = \Omega(g)$ if there is some positive constant c such that $|f(n)| \geq c|g(n)|$ for all sufficiently large n, we say $f = \Theta(G)$ if $f = O(g)$ and $f = \Omega(g)$, and finally we say $f = o(g)$ or $g = \omega(f)$ if there exists a sequence of reals c_n decreasing to zero such that $|f(n)| \leq c_n|g(n)|$ for sufficiently large n (if $g(n)$ is nonzero for all n, this is equivalent to saying $f/g \to 0$ as $n \to \infty$). If there are multiple parameters, we may use a subscript to denote which one is going to infinity (such as o_k or O_n, for instance) in case there is a chance for confusion.

Chapter 2

The Superlinear Regime

2.1 Main Results

As indicated in Chapter 1, Kwan *et al.* [30] proved the graph version of the conjecture for $\Omega_k(k) \leq \ell \leq \binom{k}{2} - \Omega_k(k)$ and the 3-uniform extension for $\ell \geq \Omega_k(k^3)$. Specifically, they prove the following two stronger results which imply the above.

Theorem 2.1.1. *For all* $0 \leq \ell \leq \binom{k}{2}$*, let* $\ell^* = \min\{\ell, \binom{k}{2} - \ell\}$*. We have*

$$\mathrm{ind}(k, \ell) \leq \log^{O(1)}(\ell^*/k)\sqrt{k/\ell^*}.$$

Theorem 2.1.2. *For all* k, ℓ *satisfying* $\ell^* = \min\{\ell, \binom{k}{3} - \ell\} = \Omega_k(k^3)$*, we have*

$$\mathrm{ind}_3(k, \ell) \leq \log^{O(1)}(k)/\sqrt{k}.$$

This chapter presents a description of their approach and a brief sketch of their proof. We conclude with a discussion of our attempts to extend their ideas to the general r-uniform framework.

2.2 Notation

Throughout Chapter 2, all asymptotics will be in the limit $n \to \infty$ unless otherwise stated.

For a sequence $\boldsymbol{x} = (x_1, x_2, \ldots, x_n)$ of zeroes and ones, we write $|\boldsymbol{x}|$ to denote the number of ones in $\boldsymbol{x}$. For any $S \subseteq [n]$, we use $\boldsymbol{x}^S$ to denote the monomial $\prod_{i \in S} x_i$.

2.3 Broad Outline of the Proofs

Given a graph G, let $A_G = (a_{xy})_{x,y}$ denote its adjacency matrix, i.e., $a_{xy} = 1$ if $xy \in E(G)$ and $a_{xy} = 0$ otherwise. Observe that we can express $X_{G,k}$ as a homogenous quadratic polynomial

$$X_{G,k} = \sum_{1 \le x < y \le n} a_{xy} \xi_x \xi_y = \frac{1}{2} \boldsymbol{\xi} A_G \boldsymbol{\xi}^T, \tag{2.1}$$

where $\boldsymbol{\xi} = (\xi_1, \ldots, \xi_n)$ is a uniformly random length-n zero-one vector with exactly k ones. Proving Theorem 2.1.1 now essentially boils down to bounding the probability $\mathbb{P}\left[\frac{1}{2}\boldsymbol{\xi} A_G \boldsymbol{\xi}^T = \ell\right]$. This suggests the usage of *qudaratic anticoncentration inequalities* for the random variable $\boldsymbol{\xi}$. Clearly, as the number of ones in $\boldsymbol{\xi}$ is constrained to be exactly equal to k, it is not a sequence of independent random variables. However, several papers [14, 36, 41, 38] have studied probabilities of the form $\mathbb{P}\left[\boldsymbol{x} A \boldsymbol{x}^T = \ell\right]$ where $\boldsymbol{x}$ is a sequence of independent zero-one random variables, and we wish to use these ideas.

For any $0 \le k \le n$, let $\mathrm{BL}(n, k)$ be the uniform distribution on binary strings $\boldsymbol{x} \in \{0,1\}^n$ with $|\boldsymbol{x}| = k$. This is precisely the distribution of $\boldsymbol{\xi}$ described above and is sometimes called the uniform distribution "on the slice" of the Boolean hypercube or the limiting distribution of the Bernoulli-Laplace model of diffusion. When $n = 2k$, a coupling argument realises $\boldsymbol{\xi}$ as a function of a uniform random permutation σ of $[n]$ and a sequence of $n/2$ i.i.d. random variables γ. Conditioning on σ, the random variable $X_{G,k}$ may now be expressed as a quadratic multilinear polynomial $f_\sigma(\gamma)$, to which one applies a polynomial anticoncentration inequality [36, Theorem 1.6] to bound the probability $\mathbb{P}[f_\sigma(\gamma) = \ell]$. The bound given by this inequality depends on the nonzero coefficients of maximum degree, i.e., the nonzero degree-2 coefficients – the more such coefficients, the stronger the bound. These coefficients can be shown to correspond to certain substructures in G. Specifically, alternating 3-paths in G will lead to nonzero degree-2 coefficients in $f_\sigma(\gamma)$. If we can prove there are many such subgraphs in G, it will yield the desired result. This is done by first concentrating the number of edges of G – we argue $X_{G,k}$ is tightly concentrated around its expectation, and so if $e(G)$ is too big or too small, then $\mathbb{P}[X_{G,k} = \ell]$ will be very small. Then, using a random greedy algorithm, we can deduce that G contains many alternating 3-paths with high probability.

The hypergraph setting, however, is much more complicated. For any r-uniform hypergraph, using the same ideas as above, it is clear that we can write $X_{G,k}$ as a homogeneous degree-r polynomial in $\boldsymbol{\xi}$,

$$X_{G,k} = \sum_{S \in \binom{[n]}{r}} a_S \boldsymbol{\xi}^S. \tag{2.2}$$

Once again, the same coupling argument can be used to express $X_{G,k}$ as some $f_\sigma(\gamma)$. However, even in the 3-uniform case with $\ell = \Omega_k(k^3)$, there could end up being very few nonzero degree-3 coefficients, which is why the same inequality as last time cannot be used directly. This is why Theorem 2.1.2 is much weaker than Theorem 2.1.1. In this case, we provide a structural characterisation of 3-graphs where $f_\sigma(\gamma)$ has very few nonzero degree-3 coefficients. Then we proceed via a modified version of the aforementioned anticoncentration inequality [36, Theorem 1.6] that accounts for nonzero degree-2 coefficients but provides a weaker bound than the original.

We now proceed to provide a brief sketch of the details of the proofs. Note that we will not be presenting the proofs in entirety or replicating the paper. We will sometimes only state key results and explain their usage in the proof.

2.4 Probabilistic Techniques and Results

We first present the probabilistic arguments that play important roles in the proof.

First, we state the coupling result to write $\boldsymbol{\xi} \in \mathrm{BL}(n, n/2)$ as a function of σ, a uniformly random permutation of $[n]$, and $\gamma \in \mathrm{Rad}^{n/2}$, a sequence of $n/2$ i.i.d. random variables from the Rademacher distribution (the uniform distribution on $\{-1, 1\}$). As mentioned previously, expressing $\boldsymbol{\xi}$ in terms of a sequence of i.i.d. random variables will allow us to use well-established anticoncentration inequalities. The following result appears in the proof of [31, Proposition 4.10].

Fact 2.4.1. *If σ is a uniformly random permutation of $[n]$ and $\gamma \in \mathrm{Rad}^{n/2}$ is a sequence of $n/2$ i.i.d. Rademacher random variables, then the distribution $\boldsymbol{\xi}$ can be obtained as follows. Set $\xi_{\sigma(i)} = 1$ for all i such that $\gamma_i = 1$, set $\xi_{\sigma(i+n/2)} = 1$ for all i such that $\gamma_i = -1$, and set $\xi_j = 0$ for all other indices j.*

Now, to actually utilise Fact 2.4.1, we need to see how it translates polynomials of $\boldsymbol{\xi} \in \mathrm{BL}(n, n/2)$ into polynomials of $\boldsymbol{\gamma} \in \mathrm{Rad}^{n/2}$.

Lemma 2.4.2. *Suppose $\boldsymbol{\xi} \in BL(n, n/2)$ and X is a random variable that is a degree-d polynomial in $\boldsymbol{\xi}$ of the form*

$$X = \sum_{S \in \binom{[n]}{d}} a_S \boldsymbol{\xi}^S.$$

We use Fact 2.4.1 to express X as a function of γ, σ. If we condition on any outcome of σ, then X is a multilinear polynomial in the γ_i of degree at most d. For any subset $I \subseteq [n/2]$ with $|I| \geq d-1$, say $I = \{i_1, \dots, i_q\}$, the coefficient g_I of γ^I is

$$\frac{1}{2^d} \sum_{\boldsymbol{b} \in \{0,1\}^q} (-1)^{|\boldsymbol{b}|} a\left(\left\{\sigma\left(i_j + b_j \frac{n}{2}\right) : 1 \leq j \leq q\right\}\right),$$

where for any $R \subseteq [n]$, $a(R)$ denotes the sum over all a_S with $S \supseteq R$.

In the context of our problem, the result above implies the following. Suppose G is a graph on $n = 2k$ vertices with adjacency matrix $A_G = (a_{ij})_{i,j}$. Then, by Equation 2.1 and Lemma 2.4.2, if we condition on any outcome of σ, then $X_{G,k}$ is a quadratic polynomial in $\{\gamma_1, \dots, \gamma_k\}$, and the coefficient of $\gamma_i\gamma_j$ is

$$\frac{1}{4}\left(a_{\sigma(i)\sigma(j)} - a_{\sigma(i+k)\sigma(j)} - a_{\sigma(i)\sigma(j+k)} + a_{\sigma(i+k)\sigma(j+k)}\right). \tag{2.3}$$

Similarly, in the 3 -graph scenario, if G is a 3-graph on $n = 2k$ vertices, by Equation 2.2 and Lemma 2.4.2, if we condition on any outcome of σ, then $X_{G,k}$ is a cubic polynomial in $\{\gamma_1, \dots, \gamma_k\}$, and the coefficient of $\gamma_i\gamma_j\gamma_q$ is

$$\begin{aligned}\frac{1}{8}\Big(&a_{\sigma(i)}a_{\sigma(j)}a_{\sigma(q)} - a_{\sigma(i+k)}a_{\sigma(j)}a_{\sigma(q)} - a_{\sigma(i)}a_{\sigma(j+k)}a_{\sigma(q)} - a_{\sigma(i)}a_{\sigma(j)}a_{\sigma(q+k)} + a_{\sigma(i+k)}a_{\sigma(j+k)}a_{\sigma(q)} \\ &+ a_{\sigma(i+k)}a_{\sigma(j)}a_{\sigma(q+k)} + a_{\sigma(i)}a_{\sigma(j+k)}a_{\sigma(q+k)} - a_{\sigma(i+k)}a_{\sigma(j+k)}a_{\sigma(q+k)}\Big).\end{aligned} \tag{2.4}$$

Next, we present the anticoncentration inequalities used for functions of a sequence of i.i.d. Rademacher random variables. For $\boldsymbol{x} = (x_1, \ldots, x_n)$, suppose we have a degree-d polynomial in $\boldsymbol{x}$ of the form

$$f(\boldsymbol{x}) = \sum_{S \in \binom{[n]}{d}} f_S \boldsymbol{x}^S.$$

Construct an auxiliary d-uniform hypergraph H on the vertex set $[n]$ with any $S \in \binom{[n]}{d}$ an edge if and only if the coefficient f_S is nonzero. The *rank* of the polynomial f is the size of the largest matching in H. The following result follows directly from [36, Theorem 1.6].

Theorem 2.4.3. *Fix $d \in \mathbb{N}$ and let $\gamma \in \mathrm{Rad}^n$. Suppose f is a degree-d polynomial of rank r. Then for any $\ell \in \mathbb{R}$,*

$$\mathbb{P}\left[f(\gamma) = \ell\right] \leq \frac{\log^{O(1)}(r)}{\sqrt{r}}.$$

The above theorem will be useful in proving Theorem 2.1.1 (and it is easy to see how it alludes to the form of the inequality we see there). However, as discussed before, in the 3-graph case, we will need to able to account for low degree coefficients as well. The trouble with Theorem 2.4.3 is that it only looks at maximum degree coefficients. Hence, we need the following result to prove Theorem 2.1.2.

Theorem 2.4.4. *Fix $d \in \mathbb{N}$ and let $\gamma \in \mathrm{Rad}^n$. Consider a degree-$d$ polynomial*

$$f(\boldsymbol{x}) = \sum_{S \in \binom{[n]}{\leq d}} f_S \boldsymbol{x}^S.$$

Let $m_d = \max\{|f_S| : |S| = d\}$ be the maximum modulus degree-d coefficient. Let H' be the $(d-1)$-uniform hypergraph with vertex set $[n]$ and edge set $\{S : |S| = d-1, |f_S| \geq rm_d\}$. If H' has a matching of size r, then for any $\ell \in \mathbb{R}$,

$$\mathbb{P}\left[f(\gamma) = \ell\right] \leq \frac{\log^{O(1)}(r)}{\sqrt{r}}.$$

Finally, we state the exponential concentration inequality we will use to concentrate $X_{G,k}$ around its expected value to give bounds on $e(G)$, which will come in handy when showing that $f_\sigma(\gamma)$ has many nonzero coefficients of maximum degree (see the discussion in Section 2.3).

Lemma 2.4.5. *Consider a function $f : \{0,1\}^n \to \mathbb{R}$ such that*

$$|f(x_1, \ldots, x_{i-1}, 0, x_{i+1}, x_n) - f(x_1, \ldots, x_{i-1}, 1, x_{i+1}, x_n)| \leq c_i$$

for all $\boldsymbol{x} = (x_1, \ldots, x_n) \in \{0,1\}^n$ and all $i \in [n]$. Let $\boldsymbol{\xi} \in \mathrm{BL}(n,k)$. Then

$$\mathbb{P}\left[f(\boldsymbol{\xi}) - \mathbb{E}f(\boldsymbol{\xi}) \geq t\right] \leq \exp\left(-\frac{t^2}{8\sum_{i=1}^{n} c_i^2}\right).$$

Proof Idea. We consider the Doob martingale $Z_i = \mathbb{E}\left[\mathbb{E}f(\boldsymbol{\xi}) \mid \xi_1, \ldots, \xi_n\right]$. We argue that $|Z_i - Z_{i-1}| \leq 2c_i$, and then the result follows from the Azuma-Hoeffding inequality. □

2.5 Proof of Theorem 2.1.1

In this section, we provide a description of how the results from Section 2.4 are used to prove Theorem 2.1.1. Again, we will focus more on the chief ideas and approaches instead of replicating the details of the proofs in [30].

Take a graph G on the vertex set $[n]$ where $n = 2k$. Observe that it suffices to prove that $\mathbb{P}\left[X_{G,k} = \ell\right] \leq \log^{O(1)}(\ell^*/k)/\sqrt{\ell^*/k}$ for all such graphs G, as $I(n,k,\ell)$ monotonically decreases with n. We may assume that $e(G) \leq \binom{n}{2}/2$ by taking the complement of G if necessary. Set $X = X_{G,k}$. As usual, if the adjacency matrix of G is $A_G = (a_{xy})_{x,y}$ and $\boldsymbol{\xi} \in \mathrm{BL}(n,k) = \mathrm{BL}(n, n/2)$, we can express $X = X(\boldsymbol{\xi})$ using Equation 2.1.

Note that $\mathbb{E}X = e(G)(k(k-1)/n(n-1)) \approx e(G)/4$. We first use Lemma 2.4.5 to show that if $e(G)$ is not of the same order as ℓ, then $\mathbb{P}\left[X = \ell\right]$ must be quite small.

Proposition 2.5.1. *For any constant $\varepsilon > 0$, if $\ell \geq (1+\varepsilon)\mathbb{E}X$ or $\ell \leq (1-\varepsilon)\mathbb{E}X$, then*

$$\mathbb{P}\left[X = \ell\right] \leq \exp\left(-\Omega\left(\frac{\varepsilon^2 \ell}{k}\right)\right).$$

Proof Sketch. Observe that X satisfies the conditions required to apply Lemma 2.4.5 with $c_x = d_G(x) = d(x)$ for all $x \in [n]$. Maximise $\sum c_x^2 = \sum d(x)^2$ by setting $d(x) = n$ for as many

x as possible, namely for $2e(G)/n$ vertices x (as $\sum d(x) = 2e(G)$). Using that $\mathbb{E}X = O(e(G))$ and applying Lemma 2.4.5 will complete the proof. $\square$

This allows us to assume from now on that $e(G) = \Omega(\ell)$ and also $\ell = \Theta(\ell^*)$. Now, let σ be a uniformly random permutation of $[n]$ and condition on its outcome. Let H be the graph on the vertex set $[k]$ with $ij \in E(H)$ if and only if

$$a_{\sigma(i)\sigma(j)} - a_{\sigma(i+k)\sigma(j)} - a_{\sigma(i)\sigma(j+k)} + a_{\sigma(i+k)\sigma(j+k)} \neq 0$$

From Lemma 2.4.2 and the discussion immediately after, we know that X is a quadratic polynomial in $\gamma \in \mathrm{Rad}^{n/2}$ and that the rank of X is precisely the size of the maximum matching in H (see Equation 2.3). Hence, we now wish to show that H has a matching of size $\Omega(\ell/k)$ with high probability. Then substituting this in Theorem 2.4.3 will complete the proof of Theorem 2.1.1.

The crucial observation we use is that if four vertices in G, say $\{x, y, x', y'\}$, form an alternating 3-path (that is, the sequence xy, yx', $x'y'$ alternates between edges and nonedges), then $a_{xy} = a_{x'y'} \neq a_{x'y}$, and hence $a_{xy} - a_{x'y} - a_{xy'} + a_{x'y'} \neq 0$. Hence, edges in H arise from alternating 3-paths in G. So our aim is to show that G has a lot of such alternating 3-paths as this will imply that H has a large matching. To do this, we use a random greedy algorithm. Set $U := \{v \in V(G) : d(v) \geq 0.9n\}$ to be the set of all high-degree vertices in G. We will consider two cases.

First, suppose that at least half the edges of G meet U. This means that $2k|U| \geq e(G)/2 = \Omega(\ell)$, and hence $|U| \geq \Omega(\ell/k)$ is large. We iteratively build a matching M in H as follows. Pick any two vertices $u, w \in U$ that have not yet been revealed/picked, and reveal $i = \sigma^{-1}(u)$ and $j = \sigma^{-1}(w)$. If $i, j \leq k$ and $\sigma(i + k)$ and $\sigma(j + k)$ have not yet been revealed, reveal them. If we now find that $ij \in E(H)$, then we add it to M. We then argue that this procedure can run for $\Omega(\ell/k)$ steps, with probability $\Omega(1)$ of adding an edge to M each time, providing a matching of size $\Omega(\ell/k)$ in H with large probability. The intuition is that since $d(u) = |N(u)|$ and $d(w) = |N(w)|$ are large, the number of unrevealed pairs $(u', w') \in N(w) \times N(u)$ is larger than $\binom{n}{2}/2 \geq e(G)$. Hence, many such pairs $u'w'$ constitute a nonedge, which means $uw'u'w$ is an alternating 3-path, so we get an edge in H.

Next, suppose at least $e(G)/2 = \Omega(\ell)$ edges are induced by $\overline{U} = [n] \setminus U$. Then, we can greedily find a matching S of size $\Omega(\ell/n) = \Omega(\ell/k)$ in $\overline{U}$. Again, we iteratively build a matching M in H. Pick an edge $uw \in S$ such that u and w have not yet been revealed, and reveal $i = \sigma^{-1}(u)$ and $j = \sigma^{-1}(w)$. If $i, j \le k$ and $\sigma(i+k)$ and $\sigma(j+k)$ have not yet been revealed, reveal them. If we now find that $ij \in E(H)$, then we add it to M. Yet again, we wish to argue that this procedure can run for $\Omega(\ell/k)$ steps, with probability $\Omega(1)$ of adding an edge to M each time. The idea is that since $d(u) = |N(u)|$ and $d(w) = |N(w)|$ are small, there are many unrevealed pairs $(u', w') \in V(G) \setminus N(u) \times V(G) \setminus N(w)$, which means $u'uww'$ forms an alternating 3-path, leading to an edge in H.

2.6 Proof of Theorem 2.1.2

The overall theme of the proof of Theorem 2.1.2 is similar to that of Theorem 2.1.1, except that certain particularly difficult cases need some extra care.

Let G be a 3-graph on the vertex set $[n]$, with $n = 2k$. As in Section 2.5, it is enough to prove that $\mathbb{P}[X_{G,k} = \ell] \le \log^{O(1)}(k)/\sqrt{k}$ for all such 3-graphs G. Also, with the same arguments as in Proposition 2.5.1, we may assume that $\min\left\{e(G), \binom{n}{3}\right\} = \Omega(n^3)$. As per usual, we may write $X = X_{G,k}$ as

$$X = \sum_{1 \le x < y < z \le n} a_{xyz} \xi_x \xi_y \xi_z,$$

where $a_{xyz} = 1$ if $xyz \in E(G)$ and $a_{xyz} = 0$ otherwise, and $\xi \in \mathrm{BL}(n,k) = \mathrm{BL}(n, n/2)$. Given this setup, it would appear that a 3-graph generalisation of the proof of Theorem 2.1.1 could be plausible to achieve. Indeed, it would suffice to show that $V(G)$ has $\Omega(n^6)$ "good" 6-tuples of vertices (x, x', y, y', z, z') such that

$$a_{xyz} - a_{x'yz} - a_{xy'z} - a_{xyz'} + a_{x'y'z} + a_{x'yz'} + a_{xy'z'} - a_{x'y'z'} \neq 0. \tag{2.5}$$

However, in the 3-graph setting, there exist examples of G such that $\min\left\{e(G), \binom{n}{3}\right\} = \Omega(n^3)$, but $V(G)$ does not contain any such good 6-tuples. We instead proceed by structurally characterising all such 3-graph G, and handle cases where there are $o(n^6)$ good 6-tuples separately.

Fix a set of six vertices x, x', y, y', z, z'. Let $\mathcal{F}$ be the set of 3-graphs on this vertex set such that the expression on the left-hand side of Equation 2.5 is nonzero. A 3-graph G is said to be $\mathcal{F}$-*free* if it contains no induced subgraph from $\mathcal{F}$. Clearly, a 3-graph is $\mathcal{F}$-free if and only if its complement is as well. Now, we consider the following family of $\mathcal{F}$-free 3-graphs. For two disjoint vertex sets A and B and a set of disjoint pairs $M \subseteq A \times B$ (imagine M as a matching in a bipartite graph with parts A and B), let $G_{A,B,M}$ be the 3-graph with vertex set $A \sqcup B$, whose edges consist of all triplets of vertices that intersect both A and B, except those triplets that contain a pair from M. It is not hard to see that $G_{A,B,M}$ is $\mathcal{F}$-free – for any 6-tuple of vertices, the expression of the left-hand side of Equation 2.5 will be zero. We claim that any $\mathcal{F}$-free 3-graph with $\min\left\{e(G), \binom{n}{3}\right\} = \Omega(n^3)$ must have this structure.

Lemma 2.6.1. *Suppose G is an $\mathcal{F}$-free 3-graph on n vertices such that $\min\left\{e(G), \binom{n}{3}\right\} = \Omega(n^3)$. For sufficiently large n, either G or its complement must be of the form $G_{A,B,M}$ for some partition $A \sqcup B$ of $V(G)$ and some set of disjoint pairs $M \subseteq A \times B$.*

The proof of Lemma 2.6.1 is highly complicated and requires tedious case-by-case analysis; hence, we omit it. The key is to first use a theorem of Fox and Sudakov on unavoidable patterns in large hypergraphs with bounded density [17, Theorem 18] to argue that either G or its complement contains a copy of the complete bipartite 3-graph $K_{5,5}^{(3)}$, and then induct on $|V(G)|$ to show that G or its complement must be of the form $G_{A,B,M}$.

Now we need to see how to use Lemma 2.6.1 to complete the proof of Theorem 2.1.2. Under the coupling Fact 2.4.1, we know that X is a function of a uniformly random permutation of $[n]$, say σ, and the i.i.d. sequence $\gamma \in \mathrm{Rad}^{n/2}$. Using Lemma 2.4.2, we know that if we condition on any outcome of σ, then X is a polynomial in $\{\gamma_1, \ldots, \gamma_k\}$ of degree at most three. The coefficient g_{ijq} of $\gamma_i\gamma_j\gamma_q$ is

$$\sum_{\boldsymbol{b} \in \{0,1\}^3} (-1)^{|\boldsymbol{b}|} a_{\sigma(i+kb_1)\sigma(j+kb_2)\sigma(q+kb_3)}$$

(note that $|g_{ijq}| \leq 4$). The coefficient g_{ij} of $\gamma_i\gamma_j$ is

$$d_G(\sigma(i)\sigma(j)) - d_G(\sigma(i+k)\sigma(j)) - d_G(\sigma(i)\sigma(j+k)) + d_G(\sigma(i+k)\sigma(j+k)).$$

Let H be the (random) 3-graph on the vertex set $[n/2] = [k]$ with an edge ijq whenever $g_{ijq} \neq 0$. First, suppose that G contains $\Omega(n^6)$ induced subgraphs from $\mathcal{F}$. Consequently, we argue that $\mathbb{E}e(H) = \Theta(n^3)$. Using a McDiarmid-type concentration inequality for random permutation σ (see [35, Section 3.2]), we get

$$\mathbb{P}\left[e(H) \leq \mathbb{E}e(H)/2\right] \leq e^{-\Omega(n)}.$$

So with high probability, $e(H) \geq \mathbb{E}e(H)/2 = \Omega(n^3)$. Hence, we can greedily find a matching of size $\Omega(n)$ in H, and then Theorem 2.4.3 provides the required result.

Now it remains to consider what happens if G has $o(n^6)$ induced subgraphs from $\mathcal{F}$. Then, using the induced hypergraph removal lemma (see [45, Theorem 6]), we can add and remove at $o(n^3)$ edges of G to obtain an $\mathcal{F}$-free 3-graph. By Lemma 2.6.1, we may assume that it is of the form $G_{A,B,M}$. As $\min\left\{e(G), \binom{n}{3}\right\} = \Omega(n^3)$, and since we've added and removed at most $o(n^3)$ edges, we deduce that $|A|, |B| \geq \Omega(n)$. Also note that $O(|M|n) = o(n^3)$ triplets across A and B can involve a pair from M. Hence, by looking at the value of

$$d_{G'}(xy) - d_{G'}(x'y) - d_{G'}(xy') + d_{G'}(x'y'),$$

for $G' = G_{A,B,\emptyset}$ and $x, y' \in A$ and $x', y \in B$, and using the fact that very few edges across A and B can contain a pair from M, we conclude that there $\Omega(n^4)$ choices for a 4-tuple (x, x', y, y') such that

$$d_G(xy) - d_G(x'y) - d_G(xy') + d_G(x'y') \geq n/2.$$

Now, set H' to be the (random) graph with vertex set $[n/2] = [k]$ with $ij \in E(H')$ if and only if $g_{ij} \geq n/2$. Similar to the previous case, we argue that $\mathbb{E}e(H') = \Theta(n^2)$ and use a concentration inequality to say that $e(H') = \Omega(n^2)$ with probability $1 - e^{-\Omega(n)}$, and hence H' must have a matching of size $\Omega(n)$. The desired result then follows from Theorem 2.4.4 with $d = 3$ and $r = \min\{m, n/8\}$.

2.7 Attempts to Generalise to r-graphs

We first attempted to extend the methods of Section 2.5 to obtain an analogous result for uniformities 3 and higher. As noted, the key observation we used was that alternating 3-paths in the graph G led to edges in the auxiliary graph H, which provided us with a large matching in H. The major difficulty in extending to 3-uniform hypergraphs is that when

X is expressed as a polynomial in $\boldsymbol{\gamma} \in \mathrm{Rad}^{n/2}$, the leading coefficients are now given by Equation 2.4 instead of Equation 2.3. The increase in complexity in this expression makes it difficult to find a simple substructure in 3-graphs that leads to edges in H, unlike alternating 3-paths in the graph scenario. So, precisely what subgraphs we would need to search for in the 3-graph setting is unclear.

We then attempted to generalise the proof of Theorem 2.1.2 in Section 2.6 to dense r-graph for $r \geq 4$ (i.e., with $\min\{e(G), \binom{n}{r} - e(G)\} = \Omega(n^r)$). As discussed, the principal result used in Section 2.6 is a structural characterisation of 3-uniform hypergraphs for which $X = X(\boldsymbol{\gamma})$ has no nonzero degree-3 coefficients. However, it is not clear whether an analogous characterisation is even true in the 4-uniform setting. Additionally, the proof of the characterisation for 3-graphs strongly uses Kőnig's theorem, whose generalisation to higher uniformities is Ryser's conjecture, which has stood open for over 50 years now. Hence, this approach likely cannot be extended to large uniformities at present.

Chapter 3

The Sparse Regime

3.1 Main Results

We now move on to the results of Martinsson *et al.* [34], who proved the following:

Theorem 3.1.1. *For every* ℓ *such that* $1 \leq \ell \leq o_k(k^{6/5})$*, we have*

$$\mathrm{ind}(k, \ell) \leq 1/e + o_k(1).$$

The paper first provides a short proof for the $\ell \leq o_k(k)$ case and then proceeds to handle the remaining part separately. We have been able to extend the former result to the r-graph setting in general.

Theorem 3.1.2. *For every uniformity* $r \geq 2$ *and for all sufficiently large* $k, \ell \in \mathbb{N}$ *such that* $1 \leq \ell \leq o_k(k)$*, we have*

$$\mathrm{ind}_r(k, \ell) \leq 1/e + o_k(1).$$

The proof of the general $o_k(k^{6/5})$ case is significantly more involved. Fortunately, we have successfully generalised some of its key results to the general r-uniform framework. In the following sections, we first present the proof of Theorem 3.1.2. We then discuss a broad outline of the approach used in [34] for Theorem 3.1.1, detail what we have been able to prove for r-graphs, and what would be needed to complete this proof.

3.2 Notation

In this chapter, we shall always interpret asymptotic statements in the limit as $k \to \infty$ (unless otherwise stated) and will drop the corresponding subscript in asymptotic notation. All results shall be proved only in this asymptotic limit, and henceforth, we shall forgo the preface that k is sufficiently large. We shall also assume $n = n(k)$ is always as large as desired to support our arguments as necessary.

We borrow notation and terminology from [34]. We say that an event holds *with high probability* (w.h.p. for short) if the probability of its occurrence approaches 1 as $k \to \infty$. For any two events $\mathcal{E} = \mathcal{E}(k)$ and $\mathcal{F} = \mathcal{F}(k)$ (which could also potentially depend on ℓ, n and G), we say that $\mathcal{E}$ is *essentially contained* in $\mathcal{F}$, and write $\mathcal{E} \lesssim \mathcal{F}$, if

$$\mathbb{P}\left(\mathcal{E} \setminus \mathcal{F}\right) = o(1).$$

Note that $\mathcal{E} \lesssim \mathcal{F}$ is equivalent to $\mathcal{E} \lesssim \mathcal{E} \cap \mathcal{F}$.

3.3 Proof of Theorem 3.1.2

We first present a simple, self-contained proof of the Edge-Statistics conjecture generalised to higher uniformities when we restrict ourselves to the regime $\ell = o(k)$. This is a straightforward extension of the proof of the same statement for graphs as presented in [34, Proposition 2.1].

Proof of Theorem 3.1.2. Fix any uniformity $r \geq 2$. Choose k and ℓ as in the theorem statement and assume $n = n(k)$ is sufficiently large. Let $G = (V, E)$ be an r-uniform hypergraph on n vertices and let $\mathbf{v} = (v_1, v_2, \ldots)$ be an infinite sequence of vertices of G picked uniformly at random from $V^{\mathbb{N}}$. We sequentially colour the vertices in $\mathbf{v}$ with two colours as follows:

(i) Colour v_1 black;

(ii) colour v_i green if and only if the subgraph of G induced by v_i and all the black vertices v_j with $j < i$ contains at least ℓ edges; otherwise colour v_i black.

Set

$$L = L(\mathbf{v}) := \min\{i \geq 1 : \text{there are } k-1 \text{ black vertices among } v_1, \ldots, v_i\}$$

and $L := \infty$ if there fewer than $k-1$ black vertices in $\mathbf{v}$. Define the random variable $Y_{G,k} = Y_{G,k}(\mathbf{v})$ to be the number of green vertices in the set $\{v_i : 1 \leq i < L\}$.

Henceforth, the overall structure of the proof is as follows. As usual, set $X = X_{G,k}$ to be the random variable corresponding to the number of edges induced by a uniformly random k-subset of $V(G)$. First, we show $\mathbb{P}[X = \ell]$ is bounded above by $\mathbb{P}[Y_{G,k} = 1]$ plus a small error term. The idea behind this is that since n is so much larger than k, the first k vertices of $\mathbf{v}$ are likely to be distinct and hence constitute a uniform random element of $V^{(k)}$. Then, since k is much larger than ℓ, we argue that the probability that the first k vertices of $\mathbf{v}$ inducing ℓ edges cannot be too much bigger than the probability that the first $k-1$ vertices of $\mathbf{v}$ induce ℓ edges as well (i.e., randomly changing v_{k-1} isn't very likely to change the fact that $e(G[v_1, \ldots, v_k]) = \ell$). We consequently deduce that this would imply that $Y_{G,k}$ is 1. Finally, we observe that upon conditioning on the first $k-1$ black vertices of $\mathbf{v}$, the distribution of $Y_{G,k}$ can be expressed as a sum of independent geometric random variables, and then show that the maximum possible value of this sum (as a function of the parameters of the geometric random variables) is $1/e$.

Let us formalise this. We begin by first showing that

$$\mathbb{P}[X = \ell] \leq \mathbb{P}[Y_{G,k} = 1] + o(1). \tag{3.1}$$

Let $\tilde{X}_k = e(G[v_1, \ldots, v_k])$, and let $\mathcal{A}$ be the event that $v_1, \ldots, v_k$ are all distinct. Assuming that n is sufficiently large as a function of k (i.e., $n \geq \omega(k^2)$), we have $\mathbb{P}[\mathcal{A}] = 1 - n \cdot \Theta(\frac{k^2}{n^2}) = 1 - o(1)$. Thus, we see that

$$\mathbb{P}[X = \ell] = \mathbb{P}[\tilde{X}_k = \ell \mid \mathcal{A}] \leq \frac{\mathbb{P}[\tilde{X}_k = \ell]}{\mathbb{P}[\mathcal{A}]} \leq \mathbb{P}[\tilde{X}_k = \ell] + o(1). \tag{3.2}$$

Next, we note that if $e(G[v_1, \ldots, v_k] = \ell)$ and v_{k-1} is an isolated vertex in $G[v_1, \ldots, v_k]$, then that would imply that $e(G[v_1, \ldots, v_{k-1}] = \ell)$ as well. Since ℓ edges can span at most $r\ell$ vertices in G, and since every possible ordering of the k vertices that constitute $v_1, \ldots, v_k$ is equally likely, it follows by symmetry that

$$\mathbb{P}[\tilde{X}_k = \tilde{X}_{k-1} = \ell] \geq \mathbb{P}[\tilde{X}_k = \ell] \cdot \frac{k - r\ell}{k} \geq \mathbb{P}[\tilde{X}_k = \ell] - o(1), \tag{3.3}$$

where the final inequality uses $\ell = o(k)$ (this is the only place where this assumption is used).

Now, we observe that $\tilde{X}_k = \tilde{X}_{k-1} = \ell$ implies that $Y_{G,k} = 1$. Indeed, $\tilde{X}_k = \tilde{X}_{k-1} \geq \ell$ means that at least one green vertex appears in $\{v_1, \dots, v_{k-1}\}$. And if there is more than one green vertex in $\{v_1, \dots, v_k\}$, then $\tilde{X}_k > \ell$. So this means that $L = k$ and $Y_{G,k} = 1$. This, along with 3.2 and 3.3, implies that $\mathbb{P}[X = \ell] \leq \mathbb{P}[Y_{G,k} = 1] + o(1)$ as desired. Thus, it suffices now to show that $\mathbb{P}[Y_{G,k} = 1] \leq 1/e$.

Let $\mathbf{u} = (u_1, \dots, u_{k-1})$ be a sequence of $k-1$ (not necessarily distinct) vertices of G. Let $U(\mathbf{u})$ be the event that $u_1, \dots, u_{k-1}$ are the first $k-1$ black vertices in $\mathbf{v}$. Define $A_i = \{v \in V : e(G[\{u_1, \dots, u_i, v\}]) \geq \ell\}$ and $p_i = \frac{1}{n}|A_i|$ for all $1 \leq i \leq k-2$. Now, observe that if $\mathbb{P}[U(\mathbf{u})]$ is nonzero, then the conditional distribution of $Y_{G,k}$ given $U(\mathbf{u})$ is given by the sum

$$\mathrm{Geom}(p_1) + \mathrm{Geom}(p_2) + \cdots + \mathrm{Geom}(p_{k-2})$$

of independent geometric distributions. Indeed, suppose that we have chosen the first t vertices of $\mathbf{v}$, say $v_1, \dots, v_t = u_i$, up to u_i, so we have selected up to the i^{th} black vertex. We proceed to pick the vertices of $\mathbf{v}$ sequentially. The probability that v_{t+1} (the next vertex in the sequence $\mathbf{v}$) is in A_i is exactly equal to p_i, and in this case, v_{t+1} would be coloured green by the definition of A_i. Furthermore, as long as we continue to pick vertices only from A_i, they are always going to be green, and the probability of picking a vertex from A_i remains p_i each time. As we have conditioned on $U(\mathbf{u})$, the first vertex we pick outside A_i must be the next black vertex u_{i+1}, and the probability of this event is $1 - p_i$. Since p_i is a function of $|A_i|$, which depends only on the black vertices chosen up to u_i, the number of green vertices picked after u_i until u_{i+1} is independent of the number of green vertices between any two consecutive black vertices of $\mathbf{u}$ up to u_i. So for any $j \geq 0$, the probability of picking j green vertices between u_i and u_{i+1} is $p_i^j(1 - p_i)$.

Thus, it follows that

$$\mathbb{P}[Y_{G,k} = 1 \mid U(\mathbf{u})] = \sum_{i=1}^{k-2} p_i \prod_{j=1}^{k-2} (1 - p_j) \leq \sum_{i=1}^{k-2} p_i e^{-\sum_{j=1}^{k-2} p_j} \leq 1/e,$$

where we use that $1 - x \leq e^{-x}$ for all real x and that $f(x) = xe^{-x}$ is maximised at $x = 1$. Since this is true for every relevant choice of $\mathbf{u}$ (otherwise $\mathbb{P}[Y_{G,k} = 1] = \mathbb{P}[Y_{G,k} = 1 \mid$

$U(\mathbf{u})] \cdot \mathbb{P}[\mathbb{U}(\mathbf{u})] = 0)$, we also have $\mathbb{P}[Y_{G,k} = 1] \leq 1/e$ unconditionally, which, along with 3.1, completes the proof. □

3.4 Probabilistic Results

We first state a purely probabilistic result that will be utilised in the proof of Theorem 3.1.1. This appears as [34, Lemma 3.1]. It tells us when we can approximate a hypergeometric random variable with a binomial distribution and how it converges in probability to a Poisson distribution. It also states that if the variance of certain hypergeometric random variables goes to infinity, then the probability of taking any single value goes to zero.

Lemma 3.4.1. *Let X be a hypergeometric random variable counting the number of successes obtained when sampling m elements without replacement from a population of size N containing Np successes. Assume $m^2/N \to 0$ and $m \to \infty$. If $mp \to \lambda < \infty$, then*

$$\max_i \left| \Pr[X = i] - \frac{\lambda^i e^{-\lambda}}{i!} \right| \to 0,$$

where the maximum is taken over all nonnegative integers. On the other hand, if $mp(1-p) \to \infty$, then $\max_i \Pr[X = i] \to 0$.

3.5 Proof of Theorem 3.1.1

Before we get into the details of the proof, we first introduce some definitions. We will let G be an r-graph on n vertices for some uniformity $r \geq 2$ and will specify when we consider G to be just a graph ($r = 2$). Assume $1 \leq \ell \leq o(k^{6/5})$. Let A denote a uniformly random k-vertex subset from $V(G)^{(k)}$. Throughout our discussion, we will use $\mathcal{E}$ to denote the event that $e(G[A]) = X = \ell$. Let $(w_k)_{k\geq 1}$ be a sequence of positive real number that diverges to infinity. We will always assume that the rate of divergence of this sequence is slow enough for all our arguments to hold. At the very least, we will require $w_k\sqrt{\ell} = o(k)$. Now, for every integer $d \geq 0$, define the event

$$\mathcal{D}_d := \{\text{all but at most } w_k\sqrt{\ell} \text{ vertices in } A \text{ have degree } d \text{ in } G[A]\}.$$

With this setup in mind, we discuss the broad outline of the proof. Observe that to prove Theorem 3.1.1, it suffices to show that $\mathcal{E}$ is essentially contained in some event $\mathcal{F}$ of probability at most $1/e + o(1)$. Our objective is to show that there exists some deterministic $d = d(G, k, \ell)$ such that $\mathcal{E} \cap \mathcal{D}_d$ is the desired event $\mathcal{F}$. There are three major components to this.

First, we show that $\mathbb{P}[\mathcal{E} \cap \mathcal{D}_d] \leq 1/e + o(1)$ for all integers $d \geq 0$, which means it now suffices to show that $\mathcal{E} \lessapprox \mathcal{E} \cap \mathcal{D}_d$ for some d. Next, we argue that if A induces exactly ℓ edges, it is very likely that almost all the k vertices have the same degree in the induced subgraph $G[A]$. The final part of the proof is dedicated to proving concentration-type results on this most common degree to show that it can take at most one deterministic value depending on G, k and ℓ, and does not depend on A. The latter two parts together will allow us to conclude that $\mathcal{E} \lessapprox \mathcal{E} \cap \mathcal{D}_d$ for some fixed d.

We have successfully extended the first two parts of the proof to r-graphs in general and will present these proofs in detail. We have not yet found a way to generalise the final part of the proof regarding concentration bounds on the most common degree. Hence, for the last part, we will only sketch the idea behind the proof used in the graph case [34]. We conclude with our attempts to concentrate the random variable corresponding to the most common degree in $G[A]$ and explanations for the limitations of these approaches.

3.5.1 Bounding the Probability of $\mathcal{E} \cap \mathcal{D}_d$

Proposition 3.5.1. *For all integers $d \geq 0$, we have*

$$\mathbb{P}\left[\mathcal{E} \cap \mathcal{D}_d\right] \leq 1/e + o(1).$$

We do not present the proof for $r = 2$ (for graphs), as those details are [34, Claim 3.2], and will only present our proof for when $r \geq 3$. We will quickly describe the proof for graphs afterwards.

We first sketch the rationale behind the proof. When $d = 0$, we proceed in a very similar fashion to Theorem 3.1.2. We use the same ideas to show that if most vertices of A are isolated in $G[A]$ (which is what D_0 implies), then if A induces ℓ edges, A minus one randomly picked vertex is likely to as well. For the case where $d \geq 1$, we first argue the distribution of $d_A(v)$ for any vertex $v \in A$ is the same as that of the random variable corresponding to the number of $(r-1)$-edges present in a uniformly random $(k-1)$-vertex subset of an $(r-1)$-graph with $(n-1)$ vertices and $d_G(v)$ edges. Then we show that $d \leq o(k)$, and consequently Theorem 3.1.2 will provide the desired bound.

Proof of Proposition 3.5.1. Assume G is an r-graph on n vertices for some $r \geq 3$. Suppose first that $d = 0$. Consider the same process described in Theorem 3.1.2, and define $Y_{G,k}$, $\tilde{X}_k$ and $\mathcal{A}$ in the same way. If we can show that

$$\mathbb{P}\left[\mathcal{E} \cap \mathcal{D}_0\right] \leq \mathbb{P}\left[Y_{G,k} = 1\right] + o(1),$$

then proceeding in the same manner as in Theorem 3.1.2 to show that $\mathbb{P}\left[Y_{G,k} = 1\right] \leq 1/e$ yields the desired result. Define $\tilde{\mathcal{D}_0}$ be the event that all but most $w_k\sqrt{\ell}$ of the vertices $v_1, \ldots, v_k$ are isolated in $G\left[\{v_1, \ldots, v_k\}\right]$. We thus have

$$\begin{aligned}
\mathbb{P}[\mathcal{E} \cap \mathcal{D}_0] &= \mathbb{P}[\tilde{\mathcal{D}_0} \cap \tilde{X}_k = \ell \mid \mathcal{A}] \\
&\leq \mathbb{P}[\tilde{\mathcal{D}_0} \cap \tilde{X}_k = \ell]/\mathbb{P}[\mathcal{A}] \\
&\leq \mathbb{P}[\tilde{\mathcal{D}_0} \cap \tilde{X}_k = \ell] + o(1) \text{ (where again we use that } \mathbb{P}[\mathcal{A}] = 1 - o(1)).
\end{aligned}$$

Like last time, since every permutation of $v_1, \ldots, v_k$ is equally likely, we further see that

$$\mathbb{P}[\tilde{X}_k = \tilde{X}_{k-1} = \ell] \geq \mathbb{P}[\tilde{\mathcal{D}_0} \cap \tilde{X}_k = \ell] - \frac{w_k\sqrt{\ell}}{k}.$$

Provided that w_k diverges slow enough, since $\sqrt{\ell} = o(k^{2/3}) = o(k)$, we will have $w_k\sqrt{\ell}/k = o(1)$. And as we have seen, $\tilde{X}_k = \tilde{X}_{k-1} = \ell$ deterministically implies that $Y_{G,k} = 1$, we obtain the desired bound on $\mathbb{P}[\mathcal{E} \cap \mathcal{D}_0]$.

Next, suppose that $d \geq 1$. Let v be a vertex chosen uniformly random from A. Then

$$\begin{aligned}
\mathbb{P}[d_A(v) = d \mid \mathcal{D}_d] &= 1 - \frac{w_k\sqrt{\ell}}{k} \\
&= 1 - o(1),
\end{aligned}$$

assuming w_k goes to infinity slowly enough. Thus, we have

$$\begin{aligned}\mathbb{P}[\mathcal{E} \cap \mathcal{D}_d] &\leq \mathbb{P}[\mathcal{D}_d] \\ &= \mathbb{P}[\mathcal{D}_d \cap d_A(v) = d]/\mathbb{P}[d_A(v) = d \mid \mathcal{D}_d] \\ &\leq \mathbb{P}[\mathcal{D}_d \cap d_A(v) = d](1 + o(1)) \\ &\leq \mathbb{P}[\mathcal{D}_d \cap d_A(v) = d] + o(1) \\ &\leq \mathbb{P}[d_A(v) = d] + o(1)\end{aligned}$$

So, it now suffices to prove that

$$\mathbb{P}[d_A(v) = d] \leq 1/e + o(1)$$

By looking sum of the degrees of vertices in $G[A]$, observe that $\mathcal{E} \cap \mathcal{D}_d$ implies that $(k - w_k\sqrt{\ell})d \leq r\ell$, and hence that $\mathbb{P}[\mathcal{E} \cap \mathcal{D}_d] = 0$ unless $kd \leq (r+1)\ell$. And since $\ell = o(k^{6/5})$, we may assume $d = o(k^{1/5})$.

Next, we note that the pair (v, A) can equivalently be generated by first uniformly randomly picking a vertex v from $V(G)$ and then choosing the remaining $(k-1)$ vertices of A uniformly randomly from the set of all $(k-1)$-vertex subsets of $V(G) \setminus \{v\}$. Now we will fix the choice of $v \in A$ and study the distribution of the random variable $d_A(v)$. To this end, we construct an auxiliary $(r-1)$-graph H_v with $V(H_v) = V(G) \setminus \{v\}$ such that for any set of $r-1$ vertices $\{x_1, \ldots, x_{r-1}\} \in V(H_v)$, we have $x_1 \ldots x_{r-1} \in E(H_v)$ if and only if $vx_1 \ldots x_{r-1} \in E(G)$. Clearly $|V(H_v)| = n-1$ and $e(H_v) = d_G(v)$. Crucially, we now interpret $d_A(v)$ as exactly the number of edges induced by a uniformly random $(k-1)$-vertex set of an $(n-1)$-vertex $d_G(v)$-edge $(r-1)$-graph. Then, by Theorem 3.1.2, we know that

$$\mathbb{P}[d_A(v) = \ell] \leq 1/e + o(1)$$

for all $1 \leq \ell \leq o(k)$. As $1 \leq d \leq o(k^{1/5}) \leq o(k)$, the above inequality provides the required upper bound on $\mathbb{P}[d_A(v) = d]$, completing the proof. $\square$

Remark 3.5.1. For the case $r = 2$ (as proved in [34, Claim 3.2]), the proof for when $d = 0$ is the same. Suppose $d \geq 1$. Observe that the random variable $d_A(v)$ follows a hypergeometric distribution with population size $n-1$, sample size $k-1$, and $d_G(v)$ successes. As $d = o(k^{1/5})$, we argue that $d_G(v)$ can't be too large, because that would make $\mathbb{E}d_A(v)$ too large and hence $d_A(v)$ can be shown to be too large using Markov's inequality. Then, since the sample size is much smaller than the population size, we use Lemma 3.4.1 to say that the hypergeometric distribution can be approximated by a binomial distribution, which converges to a Poisson distribution in probability. Using that $d \geq 1$ and optimising the Poisson parameter yields the upper bound of $1/e + o(1)$.

3.5.2 Approximate Regularity in $G[A]$

Next, we wish to show that if $X = \ell$, then it is very likely that almost all vertices in $G[A]$ have the same degree. To begin with, we define the event

$$\begin{aligned} \mathcal{D}_* &= \bigcup_{d \geq 0} \mathcal{D}_d \\ &= \{\text{all but at most } w_k\sqrt{\ell} \text{ vertices in } A \text{ have the same degree in } G[A]\} \end{aligned}$$

We will show that $\mathcal{E}$ is essentially contained in this event.

Proposition 3.5.2. *Suppose $\ell = \omega(1)$. Then we have $\mathcal{E} \subsetneq \mathcal{D}_*$.*

Proof. Set $m = k/(w_k^{1/3}\sqrt{\ell})$. Since $k = \omega(\sqrt{\ell})$, provided that w_k increases slowly enough, we may assume that $m \geq w_k$. We generate A by first choosing a uniformly random $(k-m)$-vertex subset S of $V(G)$, and then a uniformly random m-vertex subset Q of $V(G) \setminus S$. In the context of this procedure, we show that $\mathcal{E}$ is essentially contained in the following four events:

- $\mathcal{E}_1 := \{$there are no edges with at least 2 vertices in Q$\}$,
- $\mathcal{E}_2 := \{e(S) + \sum_{v \in Q} e(v, S) = \ell\}$, where $e(v, S)$ is the number of edges in G containing v along with $r - 1$ vertices in S.
- $\mathcal{E}_3 := \{$all but at most $w_k^{1/3}$ vertices $v \in Q$ have the same value of $e(v, S)\}$,
- $\mathcal{E}_4 := \{$all but at most $w_k^{1/3}$ vertices $v \in Q$ have the same degree in $G[A]\}$.

The idea behind the first one (which will directly imply the second) is that since $|Q| = m$ is much smaller than k, and since, under $\mathcal{E}$, we know that $G[A]$ contains only $\ell = o(k^{6/5})$

edges, most of them have at most one vertex in Q. To get $\mathcal{E}_3$ (and consequently $\mathcal{E}_4$), we look at the statistics of $e(v, S)$ over all $v \notin S$. In particular, if this quantity varies too much with v, then we show that the probability of $e(A)$ remaining exactly ℓ when one vertex in $A \setminus S$ is randomly replaced by another is small. Finally, to get $\mathcal{D}_*$, we will use the fact that while $|Q| = m = o(k)$, it is still large enough that the "randomness" in the degree distribution of vertices in A is reflected in that of vertices in Q.

We now begin by proving that $\mathcal{E} \precsim \mathcal{E}_1$. Let $e\left(Q_{\geq 2}\right)$ denote the number of edges with at least two vertices from Q. We generate Q by first choosing A and then picking a uniformly random m-vertex subset $Q \subseteq A$. Hence, we have

$$\mathbb{E}[e(Q_{\geq 2}) \mid X = \ell] \leq O\left(\ell \cdot \frac{\binom{m}{2}}{\binom{k}{2}}\right) = O\left(1/(w_k^{2/3})\right) = o(1),$$

where the second equality follows from the definition of m, and the last is a consequence of $w_k = \omega(1)$. Therefore, using Markov's inequality, we have

$$\begin{aligned}
\mathbb{P}[\mathcal{E} \setminus \mathcal{E}_1] &= \mathbb{P}[e(Q_{\geq 2}) \geq 1 \cap X = \ell] \\
&= \mathbb{P}[e(Q_{\geq 2}) \geq 1 \mid X = \ell] \cdot \mathbb{P}[X = \ell] \\
&\leq \mathbb{E}[e(Q_{\geq 2}) \mid X = \ell] \\
&= o(1),
\end{aligned}$$

and so $\mathcal{E} \precsim \mathcal{E}_1$.

Once we have obtained this, it follows directly from the definitions that $\mathcal{E} \precsim \mathcal{E} \cap \mathcal{E}_1 \subseteq \mathcal{E}_2$.

Next, we prove that $\mathcal{E}_2 \precsim \mathcal{E}_3$, which then implies that $\mathcal{E} \precsim \mathcal{E}_3$. Pick S as a uniformly random $(k-m)$-vertex subset of $V(G)$, and reveal it. Let d_{med} be the median value of $e(v, S)$ over all vertices $v \in V(G) \setminus S$.

CASE 1: All but at most $w_k^{1/4} n/m$ vertices $v \in V(G) \setminus S$ satisfy $e(v, S) = d_{\text{med}}$.

In this case, the expected number of vertices $v \in Q$ for which $e(v, S) \neq d_{\text{med}}$ is at most

$$|Q| \cdot \frac{w_k^{1/4} n/m}{|V(G) \setminus S|} = m \cdot \frac{w_k^{1/4} n/m}{n - k + m} = O(w_k^{1/4}) = o(w_k^{1/3}).$$

So, by Markov's inequality, we have

$$\mathbb{P}[\overline{\mathcal{E}_3}] \leq \frac{o(w_k^{1/3})}{w_k^{1/3}} = o(1),$$

and therefore $\mathcal{E}_2 \subsetneq \mathcal{E}_3$ (trivially).

CASE 2: At least $w_k^{1/4} n/m$ vertices $v \in V(G) \setminus S$ satisfy $e(v, S) \neq d_{\text{med}}$.

We will show that in this case, $\mathbb{P}[\mathcal{E}_2] = o(1)$, and so trivially $\mathcal{E}_2 \subsetneq \mathcal{E}_3$. We may suppose that at least $w_k^{1/4} n/(2m)$ vertices of $V(G) \setminus S$ satisfy $e(v, S) > d_{\text{med}}$ (the other case can be treated analogously). Call this set of vertices P, and denote $|P|$ by t, so $t \geq w_k^{1/4} n/(2m)$. Let P' be the complement of P in $V(G) \setminus S$. Set $N := |V(G) \setminus S| = n - k + m$, and hence we have $|P'| = N - t$. Now, we think of Q in the following manner. Let I be the random variable corresponding to the number of vertices $v \in Q$ picked from P' (and so the number of vertices $v \in Q$ that come from P is $m - I$). Then I is precisely a hypergeometric random variable with a population of size N containing $N - t$ successes such that we sample m elements from the population without replacement. Thus, Q consists of I uniformly randomly chosen vertices of P' and $m - I$ uniformly randomly chosen vertices of P. Formally, we let $v'_1, \ldots, v'_{N-t}$ be a random permutation of P' and $v''_1, \ldots, v''_t$ be a random permutation of P. Ultimately, we produce Q as

$$Q = \{v'_1, \ldots, v'_I, v''_1, \ldots, v''_{m-I}\}.$$

By this process, Q is nothing but a uniformly random m-vertex subset of $V(G) \setminus S$.

For $\mathcal{E}_2$ to hold, we require

$$\sum_{v \in Q} e(v, S) = \ell - e(S).$$

As the value of $e(v, S)$ for any vertex $v \in P$ is strictly larger than the corresponding value for any $v \in P'$, we conclude that for every fixed choice of the permutations of P and P', there is at most one value of I such that this equation holds. However, we know that I is a hypergeometric random variable with population size N, sample size m, and $N - t$ successes in the population. We see that $m \geq w_k = \omega(1)$, and

$$\frac{m^2}{N} \leq \frac{k^2}{n - k + m} = o(1)$$

assuming $n = \omega(k^2)$. Furthermore, $t \geq w_k^{1/4} n/(2m) = \omega(N/m)$ and $t \leq N/2$ (as d_{med} is a median by definition) imply

$$m\left(1-\frac{t}{N}\right)\left(\frac{t}{N}\right) \geq \left(\frac{m}{2}\right)\left(\omega\left(\frac{1}{m}\right)\right) = \omega(1).$$

So it follows from Lemma 3.4.1 that $\mathbb{P}[I = i] = o(1)$ for all $i \geq 0$. Thus, we have $\mathbb{P}[\mathcal{E}_3] = o(1)$ as claimed.

Finally, as $\mathcal{E} \lesssim \mathcal{E}_1$ and $\mathcal{E} \lesssim \mathcal{E}_3$, it follows directly from definitions that $\mathcal{E} \lesssim \mathcal{E}_1 \cap \mathcal{E}_3 \subseteq \mathcal{E}_4$.

Lastly, we show that $\mathcal{E}_4 \lesssim \mathcal{D}_*$, which finishes the proof. Suppose A is such that $\mathcal{D}_*$ does not occur. Condition on this event $\overline{\mathcal{D}_*}$ while leaving Q as a uniformly random m-vertex subset of A. Set d to be the median degree of a vertex in $G[A]$. The event $\overline{\mathcal{D}_*}$ implies that there exists a set of at least $w_k\sqrt{\ell}/2$ vertices in A such that either every vertex in this set has degree larger than d in $G[A]$ or smaller than d in $G[A]$. We may suppose that the former case holds as the latter may be handled similarly. Let t denote the size of this set, and so $t \geq w_k\sqrt{\ell}/2$. Let X_t be the random variable corresponding to the number of vertices of this set that get chosen in Q. Observe that X_t is a hypergeometric random variable with population size k and sample size m such that there are t successes in the population. Hence, we have

$$\mathbb{E}[X_t] = t \cdot \frac{m}{k} \geq \frac{w_k\sqrt{\ell}}{2} \cdot \frac{k}{w_k^{1/3}\sqrt{\ell}} \cdot \frac{1}{k} = \frac{w_k^{2/3}}{2} = \omega(1).$$

Further, as X_t is hypergeometric, we know that its standard deviation satisfies $\sigma(X_t) = O(\sqrt{tm/k}) = O(w_k^{1/3})$. Therefore, Chebyshev's inequality implies that w.h.p. $X_t > w_k^{1/3}$ must hold. Conversely, as $t \leq k/2$ (since d is defined to be a median), it also implies that w.h.p. $X_t \leq (1/2 + o(1))m$ (using the fact that $m \geq w_k = \omega(w_k^{1/3})$). Therefore, we may say w.h.p. that

$$w_k^{1/3} < X_t \leq (1/2 + o(1))m < m - w_k^{1/3},$$

where the final inequality again utilises $m = \omega(w_k^{1/3})$. So the number of vertices $v \in Q$ with $d_A(v) > d$ is strictly less than $m - w_k^{1/3}$, and the number of remaining vertices $v \in Q$ (with $d_A(v) \leq d$) is $m - X_t < m - w_k^{1/3}$ as well. These two inequalities imply that there is no set of $m - w_k^{1/3}$ vertices in Q with the same degree in A. Consequently,

$$\mathbb{P}[\mathcal{E}_4 \setminus \mathcal{D}_*] \leq \mathbb{P}[\mathcal{E}_4 \mid \overline{\mathcal{D}_*}] = o(1),$$

as desired. $\square$

3.5.3 Concentrating the Common Degree

Due to Proposition 3.5.2, we know that $\mathcal{E} \lesssim \mathcal{D}_*$, which means that almost all vertices in A have the same degree in $G[A]$, but the value of this most common degree could depend on A. As noted previously, we wish to use the control over the edge-statistics of A afforded by this knowledge of the degree distribution of $G[A]$ to conclude that the most frequent degree can take only one possible value $d(G, k, \ell)$, which does not depend on A.

In the r-graph setting, we have not been able to find the right method to concentrate the common degree. Hence, we will instead present the techniques and proofs used in [34] for graphs, and will henceforth assume $r = 2$ (so G is a graph). As with Chapter 2, we will not replicate the paper [34] but will focus on key ideas and provide sketches of proofs.

To begin with, we partition the vertices of G into two parts depending on their degree:

- $V_{\text{light}} = \{v \in V(G) : d_G(v) < \frac{n}{k}\ell^{1/3}\}$, the set of *light vertices*;
- $V_{\text{heavy}} = \{v \in V(G) : d_G(v) \geq \frac{n}{k}\ell^{1/3}\}$, the set of *heavy vertices*.

We first argue that there can't be too many heavy vertices.

Proposition 3.5.3. *Suppose $\ell = \omega(1)$ and G contains more than $5\ell^{2/3}n/k$ heavy vertices. Then $\mathbb{P}[\mathcal{E}] = o(1)$.*

Proof Sketch. Generate A by first picking a random set $A_1 \in V(G)$ of size $k/2$ and then another random set $A_2 \in V(G) \setminus A_1$ of size $k/2$, and finally set $A = A_1 \sqcup A_2$. The number of heavy vertices in A_1 is a hypergeometric random variable with population size n, sample size $k/2$, and with $|V_{\text{heavy}}|$ successes in the population. Consequently, Chernoff bounds for hypergeometric random variables (see [2, Appendix A] for a reference) allow us to say that $|A_1 \cap V_{\text{heavy}}|$ is concentrated around its expected value w.h.p., and this expected value will be large by our assumption. Then, for every vertex $v \in A_1 \cap V_{\text{heavy}}$, it is easy to see that $e(v, A_2)$ is hypergeometric as well (see Remark 3.5.1), and hence is concentrated around its expectation, which will be large as v is heavy. As Chernoff bounds provide an exponentially decaying tail bound, using a union bound, one can argue that all heavy vertices in A_1 have a large degree in A_2 simultaneously w.h.p.. We use that to show that w.h.p we have $X > \ell$. $\square$

Next, we bound the variance of the following random variable.

Proposition 3.5.4. *Assume that* $\ell = \omega(1)$*. Let* $Z := \sum_{v \in A \cap V_{\text{light}}} d_A(v)$*. Then either* $\mathbb{P}[\mathcal{E}] = o(1)$ *or* $\text{Var}[X - Z] \leq 30\ell^{5/3}$*.*

Proof Sketch. Set $H = e(A \cap V_{\text{heavy}})$ and $L = e(A \cap V_{\text{light}})$. Then, it is easy to see that $X - Z = H - L$. Using $(a-b)^2 \leq 2a^2 + 2b^2$, we have

$$\text{Var}[X - Z] = \text{Var}[H - L] \leq 2\text{Var}[H] + 2\text{Var}[L].$$

For any edge $e \in E(G)$, we write X_e to denote the indicator random variable for the event that both ends of e are picked in A. So we have

$$H = \sum_{e \in G[V_{\text{heavy}}]} X_e$$

and

$$L = \sum_{e \in G[V_{\text{light}}]} X_e.$$

We may express the variance of each of these random variables as a sum of covariances. The largest covariance terms arise from pairs of edges with one endvertex in common. As there are not too many heavy vertices (otherwise $\mathbb{P}[\mathcal{E}] = o(1)$), and as the degree of a light vertex is bounded, the number of such terms is not too large. Consequently, this will lead to the desired bound on $\text{Var}[X - Z]$. □

Finally, we use this to concentrate the most common degree in $G[A]$.

Proposition 3.5.5. *Assume that* $\ell = \omega(\log^3 k)$*. Then there exists some determinisic* $d = d(G, k, \ell)$ *such that* $\mathcal{E} \subsetneqq \mathcal{D}_d$*.*

Proof Sketch. Let D denote the random variable corresponding to the most common degree in $G[A]$. We first show that $\mathcal{E}$ is essentially contained in the following four events:

- $\mathcal{F}_1 = \{$every light vertex in A has degree at most $2\ell^{1/3}$ in $G[A]\}$,
- $\mathcal{F}_2 = \{$every heavy vertex in A has degree at least $\ell^{1/3}/2$ in $G[A]\}$,

- $\mathcal{F}_3 = \{X = Z + \mu \pm w_k \ell^{5/6}\}$, where $\mu = \mathbb{E}[X - Z]$,
- $\mathcal{F}_4 = \{Z = kD \pm 3w_k \ell^{5/6}\}$.

The first two follow from Chernoff bounds and a union bound, and Chebyshev's inequality implies the third. Further, by Proposition 3.5.2, we know $\mathcal{E} \subsetneq \mathcal{D}_*$. Thus, we have $\mathcal{E} \subsetneq \mathcal{E} \cap \mathcal{D}_* \cap \mathcal{F}_1 \cap \mathcal{F}_2$. To prove $\mathcal{E} \subsetneq \mathcal{F}_4$, it suffice to show $\mathcal{E} \cap \mathcal{D}_* \cap \mathcal{F}_1 \cap \mathcal{F}_2 \subseteq \mathcal{F}_4$.

Assume $\mathcal{E} \cap \mathcal{D}_* \cap \mathcal{F}_1 \cap \mathcal{F}_2$ holds. Under $\mathcal{D}_*$, we know that at least $k - w_k\sqrt{\ell}$ vertices in $G[A]$ have $d_A(v) = D$. We use $\mathcal{D}_*$ similar to how we did in Theorem 3.5.1 to argue that $D \leq 3\ell/k$. By $\mathcal{F}_2$, we know that all heavy vertices in A have degree at least $\ell^{1/3}/2 \geq \omega(\ell/k)$, so all the at least $k - w_k\sqrt{\ell}$ vertices in A of degree D must be light. Thus, we have

$$(k - w_k\sqrt{\ell})D \leq Z \leq (k - w_k\sqrt{\ell})D + w_k\sqrt{\ell} \cdot 2\ell^{1/3},$$

where the upper bound comes from $\mathcal{F}_1$. Using $D \leq 3\ell/k$ yields $\mathcal{F}_4$.

Finally, $\mathcal{E} \cap \mathcal{F}_3 \cap \mathcal{F}_4$ gives

$$D = \frac{\ell - \mu}{k} \pm \frac{w_k}{k} \cdot O(\ell^{5/6}).$$

Assuming w_k goes to infinity slowly enough, for large enough k, we see that there is (at most) on possible integer value of D satisfying this equation, which completes the proof. □

Remark 3.5.2. Observe that Propositions 3.5.2 and 3.5.5 prove that $\mathbb{P}[X = \ell] \leq 1/e + o(1)$ for $\omega(\log^3 k) \leq \ell \leq o(k^{6/5})$. This, along with Theorem 3.1.2, proves Theorem 3.1.1.

3.6 Extension to r-graphs and Further Discussion

As discussed, the only part of the proof of Theorem 3.1.1 that we have yet to successfully generalise to r-graphs is Section 3.5.3 – concentrating the most common degree of $G[A]$ to one deterministic value. In this section, we discuss some of the attempts we made to achieve this and the limitations of each approach.

The key point to note is that in the graph case, for any $v \in A$, the random variable $d_A(v)$ follows a hypergeometric distribution, and this fact is used quite strongly throughout

Section 3.5.3. However, as seen in Proposition 3.5.1 for instance, we know that $d_A(v)$ is not hypergeometric in the general r-graph setting. In order to generalise the arguments from graphs to uniform hypergraphs, we require two concentration bounds:

- Firstly, given any $v \in A$, we would require a bound on $\mathrm{Var}\left[d_A(v)\right]$ to argue that it is tightly concentrated around its expectation. Furthermore, we would ideally like to have an exponentially decaying tail beyond the standard deviation to use union bounds to simultaneously bound the degrees of all (light/heavy) vertices in A as this is useful when proving Proposition 3.5.5.
- Secondly, we need a bound on $\mathrm{Var}[X]$ or on something such as $\mathrm{Var}[X - Z]$.

We now briefly list some of the techniques we attempted and the issues we faced.

- **Direct calculations:** Both the above random variables ($d_A(v)$ and X) can be written as a sum of indicators, and one can upper bound the variance similar to the approach used in Proposition 3.5.4 by writing it out as a sum of covariances. However, the bounds obtained using this approach are too large since we cannot use Chernoff bounds like they were utilised in Proposition 3.5.4.
- **Fourier analysis of Boolean functions:** Instead of picking A as a uniformly random k-vertex subset from $V(G)^{(k)}$, we pick each vertex independently with probability k/n and use the set of picked vertices to approximate A (so the number of picked vertices follows a $\mathrm{Bin}(n, k/n)$ distribution, which means the expected number of vertices picked is k). Yet again, we may write $d_A(v)$ and X as sums of indicators and compute their variances as sums of covariances. The variance computed in this binomial approximation upper bounds the true variance (as the vertices are picked independently now, there are no negative covariance terms, unlike earlier). For instance, consider the case where G is a clique. In the original setting, it is clear that X will be exactly equal to $\binom{k}{r}$ and $d_A(v)$ will be $\binom{k-1}{r-1}$ and the variance will be zero. However, in the binomial approximation, the precise value of X and $d_A(v)$ will vary depending on the number of vertices picked, and hence these random variables will have nonzero variance. Hence, an upper bound on the variance of X and $d_A(v)$ in the binomial version of the problem will automatically provide an upper bound for the true variance.

Fourier analysis proves to be a powerful tool for calculating the variance of functions of binomial random variables (for an excellent reference, see [39]). Broadly speaking, the idea is to view the set of all Boolean functions (functions of the form $f : \{0,1\}^n \to \mathbb{R})$) as a real vector space with point-wise addition, define the inner product of two functions as the expectation of their product over a $\mathrm{Bin}(n, 1/2)$ binary string, and argue that this space has an orthonormal basis consisting of multilinear polynomials in $\boldsymbol{x} \in \{0,1\}^n$. Since any function can be expressed as a linear combination of these basis functions, the orthonormality of the basis and the definition of the inner product allow us to conveniently compute the expectation and the variance of the function in terms of the coefficients of this expansion, where the input to the function is a random binary string sampled from $\boldsymbol{x} \in \mathrm{Bin}(n, 1/2)$.

In our case, we sample from a $\mathrm{Bin}(n, k/n)$ distribution, where we have assumed $n \gg k$. Fortunately, we can still easily use the same method, the only difference being that we end up with an orthogonal basis in place of an orthonormal one. We then compute the variance exactly as described in [39, Proposition 1.13]. However, the variance bound obtained by this approach is too large to use second moment methods and Chebyshevs's inequality, which require $\mathrm{Var}[X] = o(\mathbb{E}[X]^2)$ (see [2, Chapter 4]), and hence is not good enough to concentrate X and $d_A(v)$ around their respective expectations.

- **Hypercontractivity:** Using the ideas of Chapter 2, it is easy to see that X can be expressed as a polynomial of degree r in $\boldsymbol{\xi} \in \mathrm{BL}(n,k)$, and $d_A(v)$ can similarly be written as a polynomial of degree $r-1$. We then use a hypercontractive inequality for random variables Y that are functions of $\boldsymbol{\xi}$. In particular, we use [30, Proposition 2.5] and utilise it by replicating the proof methodology of [30, Corollary 2.6]. The broad idea is to use this in conjunction with Markov's inequality as follows. For any $q > 0$, we have
$$\mathbb{P}[|Y| \geq t] = \mathbb{P}[|Y|^q \geq t^q] \leq \frac{\mathbb{E}|Y|^q}{t^q} \leq f(q,k,n)\frac{\left(\mathbb{E}\left[|Y|^2\right]\right)^{q/2}}{t^q},$$
for some function f, where the final inequality follows from the hypercontractive inequality. The aim is to optimise the value of q to get the best possible bound. In our case, as $n \gg k$, it is easy to see that the proof used in [30, Corollary 2.6] applies only when $q = 2$, and the above hypercontractivity bound reduces to Chebyshev's inequality (which can be seen by substituting $X - \mathbb{E}X$ or $d_A(v) - \mathbb{E}d_A(v)$ in place of Y). As discussed previously, we have been unable to find a good enough variance bound for X and $d_A(v)$ to use Chebyshev's inequality effectively.

- **McDiarmid-type inequalities:** The first inequality we tried using is Lemma 2.4.5. As explained in Proposition 2.5.1, when trying to concentrate X, the constants c_x are precisely equal to the degrees of the corresponding vertices $d_G(x)$. Once again, we maximise $\sum c_x^2$ by considering a hypergraph G where the degree distribution is as unbalanced as possible. Yet again, the bound obtained is too weak to obtain a meaningful concentration bound on $|X - \mathbb{E}X|$ w.h.p. (the only high probability bounds we get are trivial). This is not too hard to see because the denominator of the power of the exponential term in the bound will increase with n, which, as usual, is much larger than k, leading to a weak bound.

  We then tried using [21, Corollary 2.2], which, at least on the surface, appears better suited to our purpose (and does not involve n). However, it only provides a useful high probability bound when $e(G) \gg n^r/k^{1/2}$ – it allows us to say that in this case, we must have $\mathbb{P}[X \gg \ell] = 1 - o(1)$ and hence $\mathbb{P}[X = \ell] = o(1)$ (similar to what Propostion 2.5.1 does). So, it only allows us to rule out extremely dense hypergraphs. In other cases, the concentration bound it provides allows us to say that X lies in some interval w.h.p., but as $\ell \leq o(k^{6/5})$ is quite small in terms of k, the interval is far too large to conclude any reasonably strong result.

  Similar issues persist when trying to concentrate $d_A(v)$ as well – n is too large for the first technique, and the second one is only useful for vertices of extremely high degree.

When it comes to bounding the variance of X, it would appear that dense hypergraphs that are highly irregular (where some vertices have very large degrees and the rest have very small degrees) are the "bad" cases that lead to poor variance bounds. This will be evident from the direct calculation approach of the first method (because it is the same Proposition 3.5.4, where we saw that the largest covariance terms occur due to pairs of edges with precisely one vertex in common, so many vertices of large degree can lead to very big upper bounds on the covariance). The Fourier analytic method also points to the same conclusion. So perhaps a good idea would be to see if such dense, irregular hypergraphs can be handled in a different manner and then use the approach described in this chapter for more regular hypergraphs.

Bounding the variance of $d_A(v)$ and getting the desired exponential tails to use a union bound appears to be more challenging, and it is not clear what the right way to concentrate these random variables is.

Chapter 4

Quasirandom-Forcing Pairs of Tournaments: Introduction and Preliminaries

4.1 Introduction

A combinatorial structure is said to be *quasirandom* if it satisfies certain properties that hold asymptotically almost surely in a random structure. The study of quasirandomness was initiated by Chung, Graham and Wilson [7], Thomason [47, 48] and Rödl [44], who investigated the notion of quasirandomness in terms of different graph properties. One of the key takeaways from these studies is that several different graph properties all lead to the same notion of quasirandomness and can be equivalently used to characterise quasirandom graphs (for instance, see [2, Theorem 9.3.1]). Since then, quasirandom properties have been investigated in several different types of discrete structures, such as groups [20], hypergraphs [6, 8, 18, 19, 24, 27, 37, 46], permutations [5, 10, 29] and integers [9].

We are interested in tournament quasirandomness, a topic introduced by Chung and Graham [8], who showed that, as with graphs, a wide range of natural tournament properties can be used equivalently to describe quasirandomness. This idea, as with the previous ones, has drawn widespread attention [4, 11, 12, 13, 22, 23, 26], and the primary aim has typically been finding ways to characterise quasirandom tournaments.

We will define tournament quasirandomness in terms of the homomorphism density of subtournaments. A *homomorphism* from a directed graph D to a tournament T is a map $f : V(D) \to V(T)$ such that $f(u)f(v) \in A(T)$ whenever $uv \in A(D)$, where $A(F)$ denotes the set of arcs of a directed graph F. Let $\hom(D,T)$ denote the number of homomorphisms from D to T and define the *homomorphism density* of D in T to be

$$t(D,T) := \frac{\hom(D,T)}{v(T)^{v(D)}}$$

where $v(H)$ denotes the number of vertices in a digraph H.

We'll now see how to use this to define tournament quasirandomness. Suppose T is a random tournament on n vertices, i.e., each arc is directed randomly in one of the two possible directions independent of the directions of all other arcs and is equally likely to be oriented in either direction. Let H be a tournament on k vertices, where $n \geq k$. Then, the expected homomorphism density of H in T is

$$\mathbb{E}(t(H,T)) = \frac{(1/2)^{\binom{k}{2}} n(n-1)\cdots(n-k+1)}{n^k} = (1-o(1))(1/2)^{\binom{k}{2}}.$$

Note that we have only accounted for injective homomorphism in the numerator $\hom(H,T)$, but that is because the number of non-injective homomorphisms is clearly $o(n^k)$, and so this is all taken care of by the $o(1)$ term. This also tells that injective homomorphisms account for most homomorphisms from $V(H)$ to $V(T)$, which is what, in some sense, allows us to say we're approximately counting copies of H in T and looking at their density (up to constant factors). Simple concentration inequalities tell us that the random variable $t(H,T)$ is tightly concentrated around this expected value, and we use this observation to justify the following definition. A sequence $\{T_n\}_{n\geq 1}$ is said to be *quasirandom* if, for every $k \geq 1$ and every k-vertex tournament H,

$$\lim_{n\to\infty} t(H,T_n) = (1/2)^{\binom{k}{2}}.$$

In other words, it is quasirandom if every finite subtournament appears with roughly the same density one would expect it to appear within a random tournament. A key feature of quasirandomness is that it can be completely characterised only by the density of a few substructures, and determining these substructures is a popular theme of investigation. We say that a tournament H on k vertices is *quasirandom-forcing* if any sequence of tournaments

$\{T_n\}_{n\geq 1}$ satisfying $v(T_n) \to \infty$ and

$$\lim_{n\to\infty} t(H, T_n) = (1/2)^{\binom{k}{2}}$$

is quasirandom. That is to say, if the density of H is approaches its density in a random tournament, then that is true for all subtournaments.

The set of all quasirandom-forcing tournaments has only been determined recently. Let TT_k denote the transitive tournament on k vertices, obtained by listing all the vertices in a line and directing all possible arcs forward. Proving that TT_k is quasirandom forcing for all $k \geq 4$ is a small extension of [32, Exercise 10.44], and was reproved in [13] independently using flag algebras (see Section 5.3.2 for a brief introduction to the intuition and ideas behind this technique). Coregliano, Parente and Sato [11] also used flag algebras to obtain an example of a 5-vertex non-transitive quasirandom-forcing tournament $TT_4^{\updownarrow}$; see Figure 4.1. Subsequently, Bucić, Long and Shapira [4] proved that any quasirandom-forcing tournament on at least 7 vertices must be transitive. Finally, Hancock *et al.* [23] ruled out all other tournaments on at most 6 vertices. This completed the characterisation of all quasirandom-forcing tournaments, showing that the only ones are TT_k for $k \geq 4$ and $TT_4^{\updownarrow}$.

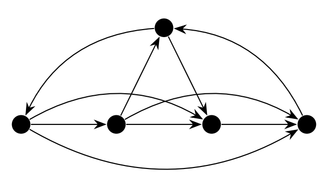

Figure 4.1: The tournament $TT_4^{\updownarrow}$

The class of quasirandom-forcing tournaments is rather small, and consequently trying to find more quasirandom-forcing structures motivates the following definition (analogous to one that was introduced in the context of quasirandom permutations in [5])

Definition 4.1.1. Let $k \geq 1$ and let S be a set of pairwise non-isomorphic tournaments on k vertices. We say that S is Σ-*forcing* if any sequence $\{T_n\}_{n\geq 1}$ of tournaments satisfying $v(T_n) \to \infty$ and

$$\lim_{n\to\infty} \sum_{H\in S} t(H, T_n) = |S| \cdot 2^{-\binom{k}{2}}$$

is quasirandom.

With this definition, it is clear that a tournament H is quasirandom-forcing if and only if $\{H\}$ is Σ-forcing.

Our goal is to extend the aforementioned studies on quasirandom-forcing tournaments to Σ-forcing pairs of tournaments, taking the first step towards understanding Σ-forcing sets in general. In particular, we will focus on non-transitive pairs of tournaments. It would also be interesting to study Σ-forcing sets that contain transitive tournaments, but the situation for such sets seems to be quite different, and we discuss it in Section 5.4.

In the next section, we introduce some preliminary notation and definitions that will come in handy. In Chapter 5, we first introduce tournament limits. We then describe in detail how we use these tools to show that most pairs of non-transitive k-vertex tournaments are not Σ-forcing. We proceed to discuss potential approaches to tackle the problem of classifying the remaining tournament pairs and finally conclude with some open problems.

4.2 Preliminaries

A lot of notation we use henceforth (such as notation regarding induced (di)graphs, asymptotic notation, etc) is shared with previous chapters and has been introduced in Section 1.2. We only present definitions and notations specific to digraphs and tournaments that we have not previously introduced.

We use standard notation for directed graphs and tournaments. A *directed graph* (or *digraph*) D is a pair $(V(D), A(D))$ such that $A(D) \subseteq V(D) \times V(D)$ consists of ordered pairs. The elements of $V(D)$ are *vertices*, and the elements of $A(D)$ are *arcs* or *directed edges*. In case there is no potential ambiguity, we sometimes denote $V(D)$ by V and $A(D)$ by A. An arc $(u, v) \in A(D)$ is often written uv for short. We let $v(D) = |V(D)|$ and $a(D) = |A(D)|$.

Given a vertex $v \in V(D)$, define its *out-neighbourhood* $N^+(v) := \{u : vu \in A(D)\}$ and its *out-degree* $d^+(v) := |N^+(v)|$. Similarly, define the *in-neighbourhood* $N^-(v) := \{u : uv \in A(D)\}$ and *in-degree* $d^-(v) := |N^-(v)|$.

A *tournament* T is a directed complete graph, i.e., for all distinct $u, v \in V(T)$, exactly one of the arcs uv or vu is in $A(T)$. The *transitive tournament* on k vertices, denoted TT_k, consists of the vertex set $\{v_1, \ldots, v_k\}$ with an arc from v_i to v_j if and only if $i < j$.

As stated previously, a *homomorphism* from a digraph D to a digraph G is a map $\varphi : V(D) \to V(G)$ such that if $uv \in A(D)$, then $\varphi(u)\varphi(v) \in A(G)$. The set of all homomorphisms from D to G is written $\operatorname{Hom}(D, G)$ and the number of such homomorphisms is $\hom(D, G) := |\operatorname{Hom}(D, G)|$. The *homomorphism density* of D in G, denoted $t(D, G)$, is defined as

$$t(D, G) = \frac{\hom(D, G)}{v(G)^{v(D)}}.$$

Intuitively, $t(D, G)$ is the proportion of all functions from $V(D)$ to $V(G)$ which are homomorphisms. Equivalently, it is the probability that a randomly chosen function from $V(D)$ to $V(G)$ is a homomorphism. An *isomorphism* is a bijective homomorphism, and an *automorphism* is an isomorphism from a digraph to itself. We let $\operatorname{Aut}(D)$ denote the group of automorphisms of D and set $\operatorname{aut}(D) := |\operatorname{Aut}(D)|$.

Chapter 5

Classifying Pairs of Tournaments

The major portion of this chapter is dedicated to showing that all pairs of non-transitive tournaments on at least seven vertices are not Σ-forcing. Following this, we briefly discuss future directions of research and potential approaches to characterise the remaining pairs. Finally, we conclude with some conjectures regarding pairs containing a transitive tournament.

5.1 Tournament Limits

In order to rule out pairs of tournaments, we first introduce *tournament limits* or *tournamentons* and translate notions such as homomorphism densities and quasirandomness into this framework. This will allow us to prove certain necessary conditions for a pair of tournaments to be Σ-forcing and hence rule out most pairs. Our discussion will be brief; for a more detailed treatment of combinatorial limits, see [33]. While most of its description is for graph limits, this theory quite easily extends to tournaments, which is what we shall present.

A *tournamenton* or a *tournament limit* is a (Lebesgue) measurable function $W : [0,1]^2 \to [0,1]$ satisfying $W(x,y) + W(y,x) = 1$ for all $x, y \in [0,1]$. The idea behind this definition is to think of a tournamenton as the continuous generalisation of the adjacency matrix of a tournament. Intuitively (and very loosely), we can think of a tournamenton W as giving the "arc" xy the "weight" $W(x,y)$ and the arc yx the weight $W(y,x) = 1 - W(x,y)$, or directing the arc from x to y with probability $W(x, y)$. We can easily convert a tournament to a tournamenton. Given a tournament T on n vertices labeled $\{v_1, \ldots, v_n\}$, define the *associated tournamenton* $W_T : [0,1]^2 \to [0,1]$ obtained by partitioning $[0,1] = I_1 \cup I_2 \cdots \cup I_n$, where $I_j = [(j-1)/n, j/n)$ for $1 \leq j \leq n-1$ and $I_n = [(n-1)/n, 1]$, such that if $(x, y) \in I_i \times I_j$, then $W_T(x, y) = 1$ if $v_i v_j \in A(T)$ and is 0 otherwise. This definition becomes easy to visualise by viewing the adjacency matrix of T as a "pixel picture" by placing a black square wherever there is a 1 and a white square for a 0.

$$\begin{pmatrix} 0 & 0 & 0 & 0 \\ 1 & 0 & 0 & 0 \\ 1 & 1 & 0 & 0 \\ 1 & 1 & 1 & 0 \end{pmatrix} \longrightarrow$$

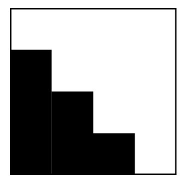

Figure 5.1: The adjacency matrix of TT_4 and the corresponding tournamenton. It takes the value 1 on the black region and 0 on the white region.

Given a digraph D with vertex set $V(D) = \{v_1, \ldots, v_k\}$ and a tournamenton W, the *homomorphism density* of D in W is defined as

$$t(D, W) := \int_0^1 \cdots \int_0^1 \prod_{v_i v_j \in A(D)} W(x_i, x_j) dx_1 \cdots dx_k. \tag{5.1}$$

It is fairly easy to see that $t(D, T) = t(D, W_T)$ for any tournament T. Another way to view $t(D, W)$ that ties into the intuition we previously described is as follows. Pick k points $x_1, \ldots, x_k$ independently from $[0,1]$ uniformly randomly. For each $v_i v_j \in A(D)$, add an arc from x_i to x_j with probability $W(x_i, x_j)$. Then $t(D, W)$ is precisely the probability that all edges of D are added during this procedure.

A sequence $\{T_n\}_{n \geq 1}$ of tournaments with $v(T_n) \to \infty$ is said to be *convergent* if $\lim_{n\to\infty} t(H, T_n)$ exists for every tournament H. A tournamenton W is the *limit* of a sequence $\{T_n\}_{n\geq 1}$ of tournaments if $\lim_{n\to\infty} t(H, T_n) = t(H, W)$ for every tournament H. A well-known result from the theory of graph limits that can be adapted to tournaments is that every convergent sequence of tournaments has a limit tournamenton (see, for instance, [22, 23]). Moreover, this limit can be shown to be unique up to reordering the vertices of the tournaments in the sequence, but this will not be necessary for us. Also, for every tournamenton W, there is a sequence of finite tournaments whose limit is W. The aforementioned pixel pictures often provide a simple mechanism to visualise the limiting tournamenton of a sequence of tournaments. For instance, the limit of the sequence $\{TT_n\}_{n\geq 1}$ can be visualised as

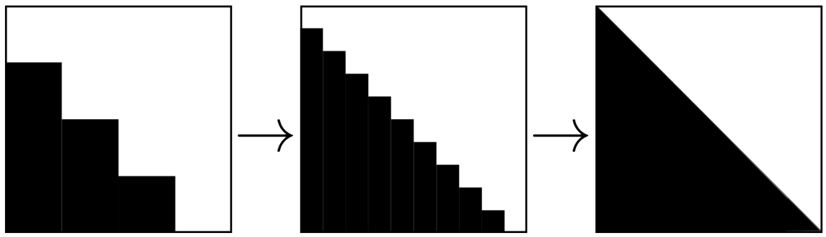

Figure 5.2: The limit of an increasing sequence of transitive tournamentons.

Formally, the limiting *transitive tournamenton* is

$$W_{TT}(x,y) := \begin{cases} 1 & \text{if } x < y, \\ 1/2 & \text{if } x = y, \\ 0 & \text{otherwise.} \end{cases} \quad . \tag{5.2}$$

Note that tournamentons are Lebesgue measurable functions and hence are defined up to equality almost everywhere, so we can set any arbitrary value for W_{TT} along the "diagonal" $\{(x,x) : x \in [0,1]\}$.

The language of tournamentons allows us to formulate a new, equivalent definition of quasirandomness (see [22, 23]).

Proposition 5.1.1. *A sequence of tournaments $\{T_n\}_{n\geq 1}$ is quasirandom if and only if its limit is the tournamenton W such that $W(x,y) = 1/2$ for all $x, y \in [0,1]$.*

One can think of the constant-1/2 tournamenton as the limit of a sequence of random tournaments of increasing order. Related to the proposition above, by modifying standard arguments from graph limits, we also obtain a new definition for a Σ-forcing set of tournaments.

Proposition 5.1.2. *Let $k \geq 1$ and let S be a set of k-vertex tournaments. Then S is Σ-forcing if every tournamenton satisfying*

$$\sum_{H\in S} t(H,W) = |S| \cdot (1/2)^{\binom{k}{2}},$$

is equal to $1/2$ *almost everywhere, that is,* $W(x,y) = 1/2$ *for almost all* $(x,y) \in [0,1]^2$.

5.2 Finding Tournamentons

By the proposition above, in order to show that a set S is not Σ-forcing, it is enough to find a tournamenton that is not almost everywhere equal to the constant-$1/2$ tournamenton but still satisfies the condition above. Our goal in this section is to find examples of such tournamentons that can be used to rule out all pairs of tournaments on at least seven vertices.

First, we simplify the problem to show that we only need to find a tournamenton with a high density of tournaments in S. This will follow using a standard intermediate value theorem argument, such as [23, Proposition 2].

Lemma 5.2.1. *Let* $k \geq 1$ *and let* S *be a set of* k*-vertex non-transitive tournaments. If there exists a tournamenton* W *such that* W *is not equal to* $1/2$ *almost everywhere and*

$$\sum_{H \in S} t(H,W) \geq |S| \cdot (1/2)^{\binom{k}{2}},$$

then S *is not* Σ*-forcing.*

Proof. Let W be the tournamenton given in the statement of the lemma, and let W_{TT} be the transitive tournamenton as described in Equation 5.2. For any $\alpha \in [0,1]$, define the tournamenton

$$W_\alpha = \begin{cases} W(x,y) & \text{if } (x,y) \in [0,\alpha]^2, \\ W_{TT}(x,y) & \text{otherwise.} \end{cases}$$

From the definitions of W and W_{TT}, we note that W_α is not equal to $1/2$ almost everywhere for all $\alpha \in [0,1]$. By assumption, we know that

$$\sum_{H \in S} t(H,W_1) = \sum_{H \in S} t(H,W) \geq |S| \cdot (1/2)^{\binom{k}{2}}.$$

On the other hand, since no tournament $H \in S$ is transitive, we will have $t(H,W_{TT}) = t(H,W_0) = 0$ for all $H \in S$. Indeed, consider any $H \in S$ with vertices labelled $\{v_1,\ldots,v_k\}$. It is an easy exercise to show that a tournament is transitive if and only if it does not have any directed cycles. Hence, H must have a directed cycle, which, without loss of generality, we may assume is $v_1 \dots v_\ell$. Referencing Equation 5.1 and using the definition of W_{TT} from Equation 5.2, it is clear that $W_{TT}(x_1x_2)\dots W_{TT}(x_{\ell-1}x_\ell)W_{TT}(x_\ell x_1) = 0$ for almost all $x_1, \dots, x_\ell \in [0,1]$, and hence $t(H, W_{TT}) = t(H, W_0) = 0$. As this holds for all $H \in S$, we have

$$\sum_{H \in S} t(H, W_0) = 0.$$

Since $\sum_{H \in S} t(H, W_\alpha)$ is a continuous function of α, by the intermediate value theorem, there exists some $\alpha \in (0, 1]$ such that

$$\sum_{H \in S} t(H, W_\alpha) = |S| \cdot (1/2)^{\binom{k}{2}}.$$

This, along with the fact that W_α is not equal to $1/2$ almost everywhere and Proposition 5.1.2, completes the proof. $\square$

One standard idea for constructing a tournamenton W that has a high density of tournaments in S, used in [4, Proof of Proposition 1.2], is to consider a "blow up" of one of the tournaments in S. Let H be a tournament on k vertices labelled $\{v_1, \dots, v_k\}$. For $n \geq k$, construct the n-vertex blow-up tournament of H, denoted $W_H^*(n)$ as follows. Let $V(W_H^*(n)) = V = \{u_1, \dots, u_n\}$ and let $V = \bigsqcup_{i \in [k]} V_i$ be a partition of V such that $|V_i| = \lfloor n/h \rfloor$ or $\lceil n/h \rceil$ for all $i \in [k]$. For every arc v_iv_j of H, direct all arcs from V_i to V_j in $W_H^*(n)$ and direct the remaining arcs arbitrarily. In terms of adjacency matrices, we obtain $W_H^*(n)$ from H by replacing the ones and zeroes in the adjacency matrix of H with appropriately sized blocks of ones and zeroes, respectively (and the remaining entries of the matrix can be decided arbitrarily). Set W_H^* to be the limiting tournamenton of the sequence of blow-ups $\{W_H^*(n)\}_{n \geq 1}$.

$$\begin{pmatrix} 0 & 0 & 0 & 0 \\ 1 & 0 & 0 & 0 \\ 1 & 1 & 0 & 0 \\ 1 & 1 & 1 & 0 \end{pmatrix}$$

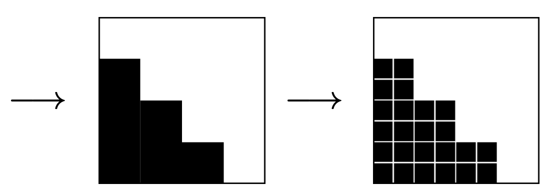

Figure 5.3: Blow-up of TT_4 on 8 vertices.

The key observation we make here is that any map $f : V(H) \to V$ with $f(v_i) \in V_i$ for all $i \in [k]$ constitutes a homomorphism, and hence we see that

$$\hom(H, W_H^*(n)) \geq (1 - o(1))n^k/k^k,$$

for all $n \geq k$. This implies that

$$t(H, W_H^*) \geq (1 - o(1))k^{-k}.$$

We will utilise this to prove the following lemma, which can then be used to rule out all 7 vertex pairs.

Lemma 5.2.2. *If S is a Σ-forcing set of non-transitive k-vertex tournaments, then*

$$|S| \geq 2^{\binom{k}{2}} k^{-k}.$$

Proof. We will prove the lemma by showing the contrapositive. Suppose $|S| < 2^{\binom{k}{2}} k^{-k}$. Fix any tournament $H_0 \in S$ and set $W = W_{H_0}^*$ to be the blow-up tournamenton. From the preceding discussion, we know that

$$\sum_{H \in S} t(H, W) \geq t(H_0, W) \geq (1 - o(1))k^{-k} \geq |S| \cdot 2^{-\binom{k}{2}},$$

where the final inequality holds for sufficiently large n by assumption. Then Lemma 5.2.1 completes the proof. □

Using that $k^{-k} > 2^{1-\binom{k}{2}}$ for $k \geq 7$, we immediately obtain the following result.

Corollary 5.2.3. *For all $k \geq 7$, there is no Σ-forcing pair of non-transitive k-vertex tournaments.*

Henceforth, we focus our efforts on non-transitive pairs of tournaments on at most 6 vertices. For $k = 2, 3$, any pair of k-vertex tournaments must contain at least one transitive tournament, and hence we only need to look at $k = 4, 5, 6$, which we discuss next.

5.3 Further Directions of Research: Characterising the Remaining Pairs

In this section, we intend to briefly explain some of the methods we aim to use to classify the remaining tournament pairs. As discussed at the end of Section 5.2, we only need to consider non-transitive pairs of tournaments on $k = 4, 5, 6$ vertices. We split this section into two parts – finding pairs that are Σ-forcing and ruling out the rest.

5.3.1 Eliminating Pairs on At Most Six Vertices

We first detail some of the strategies we wish to consider to classify non-Σ-forcing pairs. As usual, we aim to do this by employing Lemma 5.2.1 – given a pair of non-transitive k-vertex tournaments S, we wish to find a "high density" tournamenton that is not equal to $1/2$ almost everywhere.

- **Perturbing the $1/2$-tournamenton:** The approach here is similar to [23, Section 4.2]. We consider the following tournament matrix $T(x)$ which is a slight "perturbation" of the constant $1/2$ tournamenton

  $$T(x) = \begin{pmatrix} 1/2 & 1/2 + x \\ 1/2 - x & 1/2 \end{pmatrix}.$$

  Consider any $H \in S$, and then compute $t\left(H, T(x)\right)$ using 5.1 as a function of $x \in [-1/2, 1/2]$. If we can show there exists some $x \neq 0$ such that $t\left(H, T(x)\right)$ is much larger than the $(1/2)^{\binom{k}{2}}$, then it is likely that the pair S satisfies the conditions of Lemma 5.2.1. One way to do this would be to check the first and second derivatives of $t\left(H, T(x)\right)$ at $x = 0$ to see if a local minimum because, in that case, it is possible that some small nonzero x will do the trick. This can also be extended to multiple variables, for instance, by considering

  $$T(x, y, z) = \begin{pmatrix} 1/2 & 1/2 + x & 1/2 + y \\ 1/2 - x & 1/2 & 1/2 + z \\ 1/2 - y & 1/2 - z & 1/2 \end{pmatrix}.$$

Again, the aim is to find some $x, y, z \in [-1/2, 1/2]$ satisfying $(x, y, z) \neq (0, 0, 0)$ and $t(H, T(x, y, z)) > (1/2)^{\binom{k}{2}}$. In this case, we can look at the gradient and the Hessian of $t(H, T(x, y, z))$ at $(0, 0, 0)$ to determine whether it is a local minimum.

- **Monte Carlo methods:** Let M be an $n \times n$ matrix (for some fixed n) such that all entries in M are non-negative and $M + M^T$ is the all-ones matrix. Viewing M as a tournamenton, we compute the sum of densities $\sum_{H \in S} t(H, M)$. In order to find some M such that this sum is greater than $|S| \cdot (1/2)^{\binom{k}{2}}$, we execute the following iterative procedure. We pick a random row i and a random column j and increase m_{ij} by some pre-determined positive value ε and also decrease m_{ji} by ε so that $m_{ij} + m_{ji} = 1$ still holds. We then calculate the change in the sum of densities. If the change is positive, we accept it with high probability, and if it is negative, we accept it with low probability (we don't reject these changes outright to avoid getting stuck in a local minimum). We repeat this process until we get a tournamenton with the desired property or until we cross some fixed number of steps.

  There are a lot of initial conditions here that we can vary to see what works best – the value of n, the initial matrix M, the perturbation ε – so some trials are likely required to determine the best possible values of these parameters.

- **Blow-ups of small tournaments:** The idea here is similar to that of Lemma 5.2.2. However, instead of taking a blow-up of a tournament in S, we consider blow-ups of all possible tournaments on k, $k+1$, $k+2$, ... vertices. As $k \in \{4, 5, 6\}$ is quite small, it is fairly easy to use a computer to run through all tournament blow-ups on a few more vertices and make the homomorphism density calculations.

5.3.2 Finding Σ-Forcing Pairs

The key tool we use to figure out whether a pair S is Σ-forcing is the theory of *flag algebras* introduced by Razborov [42]. We will briefly outline the broad ideas underlying this technique and refer the reader to [42, 43, 15, 11] for more detailed discourse. Flag algebras provide a computational method to tackle graph-theoretic problems regarding homomorphism densities and related topics. Loosely speaking, they provide a formal setup to address the notions of adding and multiplying homomorphism densities, and they help solve problems regarding homomorphism density inequalities using techniques from sum-of-squares optimisation.

Suppose we wish to find the asymptotic maximum density of a particular tournament in a given family of tournaments (where the asymptotics are in terms of the size of the host tournament). Instead of considering every possible host tournament (or, more precisely, every possible increasing sequence of host tournaments), flag algebras allow us to concentrate only on certain specific host tournament sequences where determining the asymptotic density is computationally tractable, thus providing an upper bound on the maximum density, with the bounds provided often tight.

Given a particular member of the above family H, we consider the subfamily of all tournaments that contain a labelled embedding of H, which we shall call *H-flags*. We define an algebra over this subfamily by first considering all formal linear combinations of the tournaments over $\mathbb{R}$. When considering the density of one tournament of this subfamily in another, we will do so in a manner that preserves the labelled embedding of H. One can then show that this asymptotic density function is a linear function over this space. Furthermore, one can define a bilinear product in this space, where the product of T_1 and T_2 is a linear combination of all D in the subfamily with the coefficients corresponding to the asymptotic joint density of T_1 and T_2 in D. The space equipped with this product will form a commutative associative algebra over $\mathbb{R}$, which we call a *flag algebra* denoted $\mathcal{A}$. A crucial outcome of these definitions will be that the density of a tournament T in the product of two tournaments will simply be the product of the density T in each individual tournament, allowing us to think of this density function as an algebra homomorphism from $\mathcal{A}$ to $\mathbb{R}$. Furthermore, one can show that the set of density functions (each corresponding to the density of a particular tournament) is exactly equal to the set of all non-negative homomorphisms from $\mathcal{A}$ to $\mathbb{R}$.

When computing the density of a particular tournament, a common method is to express it in terms of the densities of other tournaments using double counting arguments by counting copies of a particular structure in different ways. For instance, given any tournament T on n vertices, a simple observation is

$$\hom(TT_3, T) + \frac{1}{3}\hom(C_3, T) = \binom{n}{3},$$

where C_3 is the directed cycle on 3 vertices. Indeed, the right-hand side counts the number of 3-vertex subsets in $V(T)$. As TT_3 and C_3 are the only two tournaments on 3 vertices (up to isomorphism), every set of three vertices in T induces one of these two tournaments. As

$\mathrm{aut}(TT_3) = 1$ and $\mathrm{aut}(C_3) = 3$, we divide the number of homomorphisms by the number of automorphisms to ensure that the left-hand side also counts each 3-vertex set only once. This lets us express the density of C_3 in terms of the density of TT_3. Consequently, providing a bound on the density of one yields a bound on that of the other.

The flag algebra framework allows us to formally write the density of one tournament in terms of sums and products of the densities of other tournaments. The goal is to find an expression that is easy to bound. In particular, it is possible to adapt ideas from sum-of-squares optimisation for polynomials to flag algebras. These techniques introduce a constraint that typically forces certain parameters to satisfy a polynomial equation which can be written as a sum of squares of other polynomials and hence must be non-negative. Hence, instead of optimising over the entire parameter space, we now have a restricted space where there are well-developed optimisation methods. For instance, if a particular linear combination of tournament densities can be bounded below by a sum of squares of densities of other tournaments, then this linear combination must be non-negative. Tools from semidefinite programming can be utilised to automate this process almost entirely, and this will directly provide bounds on homomorphism densities.

These algorithms can often provide more than just the final bounds on the densities; they can also determine conditions that must be satisfied by any homomorphism that attains these bounds, which can then be used to find the extremal constructions. This feature, along with characterisations of Σ-forcing sets like Proposition 5.1.2, can identify sets that are indeed Σ-forcing (for instance, see [11, Sections 5 and 6] for such an approach).

5.4 Open Problems

So far, we have considered Σ-forcing sets consisting of a pair of non-transitive tournaments. In contrast, from flag algebra calculations, it seems that if k is large enough, then every pair comprised of the k-vertex transitive tournament any other k-vertex tournament should be Σ-forcing.

Conjecture 5.4.1. *There exists k_0 such that if $k \geq k_0$ and H is any k-vertex tournament, then $S = \{H, TT_k\}$ is Σ-forcing.*

Furthermore, it may even be true that any linear expression of homomorphism densities of k-vertex tournaments containing a "large enough coefficient" on the density of the transitive tournament forces quasirandomness.

Conjecture 5.4.2. *For any $\alpha > 0$ there exists $k_0(\alpha)$ such that if $k \geq k_0(\alpha)$ and $c(H)$ is a real number for each tournament H on k vertices such that $\sum_{H:v(H)=k} c(H) = 1$ and $c(TT_k) \geq \alpha$, then a sequence $T_1, T_2, \ldots$ with $v(T_n) \to \infty$ is quasirandom if and only if*

$$\lim_{n\to\infty} \sum_{H:v(H)=k} c(H) \cdot t(H, T_n) = (1/2)^{\binom{k}{2}}.$$

Bibliography

[1] Noga Alon, Dan Hefetz, Michael Krivelevich, and Mykhaylo Tyomkyn. Edge-statistics on large graphs. *Combinatorics, Probability and Computing*, 29(2):163–189, 2020.

[2] Noga Alon and Joel H Spencer. *The probabilistic method*. John Wiley & Sons, 2016.

[3] József Balogh, Ping Hu, Bernard Lidickỳ, and Florian Pfender. Maximum density of induced 5-cycle is achieved by an iterated blow-up of 5-cycle. *European Journal of Combinatorics*, 52:47–58, 2016.

[4] Matija Bucić, Eoin Long, Asaf Shapira, and Benny Sudakov. Tournament quasirandomness from local counting. *Combinatorica*, 41(2):175–208, 2021.

[5] Timothy FN Chan, Daniel Král', Jonathan A Noel, Yanitsa Pehova, Maryam Sharifzadeh, and Jan Volec. Characterization of quasirandom permutations by a pattern sum. *Random Structures & Algorithms*, 57(4):920–939, 2020.

[6] Fan R. K. Chung and Ronald L. Graham. Quasi-random hypergraphs. *Random Structures & Algorithms*, 1(1):105–124, 1990.

[7] Fan R. K. Chung, Ronald L. Graham, and Richard M. Wilson. Quasi-random graphs. *Combinatorica*, 9:345–362, 1989.

[8] Fan RK Chung and Ronald L Graham. Quasi-random set systems. *Journal of the American Mathematical Society*, 4(1):151–196, 1991.

[9] Fan RK Chung and Ronald L Graham. Quasi-random subsets of zn. *Journal of Combinatorial Theory, Series A*, 61(1):64–86, 1992.

[10] Joshua N Cooper. Quasirandom permutations. *Journal of Combinatorial Theory, Series A*, 106(1):123–143, 2004.

[11] L N. Coregliano, R F. Parente, and C M. Sato. On the maximum density of fixed strongly connected subtournaments. *Electron. J. Combin.*, 26(1):Paper No. 1.44, 48, 2019.

[12] Leonardo N Coregliano and Alexander A Razborov. Natural quasirandomness properties. *Random Structures & Algorithms*, 63(3):624–688, 2023.

[13] Leonardo Nagami Coregliano and Alexander A Razborov. On the density of transitive tournaments. *Journal of Graph Theory*, 85(1):12–21, 2017.

[14] Kevin P Costello. Bilinear and quadratic variants on the littlewood-offord problem. *Israel Journal of Mathematics*, 194:359–394, 2013.

[15] Marcel K de Carli Silva, Fernando Mário de Oliveira Filho, and Cristiane Maria Sato. Flag algebras: a first glance. *arXiv preprint arXiv:1607.04741*, 2016.

[16] Jacob Fox and Lisa Sauermann. A completion of the proof of the edge-statistics conjecture. *arXiv preprint arXiv:1809.01352*, 2018.

[17] Jacob Fox and Benny Sudakov. Unavoidable patterns. *Journal of Combinatorial Theory, Series A*, 115(8):1561–1569, 2008.

[18] W Timothy Gowers. Quasirandomness, counting and regularity for 3-uniform hypergraphs. *Combinatorics, Probability and Computing*, 15(1-2):143–184, 2006.

[19] W Timothy Gowers. Hypergraph regularity and the multidimensional szemerédi theorem. *Annals of Mathematics*, pages 897–946, 2007.

[20] William T Gowers. Quasirandom groups. *Combinatorics, Probability and Computing*, 17(3):363–387, 2008.

[21] Catherine Greenhill, Mikhail Isaev, Matthew Kwan, and Brendan D McKay. The average number of spanning trees in sparse graphs with given degrees. *European Journal of Combinatorics*, 63:6–25, 2017.

[22] Andrzej Grzesik, Daniel Il'kovič, Bartłomiej Kielak, and Daniel Král'. Quasirandom-forcing orientations of cycles. *SIAM Journal on Discrete Mathematics*, 37(4):2689–2716, 2023.

[23] Robert Hancock, Adam Kabela, Taísa Martins, Roberto Parente, Fiona Skerman, Jan Volec, et al. No additional tournaments are quasirandom-forcing. *European Journal of Combinatorics*, 108:103632, 2023.

[24] Julie Haviland and Andrew Thomason. Pseudo-random hypergraphs. In *Annals of Discrete Mathematics*, volume 43, pages 255–278. Elsevier, 1989.

[25] Dan Hefetz and Mykhaylo Tyomkyn. On the inducibility of cycles. *Electronic Notes in Discrete Mathematics*, 61:593–599, 2017.

[26] Subrahmanyam Kalyanasundaram and Asaf Shapira. A note on even cycles and quasirandom tournaments. *Journal of Graph Theory*, 73(3):260–266, 2013.

[27] Yoshiharu Kohayakawa, Vojtěch Rödl, and Jozef Skokan. Hypergraphs, quasi-randomness, and conditions for regularity. *Journal of Combinatorial Theory, Series A*, 97(2):307–352, 2002.

[28] Daniel Král, Sergey Norin, and Jan Volec. A bound on the inducibility of cycles. *Journal of Combinatorial Theory Series A*, 161:359–363, 2019.

[29] Daniel Král' and Oleg Pikhurko. Quasirandom permutations are characterized by 4-point densities. *Geometric and Functional Analysis*, 23(2):570–579, 2013.

[30] Matthew Kwan, Benny Sudakov, and Tuan Tran. Anticoncentration for subgraph statistics. *Journal of the London Mathematical Society*, 99(3):757–777, 2019.

[31] Alexander E Litvak, Anna Lytova, Konstantin Tikhomirov, Nicole Tomczak-Jaegermann, and Pierre Youssef. Adjacency matrices of random digraphs: singularity and anti-concentration. *Journal of Mathematical Analysis and Applications*, 445(2):1447–1491, 2017.

[32] László Lovász. *Combinatorial problems and exercises*, volume 361. American Mathematical Soc., 2007.

[33] László Lovász. *Large networks and graph limits*, volume 60. American Mathematical Soc., 2012.

[34] Anders Martinsson, Frank Mousset, Andreas Noever, and Miloš Trujić. The edge-statistics conjecture for $\ell \ll k^{6/5}$. *Israel Journal of Mathematics*, 234(2):677–690, 2019.

[35] Colin McDiarmid. Concentration. In *Probabilistic methods for algorithmic discrete mathematics*, pages 195–248. Springer, 1998.

[36] Raghu Meka, Oanh Nguyen, and Van Vu. Anti-concentration for polynomials of independent random variables. *arXiv preprint arXiv:1507.00829*, 2015.

[37] Brendan Nagle, Vojtěch Rödl, and Mathias Schacht. The counting lemma for regular k-uniform hypergraphs. *Random Structures & Algorithms*, 28(2):113–179, 2006.

[38] Hoi H Nguyen. Inverse littlewood–offord problems and the singularity of random symmetric matrices. 2012.

[39] Ryan O'Donnell. *Analysis of boolean functions*. Cambridge University Press, 2014.

[40] Nicholas Pippenger and Martin Charles Golumbic. The inducibility of graphs. *Journal of Combinatorial Theory, Series B*, 19(3):189–203, 1975.

[41] A Razborov and E Viola. Real advantage, submitted.

[42] Alexander A Razborov. Flag algebras. *The Journal of Symbolic Logic*, 72(4):1239–1282, 2007.

[43] Alexander A Razborov. Flag algebras: an interim report. In *The Mathematics of Paul Erdős II*, pages 207–232. Springer, 2013.

[44] Vojtěch Rödl. On universality of graphs with uniformly distributed edges. *Discrete Mathematics*, 59(1-2):125–134, 1986.

[45] Vojtěch Rödl and Mathias Schacht. Generalizations of the removal lemma. *Combinatorica*, 29(4):467–501, 2009.

[46] Vojtěch Rödl and Jozef Skokan. Regularity lemma for k-uniform hypergraphs. *Random Structures & Algorithms*, 25(1):1–42, 2004.

[47] Andrew Thomason. Pseudo-random graphs. In *North-Holland Mathematics Studies*, volume 144, pages 307–331. Elsevier, 1987.

[48] Andrew Thomason. Random graphs, strongly regular graphs and pseudorandom graphs. *Surveys in combinatorics*, 123(173-195):1, 1987.

[49] Raphael Yuster. On the exact maximum induced density of almost all graphs and their inducibility. *Journal of Combinatorial Theory, Series B*, 136:81–109, 2019.