

Removal and stability for Erdős–Ko–Rado

Shagnik Das ^{*†}

Tuan Tran ^{*‡}

March 9, 2016

Abstract

A k -uniform family of subsets of $[n]$ is *intersecting* if it does not contain a disjoint pair of sets. The study of intersecting families is central to extremal set theory, dating back to the seminal Erdős–Ko–Rado theorem of 1961 that bounds the size of the largest such families. A recent trend has been to investigate the structure of set families with few disjoint pairs. Friedgut and Regev proved a general removal lemma, showing that when $\gamma n \leq k \leq (\frac{1}{2} - \gamma)n$, a set family with few disjoint pairs can be made intersecting by removing few sets.

We provide a simple proof of a removal lemma for large families, showing that families of size close to $\ell \binom{n-1}{k-1}$ with relatively few disjoint pairs must be close to a union of ℓ stars. Our lemma holds for a wide range of uniformities; in particular, when $\ell = 1$, the result holds for all $2 \leq k < \frac{n}{2}$ and provides sharp quantitative estimates.

We use this removal lemma to answer a question of Bollobás, Narayanan and Raigorodskii regarding the independence number of random subgraphs of the Kneser graph $K(n, k)$. The Erdős–Ko–Rado theorem shows $\alpha(K(n, k)) = \binom{n-1}{k-1}$. For some constant $c > 0$ and $k \leq cn$, we determine the sharp threshold for when this equality holds for random subgraphs of $K(n, k)$, and provide strong bounds on the critical probability for $k \leq \frac{1}{2}(n - 3)$.

1 Introduction

Extremal set theory, with its many connections and applications to other areas, has experienced remarkable growth in recent decades. Inspired by one of the cornerstones of the field, the celebrated Erdős–Ko–Rado theorem of 1961, a great deal of research concerns intersecting families. We say a family of sets is *intersecting* if it does not contain a pair of disjoint sets. In this paper we derive a stability result for large families that are nearly intersecting, and apply it to obtain a sparse extension of the Erdős–Ko–Rado theorem. We begin with a brief survey of relevant results.

1.1 Intersecting families and stability

We restrict our attention to k -uniform families of subsets of $[n]$. The natural extremal question is to ask how large such a family can be if it is intersecting. When $n < 2k$, there are no two disjoint sets, and hence $\binom{[n]}{k}$ itself is intersecting. For $n \geq 2k$, a natural construction is to take all sets

^{*}Freie Universität Berlin, Institut für Mathematik, Arnimallee 3, 14195 Berlin, Germany.

[†]shagnik@mi.fu-berlin.de

[‡]tuan@math.fu-berlin.de. Research supported by DFG within the Research Training Grant “Methods for Discrete Structures”.

containing some fixed element $i \in [n]$. This family, called the *star with centre i* , contains $\binom{n-1}{k-1}$ sets, and Erdős, Ko and Rado [11] showed this is best possible.

Given the extremal result, great efforts have been made to better understand the general structure of large intersecting families. Hilton and Milner [20] determined the size of the largest intersecting family that is not a subset of a star, before Frankl [13] extended this to determine the size of the largest intersecting family not containing too large a star.

In the years since these initial papers appeared, a series of stability results have been obtained. Friedgut [14] and Dinur and Friedgut [10] used spectral techniques to show, provided $k \leq (\frac{1}{2} - \gamma)n$ for some $\gamma > 0$, any intersecting family of size close to $\binom{n-1}{k-1}$ is almost entirely contained in a star. Keevash and Mubayi [23] and Keevash [22] combined these methods with combinatorial arguments to provide similar results when k is close to $\frac{1}{2}n$.

However, a recent trend in extremal set theory is to go beyond the Erdős–Ko–Rado threshold and study set families that may not be intersecting, but contain few disjoint pairs. Das, Gan and Sudakov [7] studied the supersaturation problem, determining the minimum number of disjoint pairs appearing in sufficiently sparse k -uniform families. Furthermore, a probabilistic variant of this supersaturation problem was introduced by Katona, Katona and Katona [21], and further studied by Russell [25], Russell and Walters [26] and Das and Sudakov [8].

Another direction that has been pursued is the transferral of the Erdős–Ko–Rado theorem to the sparse random setting. This study was initiated by Balogh, Bohman and Mubayi [1], who asked when the largest intersecting subfamily of a random k -uniform hypergraph is the largest star. Progress on this problem has been made in subsequent papers by Gauy, Hàn and Oliveira [17], Balogh, Das, Delcourt, Liu and Sharifzadeh [3] and Hamm and Kahn [18, 19]. An alternative version of a sparse Erdős–Ko–Rado theorem, which we shall discuss in greater detail in Section 1.3, was introduced by Bollobás, Narayanan and Raigorodskii [5].

1.2 Removal lemmas for disjoint pairs

As these new problems go beyond the Erdős–Ko–Rado threshold, we require more robust forms of stability that apply not only to intersecting families, but also to families with few disjoint pairs. This motivated the search for a *removal lemma* that would show one can remove few sets from any family with a small number of disjoint pairs to obtain an intersecting family. Such a result would be a set-theoretic analogue of the graph removal lemmas that have found a wide range of applications in extremal graph theory, details of which are in the survey of Conlon and Fox [6].

Friedgut and Regev [16] proved the first such removal lemma, stated below.

Theorem 1.1 (Friedgut–Regev). *Let $\gamma > 0$, and let k and n be positive integers satisfying $\gamma n \leq k \leq (\frac{1}{2} - \gamma)n$. Then for every $\varepsilon > 0$ there is a $\delta > 0$ such that any family $\mathcal{F} \subset \binom{[n]}{k}$ with at most $\delta |\mathcal{F}| \binom{n-k}{k}$ disjoint pairs can be made intersecting by removing at most $\varepsilon \binom{n-1}{k-1}$ sets from \mathcal{F} .*

This is a very general result that holds regardless of the size or structure of the nearest intersecting family. However, for extremal applications, one is typically interested in the case when $|\mathcal{F}| \approx \binom{n-1}{k-1}$. For example, Gauy, Hàn and Oliveira required such a lemma in [17], coupling Theorem 1.1 with known stability results to show that a family of size close to $\binom{n-1}{k-1}$ with few disjoint pairs must be close in structure to a star. They further asked if such a result also holds for $k = o(n)$. Our main theorem shows this is indeed the case. Theorem 1.2 provides a removal lemma that holds whenever \mathcal{F} has size close to a union of ℓ full stars and has relatively few disjoint pairs. Moreover, when $\ell = 1$, this holds for the full range of $2 \leq k < \frac{n}{2}$.

Theorem 1.2. *There is an absolute constant $C > 1$ such that if n, k and ℓ are positive integers satisfying $n > 2k\ell^2$, and $\mathcal{F} \subset \binom{[n]}{k}$ is a family of size $|\mathcal{F}| = (\ell - \alpha) \binom{n-1}{k-1}$ with at most $\binom{\ell}{2} + \beta \binom{n-1}{k-1} \binom{n-k-1}{k-1}$ disjoint pairs, where $\max(2\ell|\alpha|, |\beta|) \leq \frac{n-2k}{(20C)^2 n}$, then there is a family \mathcal{S} that is the union of ℓ stars satisfying $|\mathcal{F} \Delta \mathcal{S}| \leq C((2\ell - 1)\alpha + 2\beta) \frac{n}{n-2k} \binom{n-1}{k-1}$.*

Another feature of Theorem 1.2 is that, despite its simple proof, it provides quantitative control that is often sharp up to the constant. The distance from \mathcal{F} to a union of ℓ stars is measured in terms of its size (parametrised by α), the number of disjoint pairs (parametrised by β), and how close k is to $\frac{1}{2}n$. When $\ell = 1$, taking $\beta = 0$ gives a stability result for intersecting families, and the bounds sharpen those given by Keevash and Mubayi [23] and Keevash [22].

For positive β , the bounds remain sharp up to the constant. If k is bounded away from $\frac{n}{2}$, then one may take a star and add $\alpha \binom{n-1}{k-1}$ sets from another star to obtain a family of size $(1 + \alpha) \binom{n-1}{k-1}$ with $\alpha \binom{n-1}{k-1} \binom{n-k-1}{k-1}$ disjoint pairs that is $\alpha \binom{n-1}{k-1}$ -far from a star. On the other hand, if $t = n - 2k = o(n)$, consider the anti-star $\binom{[n-1]}{k}$. This has size $(1 + \frac{t}{k}) \binom{n-1}{k-1}$, contains approximately $\frac{t}{n} \binom{n-1}{k-1} \binom{n-k-1}{k-1}$ disjoint pairs, and yet is approximately $\binom{n-1}{k-1}$ -far from a star.

When $\ell \geq 2$, \mathcal{F} is much larger than the Erdős–Ko–Rado bound, and hence we would expect \mathcal{F} to contain many disjoint pairs. Das, Gan and Sudakov [7] have shown that, provided n is sufficiently large, a union of ℓ stars, which has approximately $\binom{\ell}{2} \binom{n-1}{k-1} \binom{n-k-1}{k-1}$ disjoint pairs, minimises the number of disjoint pairs in set families of this size. Theorem 1.2 provides stability for this supersaturation result, showing that families of comparable size with a similar number of disjoint pairs must be close in structure to a union of ℓ stars.

Finally, while we require $n > 2k$ when $\ell = 1$, we can do a bit better when ℓ is large: as ℓ tends to infinity, the bound on n can be lowered to $n > (\frac{1}{2} + o(1)) k\ell^2$.

1.3 Erdős–Ko–Rado for sparse Kneser subgraphs

To demonstrate the usefulness of Theorem 1.2, we shall apply it to a problem of Bollobás, Narayanan and Raigorodskii [5] regarding an extension of the Erdős–Ko–Rado theorem to the sparse random setting. To define the problem at hand, we first need to introduce the Kneser graph and its connection to the Erdős–Ko–Rado theorem.

Given integers $1 \leq k \leq \frac{1}{2}n$, the Kneser graph $K(n, k)$ is defined on the vertex set $V = \binom{[n]}{k}$, with two k -sets $F, G \in \binom{[n]}{k}$ adjacent in $K(n, k)$ if and only if $F \cap G = \emptyset$. Since edges of $K(n, k)$ denote disjoint pairs in $\binom{[n]}{k}$, it follows that independent sets of $K(n, k)$ correspond directly to intersecting families in $\binom{[n]}{k}$. Thus the Erdős–Ko–Rado theorem, viewed from the perspective of the Kneser graph, shows $\alpha(K(n, k)) = \binom{n-1}{k-1}$ when $n \geq 2k$.

Bollobás, Narayanan and Raigorodskii [5] transferred the Erdős–Ko–Rado theorem to the random setting by considering not the entire Kneser graph $K(n, k)$, but rather random subgraphs of it. Given some probability $0 \leq p \leq 1$, let $K_p(n, k)$ denote the subgraph of $K(n, k)$ where every edge is retained independently with probability p . As $K_p(n, k) \subseteq K(n, k)$, we clearly have $\alpha(K_p(n, k)) \geq \alpha(K(n, k)) = \binom{n-1}{k-1}$. They then asked for which p we have equality.

In their paper, they showed the Erdős–Ko–Rado theorem is surprisingly robust when k is not too large with respect to n . In other words, we almost surely have $\alpha(K_p(n, k)) = \binom{n-1}{k-1}$ even for very small probabilities p (and thus very sparse subgraphs of $K(n, k)$). Furthermore, they exhibited a sharp threshold for when this sparse Erdős–Ko–Rado theorem holds.

Theorem 1.3 (Bollobás–Narayanan–Raigorodskii). *Fix $\varepsilon > 0$ and suppose $2 \leq k = o(n^{1/3})$. Let*

$$p_0 = \frac{(k+1) \log n - k \log k}{\binom{n-1}{k-1}}.$$

Then, as $n \rightarrow \infty$,

$$\mathbb{P} \left(\alpha(K_p(n, k)) = \binom{n-1}{k-1} \right) \rightarrow \begin{cases} 0 & \text{if } p \leq (1 - \varepsilon)p_0 \\ 1 & \text{if } p \geq (1 + \varepsilon)p_0 \end{cases}.$$

Moreover, for $p \geq (1 + \varepsilon)p_0$, with high probability the largest independent sets are stars.

While observing that we may take $\varepsilon = O(k^{-1})$, they conjectured that the result should continue to hold provided $k = o(n)$. Partial progress was made by Balogh, Bollobás and Narayanan [2], who showed that for every $\gamma > 0$ there is some constant $c(\gamma) > 0$ such that if $k \leq (\frac{1}{2} - \gamma)n$ and $p \geq \binom{n-1}{k-1}^{-c(\gamma)}$, then $\alpha(K_p(n, k)) = \binom{n-1}{k-1}$ with high probability.

By applying Theorem 1.2, we obtain sharper results for large k , as given in the theorem below. For these larger values of k , it is convenient to present the critical probability in a different form to that of Theorem 1.3; note that p_c below is asymptotically equal to p_0 above when $k = o(n^{1/2})$.

Theorem 1.4. *There is an absolute constant $C > 0$ such that the following holds. Let k and n be integers with $\omega(1) = k \leq \frac{1}{2}(n - 3)$, let $\varepsilon = \omega(k^{-1})$, and set*

$$p_c = \frac{\log \left(n \binom{n-1}{k} \right)}{\binom{n-k-1}{k-1}}.$$

Then, as $n \rightarrow \infty$, $\mathbb{P} \left(\alpha(K_p(n, k)) = \binom{n-1}{k-1} \right) \rightarrow 0$ if $p \leq (1 - \varepsilon)p_c$.

For $k \leq \frac{n}{6C}$, if $p \geq (1 + \varepsilon)p_c$, with high probability $\alpha(K_p(n, k)) = \binom{n-1}{k-1}$ and the stars are the only maximum independent sets. For $k \leq \frac{1}{2}(n - 3)$, the same conclusion holds for $p \geq \frac{2Cn}{n-2k}p_c$.

Theorem 1.4 exhibits a sharp threshold for $k \leq \frac{n}{6C}$, thus extending Theorem 1.3 to k as large as linear in n . Furthermore, when $k \leq (\frac{1}{2} - \gamma)n$, as considered in [2], $\frac{n}{n-2k} \leq (2\gamma)^{-1}$, and so Theorem 1.4 determines the critical probability up to a constant factor. Finally, when k is close to $\frac{1}{2}n$, we find that the sparse version of the Erdős–Ko–Rado theorem still holds for very small edge probabilities; when $k = \frac{1}{2}(n - 3)$, we almost surely have $\alpha(K_p(n, k)) = \binom{n-1}{k-1}$ even for $p = \Omega(n^{-1})$.

1.4 Notation and organisation

We use standard set-theoretic and asymptotic notation. We write $[n]$ for $\{1, 2, \dots, n\}$, while $\binom{X}{k}$ is the family of all k -element subsets of a set X . A non-standard piece of notation is $\text{dp}(\mathcal{F})$, which represents the number of disjoint pairs in a set family \mathcal{F} . Given two functions f and g of some underlying parameter n , if $\lim_{n \rightarrow \infty} f(n)/g(n) = 0$, we write $f = o(g)$ and $g = \omega(f)$. Finally, all logarithms are to the base e .

In Section 2 we prove our removal lemma, Theorem 1.2. We apply this result to the sparse Erdős–Ko–Rado problem in Section 3, where we prove Theorem 1.4. The final section contains some concluding remarks and open problems.

2 The removal lemma

In this section we prove our version of the removal lemma, Theorem 1.2. Our proof combines the work of Lovász [24] on the spectrum of the Kneser graph with an analytic result of Filmus [12] regarding approximations of Boolean functions on $\binom{[n]}{k}$. Before beginning with the proof, we shall introduce the necessary terminology.

Given a family of sets $\mathcal{F} \subset \binom{[n]}{k}$, the *characteristic function* $f : \binom{[n]}{k} \rightarrow \{0, 1\}$ is a Boolean function indicating membership of the family, with $f(F) = 1$ if and only if $F \in \mathcal{F}$. We may embed $\binom{[n]}{k} \subset \{0, 1\}^n$ into the n -dimensional hypercube, and thus think of f as being defined on the k -uniform *slice* of the cube $\{(x_1, \dots, x_n) \in \{0, 1\}^n : \sum_i x_i = k\}$. A function f is *affine* if $f(x_1, x_2, \dots, x_n) = a_0 + \sum_{i \in [n]} a_i x_i$ for some constants a_i , $0 \leq i \leq n$. We will equip this space of functions with the L_2 -norm with respect to the uniform measure on $\binom{[n]}{k}$, defining

$$\|f - g\|^2 = \mathbb{E} \left[|f - g|^2 \right] = \frac{1}{\binom{[n]}{k}} \sum_{F \in \binom{[n]}{k}} |f(F) - g(F)|^2,$$

and say f and g are ε -close if $\|f - g\|^2 \leq \varepsilon$. Finally, to avail of the spectral results, which are traditionally phrased in terms of matrices and vectors, we will abuse notation and identify a function $f : \binom{[n]}{k} \rightarrow \mathbb{R}$ with the vector $f = (f(F))_{F \in \binom{[n]}{k}}$ in $\mathbb{R}^{\binom{[n]}{k}}$. Note that the L_2 -norm above arises from the standard inner product on $\mathbb{R}^{\binom{[n]}{k}}$.

The first step of our proof is the following lemma, which transfers the problem into the analytic framework set up above. The lemma shows that if a set family \mathcal{F} is as in the statement of Theorem 1.2, then its characteristic function can be approximated well by an affine function.

Lemma 2.1. *Let n , k and ℓ be positive integers satisfying $n > 2k$, and let $\mathcal{F} \subset \binom{[n]}{k}$ be a family of size $|\mathcal{F}| = (\ell - \alpha) \binom{n-1}{k-1}$ with at most $\binom{\ell}{2} + \beta$ $\binom{n-1}{k-1} \binom{n-k-1}{k-1}$ disjoint pairs. If $f : \binom{[n]}{k} \rightarrow \{0, 1\}$ is the characteristic function of \mathcal{F} , then $\|f - g\|^2 \leq ((2\ell - 1)\alpha + 2\beta) \frac{k}{n-2k}$ for some affine function $g : \binom{[n]}{k} \rightarrow \mathbb{R}$.*

To prove Lemma 2.1, we require some information on the spectrum of the Kneser graph. Let A denote the adjacency matrix of $K(n, k)$. In his celebrated paper on the Shannon capacity of graphs, Lovász [24] showed the eigenvalues of A are $\lambda_i = (-1)^i \binom{n-k-i}{k-i}$ for $0 \leq i \leq k$. Thus the largest eigenvalue is the degree of the vertices in the regular graph $K(n, k)$, $\lambda_0 = \binom{n-k}{k}$, while the smallest eigenvalue is $\lambda_1 = -\binom{n-k-1}{k-1}$. The second smallest eigenvalue is $\lambda_3 = -\binom{n-k-3}{k-3}$. Furthermore, the λ_0 -eigenspace is one-dimensional, spanned by the constant function. The $(n-1)$ -dimensional λ_1 -eigenspace is spanned by the functions $x_i - \frac{k}{n}$ (note that these functions are linearly dependent, as $\sum_i x_i \equiv k$). Hence the span of the λ_0 - and λ_1 -eigenspaces is precisely the space of affine functions. As A is a real symmetric matrix, its eigenspaces are orthogonal. Armed with these preliminaries, we can prove the lemma.

Proof of Lemma 2.1. Given the characteristic vector f of \mathcal{F} , write $f = f_0 + f_1 + f_2$, where f_0 and f_1 are the projections of f to the λ_0 - and λ_1 -eigenspaces respectively, and $f_2 = f - f_0 - f_1$. By the orthogonality of the eigenspaces, we have $\|f\|^2 = \|f_0\|^2 + \|f_1\|^2 + \|f_2\|^2$. As f is a Boolean function, $\|f\|^2 = \mathbb{E}[f^2] = \mathbb{E}[f] = |\mathcal{F}| / \binom{[n]}{k} = (\ell - \alpha) \frac{k}{n}$. Thus, solving for $\|f_1\|^2$, we find $\|f_1\|^2 =$

$(\ell - \alpha)\frac{k}{n} - \|f_0\|^2 - \|f_2\|^2$. Furthermore, since the λ_0 -eigenspace is spanned by the constant function, $f_0 \equiv \mathbb{E}[f] = (\ell - \alpha)\frac{k}{n}$, and so $\|f_0\|^2 = \mathbb{E}[f_0^2] = (\ell - \alpha)^2\frac{k^2}{n^2}$.

As A is the adjacency matrix of the Kneser graph $K(n, k)$, and f is the characteristic function of the set family \mathcal{F} , it follows that $f^T A f = 2\text{dp}(\mathcal{F})$. Using our bound on the number of disjoint pairs in \mathcal{F} ,

$$\begin{aligned} (\ell^2 - \ell + 2\beta) \binom{n-1}{k-1} \binom{n-k-1}{k-1} &\geq 2\text{dp}(\mathcal{F}) = f^T A f = f_0^T A f_0 + f_1^T A f_1 + f_2^T A f_2 \\ &\geq \lambda_0 f_0^T f_0 + \lambda_1 f_1^T f_1 + \lambda_3 f_2^T f_2. \end{aligned}$$

We divide through by $\binom{n}{k}$ to normalise, obtaining

$$\frac{(\ell^2 - \ell + 2\beta)k}{n} \binom{n-k-1}{k-1} \geq \binom{n-k}{k} \|f_0\|^2 - \binom{n-k-1}{k-1} \|f_1\|^2 - \binom{n-k-3}{k-3} \|f_2\|^2.$$

Dividing by $\binom{n-k-1}{k-1}$, substituting our expressions for $\|f_0\|^2$ and $\|f_1\|^2$, and simplifying gives

$$\begin{aligned} \frac{2\beta k}{n} &\geq \left[1 - \frac{(k-1)(k-2)}{(n-k-1)(n-k-2)} \right] \|f_2\|^2 - \frac{(2\ell-1)\alpha k}{n} + \frac{\alpha^2 k}{n} \\ &= \frac{(n-2k)(n-3)}{(n-k-1)(n-k-2)} \|f_2\|^2 - \frac{(2\ell-1)\alpha k}{n} + \frac{\alpha^2 k}{n} \geq \frac{n-2k}{n} \|f_2\|^2 - \frac{(2\ell-1)\alpha k}{n}. \end{aligned}$$

Rearranging, we deduce $\|f_2\|^2 \leq ((2\ell-1)\alpha + 2\beta)\frac{k}{n-2k}$. Recalling that $f_0 + f_1$ is spanned by the λ_0 - and λ_1 -eigenspaces, and hence affine, setting $g = f_0 + f_1$ gives the desired result. \square

Lemma 2.1 shows the characteristic function of \mathcal{F} must be close to an affine function, from which we shall deduce that \mathcal{F} itself is close to a union of stars. Note that the characteristic function g of the union of stars with centres $i \in S$ is simply $g(x_1, \dots, x_n) = \max_{i \in S} x_i$, and is thus determined only by the coordinates in S . The Friedgut–Kalai–Naor theorem [15] states that if a Boolean function $f : \{0, 1\}^n \rightarrow \{0, 1\}$ on the entire hypercube is close to an affine function, then it is close to a function determined by at most one coordinate. We shall make use of an analogous result for the k -uniform slices of the cube, due to Filmus [12].

Theorem 2.2 (Filmus). *For some constant $C > 1$ the following holds. Suppose $2 \leq k \leq \frac{1}{2}n$ and $\varepsilon < \frac{k}{128n}$. For every Boolean function $f : \binom{[n]}{k} \rightarrow \{0, 1\}$ that is ε -close to an affine function, there is some set $S \subset [n]$ of size $|S| \leq \max\left(1, \frac{Cn\sqrt{\varepsilon}}{k}\right)$ such that either f or $1 - f$ is $(C\varepsilon)$ -close to $\max_{i \in S} x_i$.*

We now have all the necessary ingredients to prove the removal lemma.

Proof of Theorem 1.2. Set $\varepsilon = ((2\ell-1)\alpha + 2\beta)\frac{k}{n-2k}$, and take C as in Theorem 2.2. By our bounds on α and β , $\varepsilon < \frac{k}{128C^2n}$. If \mathcal{F} is as in the statement of the theorem, then by Lemma 2.1 its characteristic function f is ε -close to an affine function. By Theorem 2.2, there is some $S \subset [n]$ such that f or $1 - f$ is $(C\varepsilon)$ -close to $\max_{i \in S} x_i$. Without loss of generality, we may assume $S = [s]$, where $s \leq \max\left(1, \frac{Cn\sqrt{\varepsilon}}{k}\right)$. Let $g_s = \max_{i \in [s]} x_i$, and let $\mathcal{G}_s = \binom{[n]}{k} \setminus \binom{[n] \setminus [s]}{k}$ be the family corresponding to this characteristic function.

Note that $\|f - g_s\|^2 = |\mathcal{F}\Delta\mathcal{G}_s|/\binom{n}{k}$, since for any set $F \in \binom{[n]}{k}$ we have

$$|f(F) - g_s(F)| = \begin{cases} 1 & \text{if } F \in \mathcal{F}\Delta\mathcal{G}_s \\ 0 & \text{otherwise} \end{cases}.$$

Hence we must have $|\mathcal{F}\Delta\mathcal{H}| \leq C\varepsilon\binom{n}{k}$ for $\mathcal{H} = \mathcal{G}_s$ or $\mathcal{H} = \overline{\mathcal{G}_s}$, depending on whether it is f or $1-f$ that is $(C\varepsilon)$ -close to g_s . There are six possibilities to consider:

$$\begin{array}{lll} \text{(i)} & \mathcal{H} = \mathcal{G}_s, s \leq \ell - 1 & \text{(ii)} & \mathcal{H} = \mathcal{G}_s, s \geq \ell + 1 & \text{(iii)} & \mathcal{H} = \overline{\mathcal{G}_0} \\ \text{(iv)} & \mathcal{H} = \overline{\mathcal{G}_s}, s \geq 2 & \text{(v)} & \mathcal{H} = \overline{\mathcal{G}_1} & \text{(vi)} & \mathcal{H} = \mathcal{G}_\ell \end{array}$$

Since \mathcal{G}_ℓ is the union of ℓ stars, we wish to show that (vi) must hold. We first consider the sizes of \mathcal{F} and \mathcal{H} to eliminate cases (i)-(iv). Recall that $|\mathcal{F}| = (\ell - \alpha)\binom{n-1}{k-1}$, and, by our bound on α , $\ell - \alpha \in (\ell - \frac{1}{8}, \ell + \frac{1}{8})$. Since $\|\mathcal{F} - \mathcal{H}\| \leq |\mathcal{F}\Delta\mathcal{H}| \leq C\varepsilon\binom{n}{k} < \frac{1}{8}\binom{n-1}{k-1}$, we must have $(\ell - \frac{1}{4})\binom{n-1}{k-1} \leq |\mathcal{H}| \leq (\ell + \frac{1}{4})\binom{n-1}{k-1}$.

We have $|\mathcal{G}_s| \leq s\binom{n-1}{k-1}$, which is too small if $s \leq \ell - 1$. On the other hand, observe that \mathcal{G}_s , the union of s stars, grows with s . Thus, when $s \geq \ell + 1$,

$$|\mathcal{G}_s| \geq |\mathcal{G}_{\ell+1}| \geq (\ell+1)\binom{n-1}{k-1} - \binom{\ell+1}{2}\binom{n-2}{k-2} \geq \left(\ell+1 - \frac{\ell^2 k}{n}\right)\binom{n-1}{k-1} \geq \left(\ell + \frac{1}{2}\right)\binom{n-1}{k-1},$$

which is too large. This rules out cases (i) and (ii). We also have $|\overline{\mathcal{G}_0}| = \binom{n}{k} = \frac{n}{k}\binom{n-1}{k-1} \geq 2\ell^2\binom{n-1}{k-1}$, which is again too large, ruling out case (iii) as well.

To handle case (iv), we show that $\overline{\mathcal{G}_s}$ is too large when $s \geq 2$. Since $|\overline{\mathcal{G}_s}| = \binom{n-s}{k}$ is decreasing in s , it suffices to take $s = \frac{Cn\sqrt{\varepsilon}}{k}$. We indeed have too many sets, as

$$|\overline{\mathcal{G}_s}| = \binom{n-s}{k} \geq \left(1 - \frac{sk}{n}\right)\binom{n}{k} = (1 - C\sqrt{\varepsilon})\frac{n}{k}\binom{n-1}{k-1} > \frac{3\ell^2}{2}\binom{n-1}{k-1} \geq \left(\ell + \frac{1}{2}\right)\binom{n-1}{k-1}.$$

The above argument does not immediately rule out case (v), since if $s = \max\left(1, \frac{Cn\sqrt{\varepsilon}}{k}\right) = 1$, we may not assume $s = \frac{Cn\sqrt{\varepsilon}}{k}$. However, the family $\overline{\mathcal{G}_1}$ is still too large when $\ell \geq 2$, as

$$|\overline{\mathcal{G}_1}| = \binom{n-1}{k} = \frac{n-k}{k}\binom{n-1}{k-1} \geq (2\ell^2 - 1)\binom{n-1}{k-1} > \left(\ell + \frac{1}{2}\right)\binom{n-1}{k-1}.$$

To rule out case (v) when $\ell = 1$, we consider the number of disjoint pairs in \mathcal{F} . Note that each of the $\binom{n-1}{k}$ sets in $\overline{\mathcal{G}_1}$ is disjoint from $\binom{n-k-1}{k}$ other sets in $\overline{\mathcal{G}_1}$, and hence $\text{dp}(\overline{\mathcal{G}_1}) = \frac{1}{2}\binom{n-1}{k}\binom{n-k-1}{k}$. Moreover, removing t sets from $\overline{\mathcal{G}_1}$ can account for at most $t\binom{n-k-1}{k}$ disjoint pairs. If \mathcal{F} were close to $\overline{\mathcal{G}_1}$, then $|\overline{\mathcal{G}_1} \setminus \mathcal{F}| \leq C\varepsilon\binom{n}{k}$, and so

$$\text{dp}(\mathcal{F}) \geq \text{dp}(\mathcal{F} \cap \overline{\mathcal{G}_1}) \geq \left(\frac{1}{2} - \frac{C\varepsilon n}{n-k}\right)\binom{n-1}{k}\binom{n-k-1}{k} > \left(\frac{1}{2} - 2C\varepsilon\right)\binom{n-1}{k}\binom{n-k-1}{k}.$$

On the other hand, we assumed \mathcal{F} has at most $\beta\binom{n-1}{k-1}\binom{n-k-1}{k-1}$ disjoint pairs, so we must have $\beta \geq \left(\frac{1}{2} - 2C\varepsilon\right)\frac{(n-k)(n-2k)}{k^2} > \frac{n-2k}{2n}$, contradicting our bound on β .

Thus we are only left with case (vi), where \mathcal{H} is the union of ℓ stars \mathcal{G}_ℓ , and, as required, we have $|\mathcal{F}\Delta\mathcal{G}_\ell| \leq C\varepsilon\binom{n}{k} = C((2\ell-1)\alpha + 2\beta)\frac{n}{n-2k}\binom{n-1}{k-1}$. \square

3 Independence number of random Kneser subgraphs

In this section we prove Theorem 1.4, establishing an analogue of the Erdős–Ko–Rado theorem for sparse random subgraphs of the Kneser graph. We will show that below the threshold, there is with high probability an independent *superstar*: some star \mathcal{S} and a set $F \notin \mathcal{S}$ such that $\mathcal{S} \cup \{F\}$ is independent in $K_p(n, k)$. The upper bound on the critical probability essentially follows from a union bound over all potential independent sets, where we will be able to take advantage of the fine control afforded to us by Theorem 1.2 to obtain sharp results when k is large.

Proof of Theorem 1.4. First we establish the lower bound on the critical probability. Suppose $\varepsilon = \omega(k^{-1})$ and $p \leq (1 - \varepsilon)p_c$. We wish to show that with high probability, stars can be extended to independent superstars in $K_p(n, k)$. Let \mathcal{S} be the star with centre 1, and for every $1 \notin F \in \binom{[n]}{k}$ let \mathcal{E}_F be the event that $\mathcal{S} \cup \{F\}$ is independent in $K_p(n, k)$.

Note that F is disjoint from $\binom{n-k-1}{k-1}$ sets in \mathcal{S} , and for \mathcal{E}_F to hold none of these edges can appear in $K_p(n, k)$. Thus $\mathbb{P}(\mathcal{E}_F) = (1 - p)^{\binom{n-k-1}{k-1}}$. Moreover, the events $\{\mathcal{E}_F : 1 \notin F\}$ depend on mutually disjoint sets of edges of $K(n, k)$, and are thus independent. Hence we can bound the probability that the stars are the largest independent sets of $K_p(n, k)$ by

$$\mathbb{P}\left(\alpha(K_p(n, k)) = \binom{n-1}{k-1}\right) \leq \mathbb{P}\left(\bigcap_F \overline{\mathcal{E}_F}\right) = \left(1 - (1 - p)^{\binom{n-k-1}{k-1}}\right)^{\binom{n-1}{k}}.$$

This bound is increasing in p , so it suffices to take $p = (1 - \varepsilon)p_c = \frac{(1 - \varepsilon) \log\left(n \binom{n-1}{k}\right)}{\binom{n-k-1}{k-1}}$. As $n \geq 2k + 2$, $p = O(n^{-1}) = o(\varepsilon)$, and hence $(1 - p)^{\binom{n-k-1}{k-1}} \geq \exp\left(-p(1 + p)\binom{n-k-1}{k-1}\right) \geq \left(n \binom{n-1}{k}\right)^{-(1 - \varepsilon/2)}$. Thus

$$\left(1 - (1 - p)^{\binom{n-k-1}{k-1}}\right)^{\binom{n-1}{k}} \leq \exp\left(-\binom{n-1}{k}(1 - p)^{\binom{n-k-1}{k-1}}\right) \leq \exp\left(-n^{-1}\binom{n-1}{k}^{\varepsilon/2}\right) = o(1),$$

since $\varepsilon = \omega(k^{-1})$. Hence for $p \leq (1 - \varepsilon)p_c$ we have $\alpha(K_p(n, k)) > \binom{n-1}{k-1}$ with high probability.

We now seek an upper bound on the critical probability. By monotonicity, it suffices to consider p as small as possible. To begin, we will prove the coarse threshold. Let C be the (absolute) constant from Theorem 1.2, and take $p = \zeta p_c$, where $\zeta = \frac{2Cn}{n-2k}$. For such p , we wish to show the only maximum independent sets of $K_p(n, k)$ are the stars. To this end, we define the following random variables:

$$X = |\{\text{independent superstars } \mathcal{F} : \mathcal{F} = \mathcal{S} \cup \{F\} \text{ for some star } \mathcal{S}, F \notin \mathcal{S}\}| \text{ and}$$

$$Y_i = \left| \left\{ \text{independent } \mathcal{F} : |\mathcal{F}| = \binom{n-1}{k-1}, \min_{\mathcal{S} \text{ a star}} |\mathcal{S} \setminus \mathcal{F}| = i \right\} \right|, 1 \leq i \leq \binom{n-1}{k-1}.$$

X counts the number of independent superstars in $K_p(n, k)$. If $X = 0$, the stars are all maximal independent sets. If we further have $Y_i = 0$ for all $1 \leq i \leq \binom{n-1}{k-1}$, then there are no independent sets of size $\binom{n-1}{k-1}$ that are not stars, and thus the stars are the only maximum independent sets in $K_p(n, k)$. Hence our task is to show $X + \sum_i Y_i = 0$ with high probability. By the union bound, it suffices to show $\mathbb{P}(X > 0) + \sum_i \mathbb{P}(Y_i > 0) = o(1)$.

We begin by estimating $\mathbb{P}(X > 0)$, which we can bound by $\mathbb{E}[X]$. There are n choices for the star \mathcal{S} , $\binom{n-1}{k}$ choices for the set $F \notin \mathcal{S}$, and, for each such configuration, $\binom{n-k-1}{k-1}$ edges that should not appear in $K_p(n, k)$, which occurs with probability $(1-p)^{\binom{n-k-1}{k-1}} \leq \exp\left(-\zeta p_c \binom{n-k-1}{k-1}\right)$. Thus

$$\mathbb{E}[X] \leq n \binom{n-1}{k} \exp\left(-\zeta p_c \binom{n-k-1}{k-1}\right) = \left(n \binom{n-1}{k}\right)^{1-\zeta} = o(1), \quad (1)$$

even for ζ as small as $1 + \omega(k^{-1})$.

To analyse $\mathbb{P}(Y_i > 0)$, we shall distinguish between two different cases: families that are close to a star, and families far from a star. For the first case, we assume $1 < i \leq t_1 = \frac{1}{400C} \binom{n-1}{k-1}$. The families \mathcal{F} counted by Y_i have size $\binom{n-1}{k-1}$ and $|\mathcal{F} \Delta \mathcal{S}| = 2|\mathcal{S} \setminus \mathcal{F}| \geq 2i$ for every star \mathcal{S} . By applying Theorem 1.2 with $\alpha = 0$ and $\beta = \frac{i(n-2k)}{Cn \binom{n-1}{k-1}}$, it follows that $\text{dp}(\mathcal{F}) \geq \frac{i(n-2k)}{Cn} \binom{n-k-1}{k-1}$. For \mathcal{F} to be independent in $K_p(n, k)$, none of these edges can appear, which occurs with probability $(1-p)^{\text{dp}(\mathcal{F})} \leq \left(n \binom{n-1}{k}\right)^{-\zeta i(n-2k)/(Cn)}$.

We now take a union bound over all possible choices of \mathcal{F} . We know there is some star \mathcal{S} such that $|\mathcal{S} \setminus \mathcal{F}| = i$. There are n choices for the star \mathcal{S} , $\binom{n-1}{i}$ choices for the i sets in $\mathcal{S} \setminus \mathcal{F}$, and $\binom{n-1}{i}$ choices for the i sets in $\mathcal{F} \setminus \mathcal{S}$. Hence there are at most $n \binom{n-1}{i} \binom{n-1}{i} \leq n \left(\frac{ke^2 \binom{n-1}{k}^2}{(n-k)i^2}\right)^i$ families \mathcal{F} that can be counted by Y_i . Thus we have

$$\sum_{i=1}^{t_1} \mathbb{P}(Y_i > 0) \leq \sum_{i=1}^{t_1} n \left(\frac{ke^2 \binom{n-1}{k}^2}{(n-k)i^2} \left(n \binom{n-1}{k}\right)^{-\zeta(n-2k)/(Cn)}\right)^i \leq \sum_{i=1}^{t_1} \frac{ne^{2i}}{(ni)^{2i}} = o(1). \quad (2)$$

where the second inequality follows from our choice of $\zeta = \frac{2Cn}{n-2k}$.

Finally, we bound $\mathbb{P}(Y_i > 0)$ when $i > t_1$. Applying Theorem 1.2 with $\alpha = 0$ and $\beta = \frac{n-2k}{(20C)^2 n}$, any family \mathcal{F} counted by $\sum_{i>t_1} Y_i$ must have $\text{dp}(\mathcal{F}) \geq \frac{n-2k}{(20C)^2 n} \binom{n-1}{k-1} \binom{n-1}{k-1}$. Hence the probability of such an \mathcal{F} being independent in $K_p(n, k)$ is $(1-p)^{\text{dp}(\mathcal{F})} \leq \left(n \binom{n-1}{k}\right)^{-\frac{\zeta(n-2k)}{(20C)^2 n} \binom{n-1}{k-1}}$. Recalling our choice of $\zeta = \frac{2Cn}{n-2k}$, we apply a trivial union bound, summing over all $\binom{n}{\binom{n-1}{k-1}} \leq \left(\frac{ne}{k}\right)^{\binom{n-1}{k-1}}$ families of size $\binom{n-1}{k-1}$ to find, when $k \geq 600C$,

$$\sum_{i=t_1+1}^{\binom{n-1}{k-1}} \mathbb{P}(Y_i > 0) \leq \left(\frac{ne}{k} \left(n \binom{n-1}{k}\right)^{-1/(200C)}\right)^{\binom{n-1}{k-1}} \leq \left(\frac{ne}{k} \left(\frac{k}{n}\right)^{k/(200C)}\right)^{\binom{n-1}{k-1}} = o(1). \quad (3)$$

Combining (1), (2), and (3), we find that when $p \geq \zeta p_c$, $\mathbb{P}(X > 0) + \sum_i \mathbb{P}(Y_i > 0) = o(1)$, and so for such p , the maximum independent sets in $K_p(n, k)$ are precisely the stars.

We now prove the sharp threshold result, for which we must show that the same conclusion holds when $k \leq \frac{n}{6C}$ and $\zeta = 1 + \varepsilon$, for some small $\varepsilon = \omega(k^{-1})$. As previously stated, the bound from (1) holds with this smaller value of ζ . However, bounding $\sum_i \mathbb{P}(Y_i > 0)$ requires more careful analysis. We now split the sum into three parts: $1 \leq i \leq t_0 = \frac{\varepsilon}{2} \binom{n-k-1}{k-1}$, $t_0 + 1 \leq i \leq t_1 = \frac{1}{400C} \binom{n-1}{k-1}$, and $t_1 + 1 \leq i \leq \binom{n-1}{k-1}$.

For the latter two parts, we modify slightly our analysis of the above bounds. When a family is only moderately close to a star, with $t_0 + 1 \leq i \leq t_1$, we again use the bound in (2). We begin by observing that $i > t_0 = \frac{\varepsilon}{2} \binom{n-k-1}{k-1} = \frac{\varepsilon k}{2(n-k)} \binom{n-k}{k}$, and so

$$\frac{ke^2 \binom{n-1}{k}^2}{(n-k)i^2} \leq \frac{4e^2(n-k) \binom{n-1}{k}^2}{\varepsilon^2 k \binom{n-k}{k}^2} \leq \frac{4e^2(n-k)}{\varepsilon^2 k} \left(1 + \frac{k}{n-2k}\right)^{2k} \leq \frac{4e^2(n-k)e^{k/(2C)}}{\varepsilon^2 k}.$$

Since $k \leq \frac{n}{6C}$, we may also bound $(n \binom{n-1}{k})^{-\zeta(n-2k)/(Cn)} \leq \left(\frac{k}{n}\right)^{k/(2C)}$, and so the bases of the exponential summands in (2) are at most

$$\frac{4e^2(n-k)}{\varepsilon^2 k} \left(\frac{ek}{n}\right)^{k/(2C)} \leq \frac{4e^3}{\varepsilon^2} \left(\frac{e}{6C}\right)^{k/(2C)-1} \leq \frac{4e^3}{\varepsilon^2 C^{k/(2C)-1}} \left(\frac{e}{6}\right)^2 < \frac{e^3}{\varepsilon^2 C^{k/(3C)}} < \frac{1}{2},$$

as we may assume $k \geq 6C$ and $\varepsilon \geq 10C^{-k/(6C)}$. This then implies $\sum_{i=t_0+1}^{t_1} \mathbb{P}(Y_i > 0) = o(1)$.

For families that are far from a star, we can re-estimate the upper bound in (3) to show

$$\sum_{i>t_1} \mathbb{P}(Y_i > 0) \leq \left(\frac{ne}{k} \binom{n-1}{k}\right)^{-\frac{\zeta(n-2k)}{(20C)^2 n}} \binom{n-1}{k-1} \leq \left(\frac{ne}{k} \left(\frac{k}{n}\right)^{\frac{k}{2(20C)^2}}\right)^{\binom{n-1}{k-1}} \leq \left(\frac{ek}{n}\right)^{\binom{n-1}{k-1}} = o(1),$$

assuming $k \geq (40C)^2$.

To complete the proof of the sharp threshold, we must demonstrate that we are unlikely to obtain independent families that are very close to stars, with $1 \leq i \leq t_0$. In this range, we repeat the analysis of Bollobás, Narayanan and Raigorodskii in [5], and instead consider *maximal* independent families in $K_p(n, k)$.

For $j \geq i \geq 0$, let $Z_{i,j}$ denote the number of maximal independent families \mathcal{F} such that there is some star \mathcal{S} with $|\mathcal{S} \setminus \mathcal{F}| = i$ and $|\mathcal{F} \setminus \mathcal{S}| = j$. Observe that if \mathcal{F} is a family counted by the random variable Y_i , then for any maximal independent family \mathcal{F}' containing \mathcal{F} , the family \mathcal{F}' must be counted by $Z_{i',j}$ for some $i' \leq i$ and $j \geq i$. Furthermore, if $i' = 0$, this family \mathcal{F}' contains superstars. This shows

$$\bigcup_{1 \leq i \leq t_0} \{Y_i > 0\} \subseteq \{X > 0\} \cup \bigcup_{\substack{1 \leq i \leq t_0 \\ j \geq i}} \{Z_{i,j} > 0\},$$

and so $\sum_{i=1}^{t_0} \mathbb{P}(Y_i > 0) \leq \mathbb{P}(X > 0) + \sum_{i=1}^{t_0} \sum_{j \geq i} \mathbb{P}(Z_{i,j} > 0)$. We already have $\mathbb{P}(X > 0) = o(1)$, and hence it suffices to show $\sum_{i=1}^{t_0} \sum_{j \geq i} \mathbb{E}[Z_{i,j}] = o(1)$.

Let \mathcal{F} be a maximal independent family counted by $Z_{i,j}$. Let \mathcal{S} be the corresponding star, $\mathcal{A} = \mathcal{S} \setminus \mathcal{F}$, and $\mathcal{B} = \mathcal{F} \setminus \mathcal{S}$. Thus we have $|\mathcal{A}| = i$ and $|\mathcal{B}| = j$. By virtue of \mathcal{F} being independent, all of the edges between \mathcal{B} and $\mathcal{S} \setminus \mathcal{A}$ must be missing in $K_p(n, k)$. As \mathcal{F} is maximal, each $A \in \mathcal{A}$ must have an edge to some $B \in \mathcal{B}$, for otherwise $\mathcal{F} \cup \{A\}$ would be a larger independent family. In particular, this implies \mathcal{B} is a subset of the union of the neighbourhoods in $K(n, k)$ of $A \in \mathcal{A}$.

There are thus n choices for the star \mathcal{S} , $\binom{n-1}{i}$ choices for \mathcal{A} , and at most $\binom{i \binom{n-k}{j}}{j}$ choices for \mathcal{B} . Each $A \in \mathcal{A}$ must retain at least one of its edges to \mathcal{B} , which occurs with probability at most

jp , independently for each of the i sets. Furthermore, as every $B \in \mathcal{B}$ has $\binom{n-k-1}{k-1}$ neighbours in \mathcal{S} , there are at least $j \left(\binom{n-k-1}{k-1} - i \right)$ edges between \mathcal{B} and $\mathcal{S} \setminus \mathcal{A}$ that must be missing. This gives

$$\mathbb{E}[Z_{i,j}] \leq n \binom{\binom{n-1}{k-1}}{i} \binom{\binom{n-k}{k}}{j} (jp)^i (1-p)^{j(\binom{n-k-1}{k-1}-i)} = z_{i,j}.$$

We first observe that, for $i \leq t_0 = \frac{\varepsilon}{2} \binom{n-k-1}{k-1}$ and $j \geq i$,

$$\begin{aligned} \frac{z_{i,j+1}}{z_{i,j}} &= \frac{\binom{\binom{n-k}{k}}{j+1} (j+1)^i}{\binom{\binom{n-k}{k}}{j} j^i} (1-p)^{\binom{n-k-1}{k-1}-i} \leq \frac{\left(\binom{\binom{n-k}{k}}{j} - j \right) e}{j+1} (1-p)^{\left(1-\frac{\varepsilon}{2}\right) \binom{n-k-1}{k-1}} \\ &\leq e \binom{n-k}{k} \left(n \binom{n-1}{k} \right)^{-(1+\frac{\varepsilon}{4})} = o(1), \end{aligned}$$

since $p = \zeta p_c = \frac{(1+\varepsilon) \log(n \binom{n-1}{k})}{\binom{n-k-1}{k-1}}$. Hence $\sum_{j \geq i} z_{i,j} \leq 2z_{i,i}$. Thus

$$\begin{aligned} \sum_{i=1}^{t_0} \sum_{j \geq i} \mathbb{E}[Z_{i,j}] &\leq 2 \sum_{i=1}^{t_0} z_{i,i} = 2n \sum_{i=1}^{t_0} \binom{\binom{n-1}{k-1}}{i} \binom{\binom{n-k}{k}}{i} (ip)^i (1-p)^{i(\binom{n-k-1}{k-1}-i)} \\ &\leq 2n \sum_{i=1}^{t_0} \left(e^2 \binom{n-1}{k-1} \binom{n-k}{k} p (1-p)^{\left(1-\frac{\varepsilon}{2}\right) \binom{n-k-1}{k-1}} \right)^i \\ &\leq 2n \sum_{t=1}^{t_0} \left(e^2 \binom{n-1}{k-1} \binom{n-k}{k} \frac{(1+\varepsilon) \log(n \binom{n-1}{k})}{\binom{n-k-1}{k-1}} \left(n \binom{n-1}{k} \right)^{-(1+\frac{\varepsilon}{4})} \right)^i \\ &= 2n \sum_{i=1}^{t_0} \left(\frac{(1+\varepsilon) e^2 \log(n \binom{n-1}{k})}{n \binom{n-1}{k}^{\frac{\varepsilon}{4}}} \right)^i = o(1), \end{aligned}$$

completing the proof for the sharp threshold. \square

4 Concluding remarks

In this paper, we built on the work of Filmus [12] to develop a removal lemma for large set families with few disjoint pairs. We then used this to determine the threshold for random Kneser subgraphs having the Erdős–Ko–Rado property, thus answering a question of Bollobás, Narayanan and Raigorodskii [5].

Rather than the probabilistic problem considered above, one might instead ask the corresponding extremal question: how sparse can a spanning subgraph G of $K(n, k)$ be if $\alpha(G) = \binom{n-1}{k-1}$? A lower bound can be obtained by requiring the stars to be maximal independent sets. For every set $F \in \binom{[n]}{k}$, and every element $x \notin F$, there must be an edge between F and the star \mathcal{S}_x centred at x , for otherwise $\mathcal{S}_x \cup \{F\}$ would be an independent set of size $\binom{n-1}{k-1} + 1$. As each edge $\{F, F'\}$ covers k stars, it follows that F must have degree at least $\frac{n-k}{k}$, and hence $G \subseteq K(n, k)$ must have at least $\frac{n-k}{2k} \binom{n}{k}$ edges.

Perhaps surprisingly, this simple lower bound can be tight. If k divides n , then Baranyai's Theorem [4] gives a partition of $\binom{[n]}{k}$ into perfect matchings. In the Kneser graph, this corresponds to a partition of the vertices into cliques of size $\frac{n}{k}$. Let G be the subgraph consisting only of these cliques. Any independent set in G can contain at most one vertex from each clique, and hence $\alpha(G) \leq \frac{k}{n} \binom{n}{k} = \binom{n-1}{k-1}$. Furthermore, G is $\frac{n-k}{k}$ -regular, matching the lower bound given previously.

Theorem 1.4 shows that for this bound on the independence number to hold in random graphs, they must be denser by a factor of at least $\log \binom{n-1}{k-1}$. However, these random graphs have the additional property that the only maximum independent sets are the stars, which is not the case in the construction given above. One might be interested in the extremal problem with this stricter requirement, or in the case when k does not divide n .

Returning to the random setting, Devlin and Kahn [9] have recently established threshold results when $k \sim \frac{n}{2}$. It remains to exhibit a sharp threshold around p_c for $k > \frac{n}{6C}$. We believe that, perhaps for smaller k , a more precise hitting time result may hold. Consider the random process where one removes edges from the Kneser graph $K(n, k)$ one at a time, selecting at each step an edge uniformly at random from those that remain. Is it true that, with high probability, $\alpha(G) > \binom{n-1}{k-1}$ precisely when a superstar is born? The fact that the lower bound from the sharp threshold comes from these superstars suggests this might be the case.

More generally, given how central intersecting families are to extremal set theory, we believe the removal lemma should find many other applications. In a forthcoming paper with Balogh, Liu and Sharifzadeh, we obtain some supersaturation results using the removal lemma with $\ell \geq 1$, extending the results of [7]. We hope that the lemma might prove useful for other research directions as well.

Acknowledgements We would like to thank Jozsef Balogh, Hong Liu and Maryam Sharifzadeh for suggesting the generalisation of the removal lemma to larger families, with $\ell \geq 2$. We also extend our thanks to the anonymous referees for their many helpful suggestions.

References

- [1] J. Balogh, T. Bohman and D. Mubayi, *Erdős–Ko–Rado in random hypergraphs*, *Combin. Probab. Comput.* **18.5** (2009), 629–646.
- [2] J. Balogh, B. Bollobás and B. P. Narayanan, *Transference for the Erdős–Ko–Rado theorem*, *Forum of Mathematics, Sigma* **3** (2015), e23.
- [3] J. Balogh, S. Das, M. Delcourt, H. Liu and M. Sharifzadeh, *The typical structure of intersecting families of discrete structures*, *J. Combin. Theory Ser. A* **132** (2015), 224–245.
- [4] Z. Baranyai, *On the Factorization of the Complete Uniform Hypergraph*, *Infinite and finite sets (Colloq. Keszthely, 1973; dedicated to P. Erdős on his 60th birthday)* **1** (1975), 91–108.
- [5] B. Bollobás, B. P. Narayanan and A. M. Raigorodskii, *On the stability of the Erdős–Ko–Rado theorem*, *J. Combin. Theory Ser. A* **137** (2016), 64–78.
- [6] D. Conlon and J. Fox, *Graph removal lemmas*, *Surveys in Combinatorics*, Cambridge University Press, (2013), 1–50.
- [7] S. Das, W. Gan and B. Sudakov, *The minimum number of disjoint pairs in set systems and related problems*, *Combinatorica*, to appear.

- [8] S. Das and B. Sudakov, *Most probably intersecting hypergraphs*, Electron. J. Comb. **22** (2015), P1.80.
- [9] P. Devlin and J. Kahn, *On “stability” in the Erdős–Ko–Rado theorem*, arXiv:1502.05692.
- [10] I. Dinur and E. Friedgut, *Intersecting families are essentially contained in juntas*, Combin. Probab. Comput. **18.1-2** (2009), 107–122.
- [11] P. Erdős, C. Ko and R. Rado, *Intersection theorems for systems of finite sets*, Q. J. Math. **12.1** (1961), 313–320.
- [12] Y. Filmus, *Friedgut–Kalai–Naor theorem for slices of the Boolean cube*, arXiv:1410.7834.
- [13] P. Frankl, *Erdős–Ko–Rado theorem with conditions on the maximal degree*, J. Combin. Theory Ser. A **46** (1987), 252–263.
- [14] E. Friedgut, *On the measure of intersecting families, uniqueness and stability*, Combinatorica **28.5** (2008), 503–528.
- [15] E. Friedgut, G. Kalai and A. Naor, *Boolean functions whose Fourier transform is concentrated on the first two levels*, Adv. App. Math. **29.3** (2002), 427–437.
- [16] E. Friedgut and O. Regev, *personal communication*.
- [17] M. M. Gao, H. Hàn and I. C. Oliveira, *Erdős–Ko–Rado for random hypergraphs: asymptotics and stability*, arXiv:1409.3634.
- [18] A. Hamm and J. Kahn, *On Erdős–Ko–Rado for random hypergraphs I*, arXiv:1412.5085.
- [19] A. Hamm and J. Kahn, *On Erdős–Ko–Rado for random hypergraphs II*, arXiv:1406.5793.
- [20] A. J. W. Hilton and E. C. Milner, *Some intersection theorems for systems of finite sets*, Q. J. Math. **18.1** (1967), 369–384.
- [21] G. O. H. Katona, G. Y. Katona and Z. Katona, *Most probably intersecting families of subsets*, Combin. Probab. Comput. **21.1-2** (2012), 219–227.
- [22] P. Keevash, *Shadows and intersections: stability and new proofs*, Adv. Math. **218** (2008), 1685–1703.
- [23] P. Keevash and D. Mubayi, *Set systems without a simplex or a cluster*, Combinatorica **30.2** (2010), 175–200.
- [24] L. Lovász, *On the Shannon capacity of a graph*, IEEE T. Inform. Theory **25.1** (1979), 1–7.
- [25] P. A. Russell, *Compressions and probably intersecting families*, Combin. Probab. Comput. **21.1-2** (2012), 301–313.
- [26] P. A. Russell and M. Walters, *Probably intersecting families are not nested*, Combin. Probab. Comput. **22.1** (2013), 146–160.