

INDEPENDENT SETS IN HYPERGRAPHS

JÓZSEF BALOGH, ROBERT MORRIS, AND WOJCIECH SAMOTIJ

ABSTRACT. Many important theorems and conjectures in combinatorics, such as the theorem of Szemerédi on arithmetic progressions and the Erdős-Stone Theorem in extremal graph theory, can be phrased as statements about families of independent sets in certain uniform hypergraphs. In recent years, an important trend in the area has been to extend such classical results to the so-called ‘sparse random setting’. This line of research has recently culminated in the breakthroughs of Conlon and Gowers and of Schacht, who developed general tools for solving problems of this type. Although these two papers solved very similar sets of longstanding open problems, the methods used are very different from one another and have different strengths and weaknesses.

In this paper, we provide a third, completely different approach to proving extremal and structural results in sparse random sets that also yields their natural ‘counting’ counterparts. We give a structural characterization of the independent sets in a large class of uniform hypergraphs by showing that every independent set is almost contained in one of a small number of relatively sparse sets. We then derive many interesting results as fairly straightforward consequences of this abstract theorem. In particular, we prove the well-known conjecture of Kohayakawa, Łuczak, and Rödl, a probabilistic embedding lemma for sparse graphs. We also give alternative proofs of many of the results of Conlon and Gowers and Schacht, such as sparse random versions of Szemerédi’s theorem, the Erdős-Stone Theorem, and the Erdős-Simonovits Stability Theorem, and obtain their natural ‘counting’ versions, which in some cases are considerably stronger. For example, we show that for each positive β and integer k , there are at most $\binom{\beta n}{m}$ sets of size m that contain no k -term arithmetic progression, provided that $m \geq Cn^{1-1/(k-1)}$, where C is a constant depending only on β and k . We also obtain new results, such as a sparse version of the Erdős-Frankl-Rödl Theorem on the number of H -free graphs and, as a consequence of the KLR conjecture, we extend a result of Rödl and Ruciński on Ramsey properties in sparse random graphs to the general, non-symmetric setting.

1. INTRODUCTION

A great many of the central questions in combinatorics fall into the following general framework: Given a finite set V and a collection $\mathcal{H} \subseteq \mathcal{P}(V)$ of *forbidden structures*, what can be said about sets $I \subseteq V$ that do not contain any member of \mathcal{H} ? For example, the celebrated theorem of Szemerédi [62] states that if $V = \{1, \dots, n\}$ and \mathcal{H} is the collection of k -term arithmetic progressions in $\{1, \dots, n\}$, then every set I that contains no member of \mathcal{H} satisfies $|I| = o(n)$. The archetypal problem studied in extremal graph theory, dating back

Date: May 16, 2014.

Research supported in part by: (JB) NSF CAREER Grant DMS-0745185, UIUC Campus Research Board Grant 11067, and OTKA Grant K76099; (RM) CNPq bolsa de Produtividade em Pesquisa; (WS) ERC Advanced Grant DMMCA and a Trinity College JRF.

to the work of Turán [64] and Erdős and Stone [20], is the problem of characterizing such sets I when V is the edge set of the complete graph on n vertices and \mathcal{H} is the collection of copies of some fixed graph H in K_n . In this setting, a great deal is known, not only about the maximum size of I that contains no member of \mathcal{H} , but also what the largest such sets look like, how many such sets there are, and what the structure of a typical such set is.

A collection $\mathcal{H} \subseteq \mathcal{P}(V)$ as above is usually referred to as a *hypergraph* on the vertex set V and any set $I \subseteq V$ that contains no element (*edge*) of \mathcal{H} is called an *independent set*. Therefore, one might say that a large part of extremal combinatorics is concerned with studying independent sets in various specific hypergraphs. We might add here that in many natural settings, such as the two mentioned above, the hypergraphs considered are *uniform*, that is, all edges of \mathcal{H} have the same size.

Although it might at first seem somewhat artificial to study concrete questions in such an abstract setting, the past few years have proved that taking such a general approach can be highly beneficial. The recently-proved general transference theorems of Conlon and Gowers [14] and Schacht [58] (see also [23]), which imply, among other things, sparse random analogues of the classical theorems of Szemerédi and of Erdős and Stone, were stated in the language of hypergraphs. Roughly speaking, these transference theorems say the following: Let \mathcal{H} be a hypergraph whose edges are sufficiently ‘uniformly distributed’. Then the independence number of \mathcal{H} is ‘well-behaved’ with respect to taking subhypergraphs induced by (sufficiently dense) random subsets of the vertex set. More precisely, given $p \in [0, 1]$ and a finite set V , we shall write V_p to denote the *p-random subset* of V , that is, the random subset of V in which each element of V is included with probability p , independently of all other elements. We write $\alpha(\mathcal{H})$ and $v(\mathcal{H})$ to denote the size of the largest independent set and the number of vertices in a hypergraph \mathcal{H} , respectively. The results of Conlon and Gowers [14] and Schacht [58] imply, in particular, that if the distribution of the edges of some uniform hypergraph \mathcal{H} is sufficiently ‘balanced’, then with probability tending to 1 as $v(\mathcal{H}) \rightarrow \infty$,

$$\alpha(\mathcal{H}[V(\mathcal{H})_p]) \leq p\alpha(\mathcal{H}) + o(pv(\mathcal{H})),$$

provided that p is sufficiently large.

In this work, we give an approximate structural characterization of the family of all independent sets in uniform hypergraphs whose edge distribution satisfies a certain natural boundedness condition. More precisely, we shall prove that the family $\mathcal{I}(\mathcal{H})$ of independent sets of such a hypergraph \mathcal{H} exhibits a certain clustering phenomenon. Our main result (Theorem 2.2, below) states that $\mathcal{I}(\mathcal{H})$ admits a partition into relatively few classes with the following property: all members of each class are essentially contained in a single ‘almost independent’ subset of $V(\mathcal{H})$ (i.e., one which contains only a tiny proportion of all the edges of \mathcal{H}). This somewhat abstract statement has surprisingly many deep and interesting consequences, some of which we list in the remainder of this section. We remark that Theorem 2.2 was partly inspired by the work of Kleitman and Winston [38], who implicitly considered a statement of this type in the setting of graphs (2-uniform hypergraphs) and subsequently used it to bound the number of n -vertex graphs without a 4-cycle. We also note that a result similar to Theorem 2.2 was independently proved by Saxton and Thomason [57], who also use it to derive many of the statements that we present in Sections 1.1– 1.4.

1.1. The number of sets with no k -term arithmetic progression. The celebrated theorem of Szemerédi [62] says that for every $k \in \mathbb{N}$, the largest subset of $\{1, \dots, n\}$ that contains no k -term arithmetic progression (AP) has $o(n)$ elements. It immediately follows that there are only $2^{o(n)}$ subsets of $\{1, \dots, n\}$ with no k -term AP. Our first result can be viewed as a sparse analogue of this statement.

Theorem 1.1. *For every positive β and every $k \in \mathbb{N}$, there exist constants C and n_0 such that the following holds. For every $n \in \mathbb{N}$ with $n \geq n_0$, if $m \geq Cn^{1-1/(k-1)}$, then there are at most*

$$\binom{\beta n}{m}$$

m -subsets of $\{1, \dots, n\}$ that contain no k -term AP.

We shall deduce Theorem 1.1 from our main theorem, Theorem 2.2, and a robust version of Szemerédi's theorem, see Section 4. The sparse random analogue of Szemerédi's theorem, proved by Schacht [58] and independently by Conlon and Gowers [14], follows as an easy corollary of Theorem 1.1. Following [14], we shall say that a set $A \subseteq \mathbb{N}$ is (δ, k) -Szemerédi if every subset $B \subseteq A$ with at least $\delta|A|$ elements contains a k -term AP. For the sake of brevity, let $[n] = \{1, \dots, n\}$ and recall that $[n]_p$ denotes the p -random subset of $[n]$.

Corollary 1.2. *For every $\delta \in (0, 1)$ and every $k \in \mathbb{N}$, there exists a constant C such that the following holds. If $p_n \geq Cn^{-1/(k-1)}$ for all sufficiently large n , then*

$$\lim_{n \rightarrow \infty} \mathbb{P}([n]_{p_n} \text{ is } (\delta, k)\text{-Szemerédi}) = 1.$$

We remark that Theorem 1.1 and Corollary 1.2 are both sharp up to the value of the constant C , see the discussion in Section 4, where both of these statements are proved.

Our main result has a variety of other applications in additive combinatorics, see for example [1, 2] where, jointly with Alon, we used a much simpler version of it to count sum-free sets of fixed size in various Abelian groups and the set $[n]$. In Section 4, we shall mention two other applications: generalizations of Theorem 1.1 to higher dimensions and to k -term APs whose common difference is of the form d^r . In each case, the random version (which was proved in [14, 58]) follows as an easy corollary.

1.2. Turán's problem in random graphs. The famous theorem of Erdős and Stone [20] states that the maximum number of edges in an H -free graph on n vertices, the *Turán number for H* , denoted $\text{ex}(n, H)$, satisfies

$$\text{ex}(n, H) = \left(1 - \frac{1}{\chi(H) - 1} + o(1)\right) \binom{n}{2}, \quad (1)$$

where $\chi(H)$ is the chromatic number of H . The analogue of this theorem for the Erdős-Rényi random graph $G(n, p)$ was first studied by Babai, Simonovits, and Spencer [4], who proved that *asymptotically almost surely* (a.a.s. for short), i.e., with probability tending to 1 as $n \rightarrow \infty$, the largest triangle-free subgraph of $G(n, 1/2)$ is bipartite, and by Frankl and Rödl [22], who proved that if $p \geq n^{-1/2+\varepsilon}$ then a.a.s. the largest triangle-free subgraph of $G(n, p)$ has $pn^2/8 + o(pn^2)$ edges. The systematic study of the Turán problem in $G(n, p)$

was initiated by Haxell, Kohayakawa, and Łuczak [34, 35] and by Kohayakawa, Łuczak, and Rödl [42], who posed the following problem. For a fixed graph H , determine necessary and sufficient conditions on a sequence $\mathbf{p} \in [0, 1]^{\mathbb{N}}$ of probabilities such that, a.a.s.,

$$\text{ex}(G(n, p_n), H) = \left(1 - \frac{1}{\chi(H) - 1} + o(1)\right) \binom{n}{2} p_n, \quad (2)$$

where $\text{ex}(G, H)$ denotes the maximum number of edges in an H -free subgraph of G .

By considering a random $(\chi(H) - 1)$ -partition of the vertex set of $G(n, p)$, it is straightforward to show that the inequality $\text{ex}(G(n, p), H) \geq \left(1 - \frac{1}{\chi(H) - 1} + o(1)\right) \binom{n}{2} p$ holds for every $p \in [0, 1]$. On the other hand, if the number of copies of some subgraph $H' \subseteq H$ in $G(n, p)$ is much smaller than the number of edges in $G(n, p)$, then the converse inequality cannot hold, since one can make any graph H -free by removing from it one edge from each copy of H' . This observation motivates the notion of *2-density* of H , denoted by $m_2(H)$, which is defined by

$$m_2(H) = \max \left\{ \frac{e(H') - 1}{v(H') - 2} : H' \subseteq H \text{ with } v(H') \geq 3 \right\}. \quad (3)$$

It now follows easily that for every graph H with maximum degree at least 2 and every $\delta \in (0, 1/(\chi(H) - 1))$, there exists a positive constant c such that if $p_n \leq cn^{-1/m_2(H)}$, then a.a.s.

$$\text{ex}(G(n, p_n), H) > \left(1 - \frac{1}{\chi(H) - 1} + \delta\right) \binom{n}{2} p_n.$$

It was conjectured by Haxell, Kohayakawa, and Łuczak [34] and Kohayakawa, Łuczak, and Rödl [42] that the above simple argument, removing an arbitrary edge from each copy of H' in $G(n, p)$, is the main obstacle that prevents (2) from holding asymptotically almost surely. The conjecture, often referred to as Turán's theorem for random graphs, has attracted considerable attention in the past fifteen years. Numerous partial results and special cases had been established by various researchers [24, 28, 31, 34, 35, 42, 44, 61] before the conjecture was finally proved by Conlon and Gowers [14] (under the assumption that H is *strictly 2-balanced*¹) and by Schacht [58].

Theorem 1.3. *For every graph H with $\Delta(H) \geq 2$ and every positive δ , there exists a positive constant C such that if $p_n \geq Cn^{-1/m_2(H)}$, then a.a.s.*

$$\text{ex}(G(n, p_n), H) \leq \left(1 - \frac{1}{\chi(H) - 1} + \delta\right) \binom{n}{2} p_n.$$

Our methods give yet another proof of Theorem 1.3. In fact, we shall deduce from our main result, Theorem 2.2, a version of the general transference theorem of Schacht [58, Theorem 3.3], which easily implies Theorem 1.3 for such graphs H . Our version of Schacht's transference theorem, Theorem 5.2, is stated and proved in Section 5. We then, in Section 7, use it to derive a natural generalization of Theorem 1.3 to t -uniform hypergraphs, Theorem 7.2, which was also first proved in [14] and [58].

¹A graph H is 2-balanced if the maximum in (3) is achieved with $H' = H$, that is, if $m_2(H) = \frac{e(H)-1}{v(H)-2}$. It is strictly 2-balanced if $m_2(H) > m_2(H')$ for every proper subgraph $H' \subsetneq H$.

Remark 1.4. In the original version of this paper, we only proved the results concerning H -free graphs under the additional assumption that H is 2-balanced. However, a simple modification of our method (permitting multiple edges in our hypergraphs, as in [57]) allowed us to remove this condition. We would like to thank David Saxton for pointing this out.

Our methods also yield the following sparse random analogue of the famous stability theorem of Erdős and Simonovits [16, 60], originally proved by Conlon and Gowers [14] in the case when H is strictly 2-balanced and then extended to arbitrary H by Samotij [56], who adapted the argument of Schacht [58] for this purpose.

Theorem 1.5. *For every graph H with $\Delta(H) \geq 2$ and every positive δ , there exist positive constants C and ε such that if $p_n \geq Cn^{-1/m_2(H)}$, then a.a.s. the following holds. Every H -free subgraph of $G(n, p_n)$ with at least*

$$\left(1 - \frac{1}{\chi(H) - 1} - \varepsilon\right) \binom{n}{2} p_n$$

edges may be made $(\chi(H) - 1)$ -partite by removing from it at most $\delta n^2 p_n$ edges.

As with Theorem 1.3, we shall in fact deduce Theorem 1.5 from a more general statement, Theorem 6.2, which is a version of the general transference theorem for stability results proved in [56]. Theorem 6.2 is stated and proved in Section 6; in Section 7, we use it to derive Theorem 1.5.

1.3. The typical structure of H -free graphs. Let H be an arbitrary non-empty graph. For an integer n , denote by $f_n(H)$ the number of labelled H -free graphs on the vertex set $[n]$. Since every subgraph of an H -free graph is also H -free, it follows that $f_n(H) \geq 2^{\text{ex}(n, H)}$. Erdős, Frankl, and Rödl [17] proved that this crude lower bound is in a sense tight, namely that

$$f_n(H) = 2^{\text{ex}(n, H) + o(n^2)}. \tag{4}$$

Our next result can be viewed as a ‘sparse version’ of (4). Such a statement was already considered by Łuczak [46], who derived it from the so-called KLR conjecture, which we discuss in the next subsection. For integers n and m with $0 \leq m \leq \binom{n}{2}$, let $f_{n,m}(H)$ be the number of labelled H -free graphs on the vertex set $[n]$ that have exactly m edges. The following theorem refines (4) to n -vertex graphs with m edges.

Theorem 1.6. *For every graph H and every positive δ , there exists a positive constant C such that the following holds. For every $n \in \mathbb{N}$, if $m \geq Cn^{2-1/m_2(H)}$, then*

$$\binom{\text{ex}(n, H)}{m} \leq f_{n,m}(H) \leq \binom{\text{ex}(n, H) + \delta n^2}{m}.$$

In fact, we shall deduce from our main result, Theorem 2.2, a ‘counting version’ of the general transference theorem of Schacht [58, Theorem 3.3], which easily implies Theorem 1.6. This ‘counting version’ of Schacht’s theorem (which refines and, in some respects, strengthens the main results of [14, 58]) is stated and proved in Section 5. We then use it to derive Theorem 1.6 in Section 8. We remark that (4) was refined in a different sense by Balogh, Bollobás, and Simonovits [5], who showed that $f_n(H) = 2^{\text{ex}(n, H) + O(n^{2-c(H)})}$, where $c(H)$ is

some positive constant, and also gave a very precise structural description of almost all H -free graphs. We would also like to point out that our proof of Theorem 1.6 does not use Szemerédi's regularity lemma, unlike the proof given in [46] or the proofs of Erdős, Frankl, and Rödl [17] and Balogh, Bollobás, and Simonovits [5].

The result of Erdős, Frankl, and Rödl has, in some cases, a structural counterpart that significantly strengthens (4). For example, Erdős, Kleitman, and Rothschild [18] proved that *almost all* triangle-free graphs are bipartite, that is, that with probability tending to 1 as $n \rightarrow \infty$, a graph selected uniformly at random from the family of all triangle-free graphs on the vertex set $[n]$ is bipartite or, in other words (since clearly every bipartite graph is triangle-free), $f_n(K_3)$ is asymptotic to the number of bipartite graphs on the vertex set $[n]$. Extending this result, Osthus, Prömel, and Taraz [49] proved that if $m \geq Cn^{3/2}\sqrt{\log n}$ for some $C > \sqrt{3}/4$, then almost all n -vertex triangle-free graphs with m edges are bipartite. The corresponding result for K_{r+1} -free graphs was proved recently in [7].

Our next result, which is a strengthening of Theorem 1.6, is an approximate version of this statement for an arbitrary graph H . Such a statement was also considered by Łuczak [46], who derived it from the KLR conjecture. Following [46], given a positive real δ and an integer k , let us say that a graph G is (δ, k) -partite if G can be made k -partite by removing from it at most $\delta e(G)$ edges.

Theorem 1.7. *For every graph H with $\chi(H) \geq 3$, and every positive δ , there exists a positive constant C such that the following holds. If $m \geq Cn^{2-1/m_2(H)}$, then almost all H -free graphs with n vertices and m edges are $(\delta, \chi(H) - 1)$ -partite.*

As with Theorem 1.6, we shall in fact deduce Theorem 1.7 from a ‘counting version’ of the general transference theorem for stability results proved in [56]. Our version of it, Theorem 6.3, is stated and proved in Section 6. In Section 8, we use it to derive Theorem 1.7. Once again, our proof does not use the regularity lemma, unlike that in [46]. Finally, we would like to mention that, as observed by Łuczak [46], Theorem 1.7 has the following elegant corollary.

Corollary 1.8. *For every graph H with $\chi(H) \geq 3$ and every positive ε , there exist positive constants C and n_0 such that the following holds. For every $n \in \mathbb{N}$ with $n \geq n_0$ and every $m \in \mathbb{N}$ with $Cn^{2-1/m_2(H)} \leq m \leq n^2/C$,*

$$\left(\frac{\chi(H) - 2}{\chi(H) - 1} - \varepsilon \right)^m \leq \mathbb{P}(G_{n,m} \not\cong H) \leq \left(\frac{\chi(H) - 2}{\chi(H) - 1} + \varepsilon \right)^m, \quad (5)$$

where $G_{n,m}$ is a uniformly selected random n -vertex graph with m edges.

Note that (5) does not hold if m is too large; for example, if $m > n^2/4$ then $\mathbb{P}(G_{n,m} \not\cong K_3) = 0$. We remark that a great deal more is known about the structure of a typical H -free graph (drawn uniformly at random from the set of all n -vertex H -free graphs), see [6] and the references therein for more details.

1.4. The KLR conjecture. The celebrated Szemerédi regularity lemma [63], which is considered to be one of the most important and powerful tools in extremal graph theory, says

that the vertex set of every graph may be divided into a bounded number of parts of approximately the same size in such a way that most of the bipartite subgraphs induced between pairs of parts of the partition satisfy a certain pseudo-randomness condition termed ε -regularity. The strength of the regularity lemma lies in the fact that it may be combined with the so-called *embedding lemma* to show that a graph contains particular subgraphs. The combination of the regularity and embedding lemmas allows one to prove many well-known theorems in extremal graph theory, such as the theorem of Erdős and Stone [20] and the stability theorem of Erdős and Simonovits [16, 60], both mentioned in Section 1.2.

For sparse graphs, that is, n -vertex graphs with $o(n^2)$ edges, the original version of the regularity lemma is vacuous since if the vertex set of a sparse graph is partitioned into a bounded number of parts, then all induced bipartite subgraphs thus obtained are trivially ε -regular, provided that n is sufficiently large. However, it was independently observed by Kohayakawa [39] and Rödl (unpublished) that the notion of ε -regularity may be extended in a meaningful way to graphs with density tending to zero. Moreover, with this more general notion of regularity, they were also able to prove an associated regularity lemma which applies to a large class of sparse graphs, including (a.a.s.) the random graph $G(n, p)$.

Given a $p \in [0, 1]$ and a positive ε , we say that a bipartite graph between sets V_1 and V_2 is (ε, p) -regular if for every $W_1 \subseteq V_1$ and $W_2 \subseteq V_2$ with $|W_1| \geq \varepsilon|V_1|$ and $|W_2| \geq \varepsilon|V_2|$, the density $d(W_1, W_2)$ of edges between W_1 and W_2 satisfies

$$|d(W_1, W_2) - d(V_1, V_2)| \leq \varepsilon p.$$

A partition of the vertex set of a graph into r parts V_1, \dots, V_r is said to be (ε, p) -regular if $||V_i| - |V_j|| \leq 1$ for all i and j and for all but at most εr^2 pairs (V_i, V_j) , the graph induced between V_i and V_j is (ε, p) -regular. The class of graphs to which the Kohayakawa-Rödl regularity lemma applies are the so-called upper-uniform graphs. Given positive η and K , we say that an n -vertex graph G is (η, p, K) -upper-uniform if for all $W \subseteq V(G)$ with $|W| \geq \eta n$, the density of edges within W satisfies $d(W) \leq Kp$. This condition is satisfied by many natural classes of graphs, including (a.a.s.) all subgraphs of random graphs of density p . The sparse regularity lemma of Kohayakawa [39] and Rödl says the following.

The sparse Szemerédi regularity lemma. *For all positive ε , K , and r_0 , there exist a positive constant η and an integer R such that for every $p \in [0, 1]$, the following holds. Every (ε, p, K) -upper-uniform graph with at least r_0 vertices admits an (ε, p) -regular partition of its vertex set into r parts, for some $r \in \{r_0, \dots, R\}$.*

We remark that a version of this theorem avoiding the need for the upper-uniformity assumption was recently proved by Scott [59].

The aforementioned embedding lemma roughly says that if we start with an arbitrary graph H , replace its vertices by large independent sets and its edges by ε -regular bipartite graphs with density much larger than ε , then this blown-up graph will contain a copy of H . To make it more precise, let H be a graph on the vertex set $\{1, \dots, v(H)\}$, let ε and p be as above, and let n and m be integers satisfying $0 \leq m \leq n^2$. Let us denote by $\mathcal{G}(H, n, m, p, \varepsilon)$ the collection of all graphs G constructed in the following way. The vertex set of G is a disjoint union $V_1 \cup \dots \cup V_{v(H)}$ of sets of size n , one for each vertex of H . For each edge $\{i, j\}$ of H , we add to G an (ε, p) -regular bipartite graph with m edges between the sets V_i and

V_j . These are the only edges of G . With this notation in hand, we can state the embedding lemma. Given any graph G as above, we define *canonical copies of H* to be all copies of H in G in which (the image of) each vertex $i \in V(H)$ lies in the set $V_i \subseteq V(G)$.

The embedding lemma. *For every graph H and every positive d , there exist a positive ε and an integer n_0 such that for every n and m with $n \geq n_0$ and $m \geq dn^2$, every $G \in \mathcal{G}(H, n, m, 1, \varepsilon)$ contains a canonical copy of H .*

One might hope that a similar statement holds when one replaces 1 by an arbitrary p and the assumption $m \geq dn^2$ by $m \geq pdn^2$, even if p is a decreasing function of n . However, for an arbitrary function p , this is too much to hope for. Indeed, consider the random ‘blow-up’ of H , that is, the random graph G obtained from H by replacing each vertex of H by an independent set of size n and each edge of H by a random bipartite graph with pn^2 edges. With high probability, the number of canonical copies of H in G will be about $p^{\varepsilon(H)} n^{v(H)}$ and hence if $p^{\varepsilon(H)} n^{v(H)} \ll pn^2$, then one can remove all copies of H from G by deleting a tiny proportion of all edges. Since in the above argument one may replace H with an arbitrary subgraph $H' \subseteq H$, it follows easily² that if $p \ll n^{-1/m_2(H)}$, then there are graphs in $\mathcal{G}(H, n, pn^2, p, \varepsilon)$ that do not contain any canonical copies of H .

As in the case of Turán’s theorem for random graphs, see Section 1.2, one might still hope that if $p \geq Cn^{-1/m_2(H)}$ for some large constant C , then the natural sparse analogue of the embedding lemma discussed above holds. However, it was observed by Łuczak (see [32, 43]) that, somewhat surprisingly, for any graph H which contains a cycle and any function p satisfying $p = o(1)$, there are graphs in $\mathcal{G}(H, n, pn^2, p, \varepsilon)$ with no canonical copy of H . Nevertheless, it still seemed likely that such atypical graphs comprise so tiny a proportion of $\mathcal{G}(H, n, m, p, \varepsilon)$ that they do not appear in $G(n, p)$ asymptotically almost surely.

This was formalized in the following conjecture of Kohayakawa, Łuczak, and Rödl [42], usually referred to as the *KLR conjecture*. Given a graph H , integers m and n , a $p \in [0, 1]$, and a positive ε , let $\mathcal{G}^*(H, n, m, p, \varepsilon)$ denote the collection of graphs in $\mathcal{G}(H, n, m, p, \varepsilon)$ that contain no canonical copy of H . We will prove the conjecture in Section 9.

Theorem 1.9 (The KLR conjecture). *For every graph H and every positive β , there exist positive constants C , n_0 , and ε such that the following holds. For every $n \in \mathbb{N}$ with $n \geq n_0$ and $m \in \mathbb{N}$ with $m \geq Cn^{2-1/m_2(H)}$,*

$$|\mathcal{G}^*(H, n, m, m/n^2, \varepsilon)| \leq \beta^m \binom{n^2}{m}^{\varepsilon(H)}.$$

The KLR conjecture has been one of the central open questions in extremal graph theory and has attracted substantial attention from many researchers over the past fifteen years. It has been verified in several special cases. It is easy to see that it holds for all graphs H which do not contain a cycle. The cases $H = K_3$, K_4 , and K_5 were resolved in [41], [30], and [31], respectively. The case $H = C_\ell$ has also been resolved, but here the history is somewhat more complex. A proof under some extra technical assumptions was given in [40]. Those extra assumptions were later removed in [29] and, independently, in [12]. We remark here that in parallel to this work, Conlon, Gowers, Samotij, and Schacht [15] have proved a sparse

²Note that we also replace p with some $p' = (1 + o(1))p$, and that the removal of $o(pn^2)$ edges does not affect the ε -regularity conditions.

analogue of the *counting lemma* for subgraphs of the random graph $G(n, p)$, which may be viewed as a version of the KLR conjecture that is stronger in some aspects and weaker in other aspects.

It is well-known that Theorem 1.9 easily implies Turán's theorem for random graphs, Theorem 1.3, and also its stability version, Theorem 1.5. In fact, this was the original motivation behind the KLR conjecture, see [42]. Moreover, it was proved by Łuczak [46] that Theorem 1.9 implies Theorems 1.6 and 1.7. The work of Conlon and Gowers [14] and Schacht [58] (see also [56]), as well as this work, have shown that one does not need to appeal to the sparse regularity lemma and to the KLR conjecture in order to prove such extremal statements in random graphs. Nevertheless, there are still many beautiful corollaries of the conjecture that cannot (yet) be proved by other means. For discussion and derivation of some of them, we refer the reader to [15]. Here, we present only one corollary of the KLR conjecture, the threshold for asymmetric Ramsey properties of random graphs, which does not follow from the version of the conjecture proved in [15]. The deduction of this result from the KLR conjecture is essentially due to Kohayakawa and Kreuter [40].

1.5. Ramsey properties of random graphs. Let H be a fixed graph and let r be a positive integer. For an arbitrary graph G , we write $G \rightarrow (H)_r$ if every r -coloring of the edges of G contains a monochromatic copy of H . It follows from the classical result of Ramsey [51] that $K_n \rightarrow (H)_r$, provided that n is sufficiently large. Ramsey properties of random graphs were first investigated by Frankl and Rödl [22] and since then much effort has been devoted to their study. Most notably, Rödl and Ruciński [52, 53] established the following general threshold result.

Theorem 1.10. *For every graph H that is not a forest, and every positive integer r , there exist positive constants c and C such that*

$$\lim_{n \rightarrow \infty} \mathbb{P}(G(n, p_n) \rightarrow (H)_r) = \begin{cases} 1 & \text{if } p_n \geq Cn^{-1/m_2(H)}, \\ 0 & \text{if } p_n \leq cn^{-1/m_2(H)}. \end{cases}$$

In the above discussion, a copy of the same graph H is forbidden in each of the r color classes. A natural generalization of Theorem 1.10 would determine thresholds for so-called *asymmetric Ramsey properties*. For any graphs G, H_1, \dots, H_r , we write $G \rightarrow (H_1, \dots, H_r)$ if for every coloring of the edges of G with colors $1, \dots, r$, there exists, for some $i \in [r]$, a copy of H_i all of whose edges have color i . In the context of asymmetric Ramsey properties of random graphs, the following generalization of the 2-density $m_2(\cdot)$ was introduced in [40]. For two graphs H_1 and H_2 , define³

$$m_2(H_1, H_2) = \max \left\{ \frac{e(H'_1)}{v(H'_1) - 2 + 1/m_2(H_2)} : H'_1 \subseteq H_1 \text{ with } v(H'_1) \geq 3 \right\}. \quad (6)$$

Kohayakawa and Kreuter [40] formulated the following conjecture and proved it in the case when all H_i are cycles.

³To motivate this definition, set $p = n^{-1/m_2(H_1, H_2)}$ and observe that the edges of $G(n, p)$ which are contained in a copy of each subgraph $H'_1 \subseteq H_1$ have density roughly $n^{-1/m_2(H_2)}$.

Conjecture 1.11. *Let H_1, \dots, H_r be graphs with $1 < m_2(H_r) \leq \dots \leq m_2(H_1)$. Then there exist constants c and C such that*

$$\lim_{n \rightarrow \infty} \mathbb{P}(G(n, p_n) \rightarrow (H_1, \dots, H_r)) = \begin{cases} 1 & \text{if } p_n \geq Cn^{-1/m_2(H_1, H_2)}, \\ 0 & \text{if } p_n \leq cn^{-1/m_2(H_1, H_2)}. \end{cases}$$

More accurately, the above conjecture was stated in [40] only in the case $r = 2$, but the above generalization is quite natural.⁴ There had been little progress on Conjecture 1.11 until quite recently, when the 0-statement was proved by Marciniszyn, Skokan, Spöhel, and Steger [47] in the case where all of the H_i are cliques, and the 1-statement in the case $r = 2$ was established⁵ by Kohayakawa, Schacht, and Spöhel [45] under very mild extra assumptions on H_1 and H_2 . It was observed in [47, Theorem 31] that, using Theorem 1.9, the approach of Kohayakawa and Kreuter [40], which employs the sparse regularity lemma, can be adapted to yield a proof of the 1-statement in Conjecture 1.11 for the following class of graphs.

Theorem 1.12. *Let H_1, \dots, H_r be graphs with $1 < m_2(H_r) \leq \dots \leq m_2(H_1)$ and such that H_1 is strictly 2-balanced. Then there exists a constant C such that if $p_n \geq Cn^{-1/m_2(H_1, H_2)}$, then a.a.s.*

$$G(n, p_n) \rightarrow (H_1, \dots, H_r).$$

For the deduction of Theorem 1.12 from Theorem 1.9, see [40] and [47, Section 4].

1.6. Outline of the paper. The remainder of this paper is organized as follows. In Section 2, we state and discuss our main result, Theorem 2.2, which we then prove in Section 3. In Section 4, we discuss the applications of Theorem 2.2 in the context of subsets of $[n]$ with no k -term arithmetic progressions. In particular, we prove Theorem 1.1 and use it to derive Corollary 1.2. In Section 5, we prove two versions of the general transference theorem of Schacht [58, Theorem 3.3] (obtained independently, in a slightly different form, by Conlon and Gowers [14]) – a ‘random’ version suited for extremal problems in sparse random discrete structures and its ‘counting’ counterpart that generalizes Theorem 1.1. In Section 6, we prove ‘random’ and ‘counting’ versions of the general stability result of Conlon and Gowers [14] in a form that is easily comparable with [56, Theorem 3.4]. In Section 7, we discuss several applications of Theorem 2.2 in the context of the Turán problem in sparse random graphs. In particular, using the results of Sections 5 and 6 we give new proofs of the sparse random analogues (stated above) of the classical theorems of Erdős and Stone, and Erdős and Simonovits, see Section 1.2. In Section 8, we discuss applications of Theorem 2.2 to the problem of describing the typical structure of a sparse graph without a forbidden subgraph. In particular, we prove sparse analogues of classical theorems of Erdős, Frankl, and Rödl and Erdős, Kleitman, and Rothschild, see Section 1.3. Finally, in Section 9, we use Theorem 2.2 to prove the KLR conjecture for every graph H .

⁴To see why the graphs H_3, \dots, H_r do not appear in the threshold, replace each of H_2, \dots, H_r by the disjoint union $H' = H_2 \cup \dots \cup H_r$, and note that $m_2(H') = m_2(H_2)$, see [47].

⁵In their concluding remarks, the authors of [45] moreover claim that their method can be extended to the setting with more than two colours, using ideas from [53].

2. THE MAIN THEOREM

In this section, we present the main result of this paper, Theorem 2.2, which gives a structural characterization of the collection of all independent sets in a large class of uniform hypergraphs. Let us stress here that all of the hypergraphs we consider are allowed to have multiple edges; moreover, we shall always count edges with multiplicities.

We start with an important definition. Recall that a family of sets $\mathcal{F} \subseteq \mathcal{P}(V)$ is called *increasing* (or an *upset*) if it is closed under taking supersets, that is, if for every $A, B \subseteq V$, $A \in \mathcal{F}$ and $A \subseteq B$ imply that $B \in \mathcal{F}$.

Definition 2.1. Let \mathcal{H} be a uniform hypergraph with vertex set V , let \mathcal{F} be an increasing family of subsets of V and let $\varepsilon \in (0, 1]$. We say that \mathcal{H} is $(\mathcal{F}, \varepsilon)$ -dense if

$$e(\mathcal{H}[A]) \geq \varepsilon e(\mathcal{H})$$

for every $A \in \mathcal{F}$.

A moment of thought reveals that for an arbitrary hypergraph \mathcal{H} and $\varepsilon \in (0, 1]$, it is extremely simple to find families $\mathcal{F} \subseteq \mathcal{P}(V(\mathcal{H}))$ for which \mathcal{H} is $(\mathcal{F}, \varepsilon)$ -dense. To this end, let

$$\mathcal{F}_\varepsilon = \{A \subseteq V(\mathcal{H}) : e(\mathcal{H}[A]) \geq \varepsilon e(\mathcal{H})\}$$

and note that \mathcal{F}_ε is increasing and \mathcal{H} is $(\mathcal{F}_\varepsilon, \varepsilon)$ -dense. In fact, the families \mathcal{F} for which \mathcal{H} is $(\mathcal{F}, \varepsilon)$ -dense are precisely all increasing subfamilies of \mathcal{F}_ε .

In this work, we will be interested in upsets that admit a much more ‘constructive’ description than that of \mathcal{F}_ε . Many such families arise naturally in the study of extremal and structural problems in combinatorics. For example, consider the k -uniform hypergraph \mathcal{H}_1 on the vertex set $[n]$ whose edges are all k -term arithmetic progressions in $[n]$ and let \mathcal{F}_1 be the collection of all subsets of $[n]$ with at least δn elements. Clearly, \mathcal{F}_1 is an upset and it follows from the famous theorem of Szemerédi [62] that \mathcal{H}_1 is $(\mathcal{F}_1, \varepsilon)$ -dense for some positive ε depending only on δ and k , see Section 4. Similarly, consider the 3-uniform hypergraph \mathcal{H}_2 on the vertex set $E(K_n)$ whose edges are edge sets of all copies of K_3 in the complete graph K_n and let \mathcal{F}_2 be the family of all n -vertex graphs (subgraphs of K_n) with at least $(1/2 - \varepsilon)\binom{n}{2}$ edges such that every 2-coloring of its vertices yields at least δn^2 monochromatic edges. Again, \mathcal{F}_2 is increasing and it follows from the stability theorem of Erdős and Simonovits [16, 60] and the triangle removal lemma of Ruzsa and Szemerédi [55] that \mathcal{H}_2 is $(\mathcal{F}_2, \varepsilon)$ -dense, provided that ε is sufficiently small as a function of δ .

Our main result roughly says the following. If \mathcal{H} is a uniform hypergraph that is $(\mathcal{F}, \varepsilon)$ -dense for some family \mathcal{F} and whose edge distribution satisfies certain natural boundedness conditions, then the collection $\mathcal{I}(\mathcal{H})$ of all independent sets in \mathcal{H} admits a partition into relatively few classes such that all independent sets in one class are essentially contained in a single set $A \notin \mathcal{F}$. Before we state the result, we first need to quantify the above boundedness conditions for the edge distribution of a hypergraph. Given a hypergraph \mathcal{H} , for each $T \subseteq V(\mathcal{H})$, we define⁶

$$\deg_{\mathcal{H}}(T) = |\{e \in \mathcal{H} : T \subseteq e\}|,$$

⁶We emphasize that if \mathcal{H} has multiple edges, then $\{e \in \mathcal{H} : T \subseteq e\}$ should be thought of as a multi-set. In other words, $\deg_{\mathcal{H}}(T)$ is the number of edges of \mathcal{H} , counted with multiplicities, which contain T .

and let

$$\Delta_\ell(\mathcal{H}) = \max \{ \deg_{\mathcal{H}}(T) : T \subseteq V(\mathcal{H}) \text{ and } |T| = \ell \}.$$

Recall that $\mathcal{I}(\mathcal{H})$ denotes the family of all independent sets in \mathcal{H} . The following theorem is our main result.

Theorem 2.2. *For every $k \in \mathbb{N}$ and all positive c and ε , there exists a positive constant C such that the following holds. Let \mathcal{H} be a k -uniform hypergraph and let $\mathcal{F} \subseteq \mathcal{P}(V(\mathcal{H}))$ be an increasing family of sets such that $|A| \geq \varepsilon v(\mathcal{H})$ for all $A \in \mathcal{F}$. Suppose that \mathcal{H} is $(\mathcal{F}, \varepsilon)$ -dense and $p \in (0, 1)$ is such that, for every $\ell \in [k]$,*

$$\Delta_\ell(\mathcal{H}) \leq c \cdot p^{\ell-1} \frac{e(\mathcal{H})}{v(\mathcal{H})}.$$

Then there exists a family $\mathcal{S} \subseteq \binom{V(\mathcal{H})}{\leq Cp \cdot v(\mathcal{H})}$ and functions $f: \mathcal{S} \rightarrow \overline{\mathcal{F}}$ and $g: \mathcal{I}(\mathcal{H}) \rightarrow \mathcal{S}$ such that for every $I \in \mathcal{I}(\mathcal{H})$,

$$g(I) \subseteq I \quad \text{and} \quad I \setminus g(I) \subseteq f(g(I)).$$

Roughly speaking, if \mathcal{H} satisfies certain technical conditions, then each independent set I in \mathcal{H} can be labelled with a small subset $g(I)$ in such a way that all sets labelled with some $S \in \mathcal{S}$ are essentially contained in a single set $f(S)$ that contains very few edges of \mathcal{H} . We remark that the constant C in the theorem has only a polynomial dependence on ε . Unfortunately, however, in most of our applications ε will have a tower-type dependence on some other parameter.

Theorem 2.2 will be proved in Section 3. We end this section with a short informal discussion of its consequences. As we have already mentioned, Theorem 2.2 combined with some classical extremal results on discrete structures has strikingly strong implications. Let us briefly explain why this is so. Many classical extremal problems ask for an estimate on the number of independent sets (of a certain size) in some auxiliary uniform hypergraph. If applicable, Theorem 2.2 implies that all such independent sets are almost contained in one of very few sets that are almost independent, that is, contain a small number of copies of some forbidden substructure. If we know a good characterization of sets that are almost independent in the above sense, which is often the case, we can easily obtain an upper bound on the number of independent sets. For example, consider the problem of counting subsets of $[n]$ with no k -term AP and recall the definition of \mathcal{H}_1 and \mathcal{F}_1 from the beginning of this section. Theorem 2.2, applied to this pair, implies that every subset of $[n]$ with no k -term AP is essentially contained in one of at most $\binom{n}{O(n^{1-1/(k-1)})}$ sets of size at most δn each, where δ is an arbitrarily small positive constant. This easily implies that if $m \gg n^{1-1/(k-1)}$, then there are at most $\binom{2\delta n}{m}$ sets of size m with no k -term AP. For more details, we refer the reader to Section 4.

3. PROOF OF THE MAIN THEOREM

In this section, we shall prove Theorem 2.2. The main ingredient in the proof is the following proposition, which (roughly) says that Theorem 2.2 holds in the special case when

\mathcal{F} is the family of all subsets of $V(\mathcal{H})$ with at least $(1 - \delta)v(\mathcal{H})$ elements. Theorem 2.2 follows by applying Proposition 3.1 a constant number of times.

Proposition 3.1. *For every integer k and positive c , there exists a positive δ such that the following holds. Let $p \in (0, 1)$ and suppose that \mathcal{H} is a k -uniform hypergraph such that, for every $\ell \in [k]$,*

$$\Delta_\ell(\mathcal{H}) \leq c \cdot p^{\ell-1} \frac{e(\mathcal{H})}{v(\mathcal{H})}.$$

Then there exist a family $\mathcal{S} \subseteq \binom{V(\mathcal{H})}{\leq (k-1)p \cdot v(\mathcal{H})}$ and functions $f_0: \mathcal{S} \rightarrow \mathcal{P}(V(\mathcal{H}))$ and $g_0: \mathcal{I}(\mathcal{H}) \rightarrow \mathcal{S}$ such that for every $I \in \mathcal{I}(\mathcal{H})$,

$$g_0(I) \subseteq I \subseteq f_0(g_0(I)) \cup g_0(I) \quad \text{and} \quad |f_0(g_0(I))| \leq (1 - \delta)v(\mathcal{H}).$$

Moreover, if for some $I, I' \in \mathcal{I}(\mathcal{H})$, $g_0(I) \subseteq I'$ and $g_0(I') \subseteq I$, then $g_0(I) = g_0(I')$.

The final line of Proposition 3.1 states that the labelling function g_0 exhibits a certain consistency. This property of g_0 , which may look somewhat puzzling, will be crucial in the proof of Theorem 2.2.

In order to prove Proposition 3.1, given an independent set $I \in \mathcal{I}(\mathcal{H})$, we shall construct a sequence (B_{k-1}, \dots, B_q) of subsets of I with $|B_{k-1}|, \dots, |B_q| \leq pv(\mathcal{H})$, for some $q \in [k-1]$, and use it to define a sequence $(\mathcal{H}_{k-1}, \dots, \mathcal{H}_r)$, where $r \in \{q, q+1\}$, of hypergraphs such that the following holds for each $i \in \{r, \dots, k-1\}$:

- (a) \mathcal{H}_i is an i -uniform hypergraph on the vertex set $V(\mathcal{H})$,
- (b) I is an independent set in \mathcal{H}_i ,
- (c) $\Delta_1(\mathcal{H}_i) \leq O(e(\mathcal{H}_i)/v(\mathcal{H}_i))$, and
- (d) $e(\mathcal{H}_i) \geq \Omega(p^{k-i}e(\mathcal{H}))$.

We shall be able to do it in such a way that in the end, there will be a set $A \subseteq V(\mathcal{H})$ of size at most $(1 - \delta)v(\mathcal{H})$ such that the remaining elements of I (i.e., the set $I \setminus S$, where $S = B_{k-1} \cup \dots \cup B_q$) must all lie inside A . If $r = 1$, then we will simply let A be the set of non-edges of the 1-uniform hypergraph \mathcal{H}_1 ; in this case, the upper bound on $|A|$ will follow from (c) and (d). If $r > 1$, then we will obtain an appropriate A while trying (and failing) to construct the hypergraph \mathcal{H}_{r-1} using the hypergraph \mathcal{H}_r and the set B_{r-1} . Crucially, this set A will depend solely on S , that is, if for some pair $I, I' \in \mathcal{I}(\mathcal{H})$ our procedure generates (S, A) and (S', A') , respectively, and if $S = S'$, then also $A = A'$. This will allow us to set $g_0(I) = S$ and $f_0(S) = A$.

3.1. The Algorithm Method. For the remainder of this section, let us fix k , c , p , and \mathcal{H} as in the statement of Proposition 3.1. Without loss of generality, we may assume that $c \geq 1$. Let I be an independent set in \mathcal{H} . We shall describe a procedure of choosing the sets $B_i \subseteq I$ and constructing the hypergraphs \mathcal{H}_i as above. This procedure, which we shall term the *Scythe Algorithm*, lies at the heart of the proof of Proposition 3.1.

The general strategy used in the Scythe Algorithm, that of selecting a small set S of high-degree vertices and using it to define a set A such that $S \subseteq I \subseteq A \cup S$, dates back to the work of Kleitman and Winston [38], who used it to bound the number of independent sets in graphs satisfying the following local density condition: all sufficiently large vertex

sets induce subgraphs with many edges. Recently, Balogh and Samotij [10, 11] refined the ideas of Kleitman and Winston and obtained a bound on the number of independent sets in uniform hypergraphs satisfying a similar local density condition. Even more recently, Alon, Balogh, Morris and Samotij [1] used similar ideas to bound the number of independent sets in ‘almost linear’ 3-uniform hypergraphs satisfying a more general density condition termed (α, \mathcal{B}) -stability, see Definition 6.1. Here, we combine, generalize, and refine all of the above approaches and make them work in the general setting of $(\mathcal{F}, \varepsilon)$ -dense uniform hypergraphs.

At each step of the Scythe Algorithm, we shall order the vertices of a certain subhypergraph of \mathcal{H} with respect to their degrees in that subhypergraph. For the sake of brevity and clarity of the presentation, let us make the following definition.

Definition 3.2 (Max-degree order). Given a hypergraph \mathcal{G} , we define the *max-degree order* on $V(\mathcal{G})$ as follows:

- (1) Fix an arbitrary total ordering of $V(\mathcal{G})$.
- (2) For each $j \in \{1, \dots, v(\mathcal{G})\}$, let u_j be the maximum-degree vertex in the hypergraph $\mathcal{G}[V(\mathcal{G}) \setminus \{u_1, \dots, u_{j-1}\}]$; ties are broken by giving preference to vertices which come earlier in the order chosen in (1).
- (3) The max-degree order on $V(\mathcal{G})$ is $(u_1, \dots, u_{v(\mathcal{G})})$.

Finally, we write $W(u)$ to denote the initial segment of the max-degree order on $V(\mathcal{G})$ that ends with u , i.e., for every j , we let $W(u_j) = \{u_1, \dots, u_j\}$.

We remark here that the only property of the max-degree order that will be important for us is that for every $j \in \{1, \dots, v(\mathcal{G})\}$, the degree of the vertex u_j in the hypergraph $\mathcal{G}[V(\mathcal{G}) \setminus W(u_{j-1})]$ is at least as large as the average degree of this hypergraph.

We next define the numbers Δ_ℓ^i , where $1 \leq \ell \leq i \leq k$, which will play a crucial role in the description and the analysis of the algorithm.

Definition 3.3. For every $\ell \in [k]$, let $\Delta_\ell^k = \Delta_\ell(\mathcal{H})$ and for all $i \in [k-1]$ and $\ell \in [i]$, let

$$\Delta_\ell^i = \max \left\{ 2 \cdot \Delta_{\ell+1}^{i+1}, p \cdot \Delta_\ell^{i+1} \right\}. \quad (7)$$

We use the numbers Δ_ℓ^i to define the following families of sets with high degree.

Definition 3.4. Given an $i \in [k]$, an i -uniform hypergraph \mathcal{G} and an $\ell \in [i]$, let

$$M_\ell^i(\mathcal{G}) = \left\{ T \in \binom{V(\mathcal{G})}{\ell} : \deg_{\mathcal{G}}(T) \geq \frac{\Delta_\ell^i}{2} \right\}.$$

Let $b = pv(\mathcal{H})$ and for each $i \in [k]$, let $c_i = (ck2^{k+1})^{i-k}$.

Properties. The key properties that we would like the constructed hypergraph \mathcal{H}_i to possess are:

- (P1) \mathcal{H}_i is i -uniform and $V(\mathcal{H}_i) = V(\mathcal{H})$,
- (P2) I is an independent set in \mathcal{H}_i ,
- (P3) $\Delta_\ell(\mathcal{H}_i) \leq \Delta_\ell^i$ for each $\ell \in [i]$,
- (P4) $e(\mathcal{H}_i) \geq c_i p^{k-i} e(\mathcal{H})$.

Set $\mathcal{H}_k = \mathcal{H}$ and note that (P1)–(P4) are vacuously satisfied for $i = k$. The main step of the Scythe Algorithm will be a procedure that, given \mathcal{H}_{i+1} and I satisfying (P1)–(P4), outputs a set $B_i \subseteq I$ of cardinality at most b , a set $A_i \subseteq V(\mathcal{H})$ with the property that $I \setminus B_i \subseteq A_i$, and a hypergraph \mathcal{H}_i satisfying (P1)–(P3). Moreover, if the constructed \mathcal{H}_i does not satisfy (P4), then we have $|A_i| \leq (1 - c_i)v(\mathcal{H})$. Crucially, these A_i and \mathcal{H}_i depend solely on B_i and \mathcal{H}_{i+1} , that is, if on two inputs (\mathcal{H}_{i+1}, I) and (\mathcal{H}_{i+1}, I') , the procedure outputs the same set B_i , it also outputs the same A_i and \mathcal{H}_i .

The Scythe Algorithm. Given an $(i + 1)$ -uniform hypergraph \mathcal{H}_{i+1} and an independent set $I \in \mathcal{I}(\mathcal{H}_{i+1})$, set $\mathcal{A}_{i+1}^{(0)} = \mathcal{H}_{i+1}$ and let $\mathcal{H}_i^{(0)}$ be the empty hypergraph on the vertex set $V(\mathcal{H})$. For $j = 0, \dots, b - 1$, do the following:

- (1) If $I \cap V(\mathcal{A}_{i+1}^{(j)}) = \emptyset$, then set $\mathcal{H}_i = \mathcal{H}_i^{(0)}$, $A_i = \emptyset$, and $B_i = \{u_0, \dots, u_{j-1}\}$ and STOP.
- (2) Let u_j be the first vertex of I in the max-degree order on $V(\mathcal{A}_{i+1}^{(j)})$.
- (3) Let $\mathcal{H}_i^{(j+1)}$ be the hypergraph on the vertex set $V(\mathcal{H})$ defined by:

$$\mathcal{H}_i^{(j+1)} = \mathcal{H}_i^{(j)} \cup \left\{ D \in \binom{V(\mathcal{H})}{i} : D \cup \{u_j\} \in \mathcal{A}_{i+1}^{(j)} \right\}.$$

- (4) Let $\mathcal{A}_{i+1}^{(j+1)}$ be the hypergraph on the vertex set $V(\mathcal{A}_{i+1}^{(j)}) \setminus W(u_j)$ defined by:⁷

$$\mathcal{A}_{i+1}^{(j+1)} = \left\{ D \in \mathcal{A}_{i+1}^{(j)} : D \cap W(u_j) = \emptyset \text{ and } T \not\subseteq D \text{ for every } T \in \bigcup_{\ell=1}^i M_\ell^i(\mathcal{H}_i^{(j+1)}) \right\}.$$

Finally, set $\mathcal{H}_i = \mathcal{H}_i^{(b)}$, $A_i = V(\mathcal{A}_{i+1}^{(b)})$, and $B_i = \{u_0, \dots, u_{b-1}\}$.

We shall now establish various properties of the Scythe Algorithm. We begin by making some basic (but key) observations.

Lemma 3.5. *The following hold for every $i \in [k - 1]$:*

- (a) \mathcal{H}_i is i -uniform and $V(\mathcal{H}_i) = V(\mathcal{H})$.
- (b) If $I \in \mathcal{I}(\mathcal{H}_{i+1})$, then $I \in \mathcal{I}(\mathcal{H}_i)$.
- (c) $B_i \subseteq I \subseteq A_i \cup B_i$.
- (d) The hypergraph \mathcal{H}_i and the set A_i depend only on \mathcal{H}_{i+1} and the set B_i .

Proof. Property (a) is trivial. To see (b), simply observe that each edge of \mathcal{H}_i is of the form $D \setminus \{u\}$ for some $D \in \mathcal{H}_{i+1}$ and $u \in I$. Thus, if I contains an edge of \mathcal{H}_i , it must also contain an edge of \mathcal{H}_{i+1} . To see (c), observe that for each j , u_j is the first vertex of I in the max-degree order on $V(\mathcal{A}_{i+1}^{(j)})$ and hence $W(u_j) \cap I = \{u_j\}$. It follows that $B_i \subseteq I$ and that $I \setminus A_i = B_i$. Note in particular that if $A_i = \emptyset$, then $I \cap V(\mathcal{A}_{i+1}^{(j)}) = \emptyset$ for some $j \in \{0, \dots, b\}$, which implies that $B_i = I$. Finally, to prove (d), observe that all steps of the Scythe Algorithm are deterministic and that every element of I that we need to observe in order to define A_i and \mathcal{H}_i is placed in B_i . More precisely, note that while choosing the vertex u_j , we only need to know the first vertex of I in the max-degree order on $V(\mathcal{A}_{i+1}^{(j)})$; the remaining vertices remain unobserved. Since we have $W(u_j) \cap B_i = W(u_j) \cap I = \{u_j\}$,

⁷We emphasize that $W(u_j)$ is defined relative to the max-degree order on $V(\mathcal{A}_{i+1}^{(j)})$.

this information can be recovered from B_i . Thus, at each step, the hypergraph $\mathcal{H}_i^{(j+1)}$ can be recovered from $\mathcal{H}_i^{(j)}$ and B_i , and the hypergraph $\mathcal{A}_{i+1}^{(j+1)}$ can be recovered from $\mathcal{A}_{i+1}^{(j)}$, $\mathcal{H}_i^{(j+1)}$ and B_i . Hence, a trivial inductive argument proves that, if the algorithm does not stop in step (1), for each $j \in \{0, \dots, b\}$, the hypergraphs $\mathcal{H}_i^{(j)}$ and $\mathcal{A}_{i+1}^{(j)}$ are determined by \mathcal{H}_{i+1} and the set B_i , as required. Finally, the algorithm stops in step (1) if and only if $|B_i| < b$. If this happens, then \mathcal{H}_i and A_i are empty. \square

We next show that the Scythe Algorithm exhibits a certain ‘consistency’ while generating its output. This property will be important in the proof of Proposition 3.1.

Lemma 3.6. *Suppose that on inputs (\mathcal{H}_{i+1}, I) and (\mathcal{H}_{i+1}, I') , the Scythe Algorithm outputs $(A_i, B_i, \mathcal{H}_i)$ and $(A'_i, B'_i, \mathcal{H}'_i)$, respectively. If $B_i \subseteq I'$ and $B'_i \subseteq I$, then $(A_i, B_i, \mathcal{H}_i) = (A'_i, B'_i, \mathcal{H}'_i)$.*

Proof. By Lemma 3.5, it suffices to show that $B_i = B'_i$. Let us first consider the (degenerate) case when $\min\{|B_i|, |B'_i|\} < b$. Without loss of generality, we may assume that $|B_i| < b$. This means that, while running on (\mathcal{H}_{i+1}, I) , the Scythe Algorithm stopped in step (1). By Lemma 3.5, it follows that $B_i = I$ and hence $B'_i \subseteq B_i$, which means that $|B'_i| < b$ and therefore $B'_i = I'$. Hence, $B_i = B'_i$, as claimed. On the other hand, if $|B_i| = |B'_i| = b$ and $B_i \neq B'_i$, then there must exist some j such that $u_j \neq u'_j$. Let j be the smallest such index. Note that by the minimality of j , we have $\mathcal{A}_{i+1}^{(j)} = (\mathcal{A}_{i+1}^{(j)})' = \mathcal{A}$. Since $u_j \neq u'_j$, one of these vertices comes earlier in the max-degree order on $V(\mathcal{A})$; without loss of generality, we may suppose that it is u_j . Since $B_i \subseteq I'$, it follows that $u_j \in I'$ and hence the Algorithm, while running on the input (\mathcal{H}_{i+1}, I') , would not pick u'_j in step j , a contradiction. This shows that in fact $B_i = B'_i$, as required. \square

The next lemma shows that if \mathcal{H}_{i+1} satisfies (P3), then so does \mathcal{H}_i . The lemma follows easily from the definitions of Δ_ℓ^i and $M_\ell^i(\mathcal{G})$.

Lemma 3.7. *If $\Delta_{\ell+1}(\mathcal{H}_{i+1}) \leq \Delta_{\ell+1}^{i+1}$ for some $\ell \in [i]$, then $\Delta_\ell(\mathcal{H}_i) \leq \Delta_\ell^i$.*

Proof. The crucial observation is that if

$$\deg_{\mathcal{H}_i^{(j)}}(T) \geq \frac{\Delta_\ell^i}{2}$$

for some $T \in \binom{V(\mathcal{H})}{\ell}$ and $j \in [b]$, then all edges containing T are removed from $\mathcal{A}_{i+1}^{(j)}$ and hence no more such edges are added to \mathcal{H}_i . It follows that $\deg_{\mathcal{H}_i}(T) = \deg_{\mathcal{H}_i^{(j)}}(T)$. Moreover, when we extend $\mathcal{H}_i^{(j-1)}$ to $\mathcal{H}_i^{(j)}$, then we only add to it sets D such that $D \cup \{u_j\} \in \mathcal{A}_{i+1}^{(j-1)} \subseteq \mathcal{H}_{i+1}$ and hence

$$\deg_{\mathcal{H}_i^{(j)}}(T) - \deg_{\mathcal{H}_i^{(j-1)}}(T) \leq \deg_{\mathcal{H}_{i+1}}(T \cup \{u_j\}) \leq \Delta_{|T|+1}(\mathcal{H}_{i+1}).$$

It follows that

$$\Delta_\ell(\mathcal{H}_i) \leq \frac{\Delta_\ell^i}{2} + \Delta_{\ell+1}(\mathcal{H}_{i+1}) \leq \frac{\Delta_\ell^i}{2} + \Delta_{\ell+1}^{i+1} \leq \Delta_\ell^i$$

where the last inequality follows from (7). \square

Next, let us establish an easy bound on the numbers Δ_1^i .

Lemma 3.8. $\Delta_1^i \leq c 2^k p^{k-i} \frac{e(\mathcal{H})}{v(\mathcal{H})}$ for every $i \in \{1, \dots, k\}$.

Proof. To prove the lemma, simply note that, by the definition of Δ_ℓ^i , for every $i \in [k]$ and every $\ell \in [i]$,

$$\Delta_\ell^i = 2^d p^{k-i-d} \Delta_{d+\ell}(\mathcal{H}) \quad \text{for some } d \in \{0, \dots, k-i\}. \quad (8)$$

One easily proves (8) by induction on $k-i$. Intuitively, d in (8) is the number of times that the first term in the maximum in (7) is larger than the second term when following the recursive definition of Δ_ℓ^i back to $\Delta_{d+\ell}^k$.

Since $\Delta_\ell(\mathcal{H}) \leq c \cdot p^{\ell-1} \frac{e(\mathcal{H})}{v(\mathcal{H})}$, as in the statement of Proposition 3.1, it follows from (8) that

$$\Delta_1^i \leq \max_{0 \leq d \leq k-i} \{2^d p^{k-i-d} \Delta_{d+1}(\mathcal{H})\} \leq \max_{0 \leq d \leq k-i} \left\{ 2^d p^{k-i-d} \cdot c p^d \cdot \frac{e(\mathcal{H})}{v(\mathcal{H})} \right\} \leq c \cdot 2^k p^{k-i} \frac{e(\mathcal{H})}{v(\mathcal{H})},$$

as required. \square

Finally, we show that if \mathcal{H}_{i+1} satisfies (P3) and (P4), then either \mathcal{H}_{i+1} also satisfies (P4) or we have $|A_i| \leq (1 - c_i)v(\mathcal{H})$. Recall that $c_i = (ck2^{k+1})^{i-k}$.

Lemma 3.9. Let $i \in [k-1]$ and suppose that $e(\mathcal{H}_{i+1}) \geq c_{i+1} p^{k-(i+1)} e(\mathcal{H})$ and that $\Delta_\ell(\mathcal{H}_{i+1}) \leq \Delta_\ell^{i+1}$ for every $\ell \in [i+1]$. Then either

$$e(\mathcal{H}_i) \geq \frac{p}{c \cdot 2^{k+1} k} e(\mathcal{H}_{i+1}) \geq c_i p^{k-i} e(\mathcal{H}) \quad (9)$$

or $|A_i| \leq (1 - c_i)v(\mathcal{H})$.

Proof. If the Scythe Algorithm stops in step (1), then $|A_i| = 0$ and there is nothing to prove. Hence, we may assume that steps (2)–(4) are executed b times. Note that, for each $j \in \{0, \dots, b-1\}$, we have

$$e(\mathcal{H}_i^{(j+1)}) - e(\mathcal{H}_i^{(j)}) = \deg_{\mathcal{A}_{i+1}^{(j)}}(u_j). \quad (10)$$

By the definition of the max-deg order, the right-hand side of (10) is at least the average degree of the hypergraph $\tilde{\mathcal{A}}_{i+1}^{(j)}$, the subhypergraph of $\mathcal{A}_{i+1}^{(j)}$ induced by the set $(V(\mathcal{A}_{i+1}^{(j)}) \setminus W(u_j)) \cup \{u_j\}$. Therefore, by the definition of $\mathcal{A}_{i+1}^{(j+1)}$, we have

$$e(\mathcal{H}_i^{(j+1)}) - e(\mathcal{H}_i^{(j)}) \geq \frac{(i+1)e(\tilde{\mathcal{A}}_{i+1}^{(j)})}{v(\tilde{\mathcal{A}}_{i+1}^{(j)})} \geq \frac{(i+1)e(\mathcal{A}_{i+1}^{(j+1)})}{v(\mathcal{H})}.$$

Hence, if $(i+1)e(\mathcal{A}_{i+1}^{(j+1)}) \geq e(\mathcal{H}_{i+1})$ for every $j \in \{0, \dots, b-1\}$, then

$$e(\mathcal{H}_i) \geq \sum_{j=0}^{b-1} \frac{(i+1)e(\mathcal{A}_{i+1}^{(j+1)})}{v(\mathcal{H})} \geq b \cdot \frac{e(\mathcal{H}_{i+1})}{v(\mathcal{H})} = p \cdot e(\mathcal{H}_{i+1}),$$

since $b = p \cdot v(\mathcal{H})$, as required. Thus, we may assume that for some j ,

$$e(\mathcal{A}_{i+1}^{(b)}) \leq e(\mathcal{A}_{i+1}^{(j+1)}) < \frac{e(\mathcal{H}_{i+1})}{i+1}. \quad (11)$$

Intuitively, (11) means that while running the Scythe Algorithm on \mathcal{H}_{i+1} and I , many edges are removed from \mathcal{A}_{i+1} (that is, \mathcal{H}_{i+1}) in step (4). This may happen for one of the following two reasons: either many of the initial segments $W(u_j)$ are long or one of the families $M_\ell^i(\mathcal{H}_i)$ of sets with high degree in \mathcal{H}_i is large.

Claim. Either

$$\sum_{j=0}^{b-1} |W(u_j)| \geq \frac{1}{4\Delta_1^{i+1}} \cdot e(\mathcal{H}_{i+1})$$

or for some $\ell \in [i]$,

$$|M_\ell^i(\mathcal{H}_i)| \geq \frac{1}{2(i+1)\Delta_\ell^{i+1}} \cdot e(\mathcal{H}_{i+1}).$$

Proof of claim. Recall that $\mathcal{A}_{i+1}^{(0)} = \mathcal{H}_{i+1}$ and observe that for every $j \in \{0, \dots, b-1\}$,

$$e(\mathcal{A}_{i+1}^{(j)}) - e(\mathcal{A}_{i+1}^{(j+1)}) \leq |W(u_j)| \cdot \Delta_1(\mathcal{H}_{i+1}) + \sum_{\ell=1}^i \left| M_\ell^i(\mathcal{H}_i^{(j+1)}) \setminus M_\ell^i(\mathcal{H}_i^{(j)}) \right| \cdot \Delta_\ell(\mathcal{H}_{i+1}). \quad (12)$$

Inequality (12) follows since in step (4) of the Scythe Algorithm, we remove from $\mathcal{A}_{i+1}^{(j)}$ only the edges that contain either a vertex of $W(u_j)$ or a member of $M_\ell^i(\mathcal{H}_i^{(j+1)})$ for some $\ell \in [i]$. Thus, since $\Delta_\ell(\mathcal{H}_{i+1}) \leq \Delta_\ell^{i+1}$ for every $\ell \in [i]$, summing (12) over all j , we get

$$e(\mathcal{H}_{i+1}) - e(\mathcal{A}_{i+1}^{(b)}) \leq \sum_{j=0}^{b-1} |W(u_j)| \cdot \Delta_1^{i+1} + \sum_{\ell=1}^i \left| M_\ell^i(\mathcal{H}_i^{(b)}) \right| \cdot \Delta_\ell^{i+1}.$$

Since we assumed that $e(\mathcal{A}_{i+1}^{(b)}) < e(\mathcal{H}_{i+1})/(i+1)$, see (11), and $\mathcal{H}_i = \mathcal{H}_i^{(b)}$, it follows that if

$$\sum_{j=0}^{b-1} |W(u_j)| \cdot \Delta_1^{i+1} < \frac{e(\mathcal{H}_{i+1})}{4} \leq \frac{i}{2(i+1)} \cdot e(\mathcal{H}_{i+1}),$$

then

$$|M_\ell^i(\mathcal{H}_i)| \cdot \Delta_\ell^{i+1} \geq \frac{1}{2(i+1)} \cdot e(\mathcal{H}_{i+1}) \quad \text{for some } \ell \in [i],$$

as claimed. \square

Finally, let us deal with the two cases implied by the claim. In the remainder of the proof, we will show that if $M_\ell^i(\mathcal{H}_i)$ is large for some $\ell \in [i]$, then $e(\mathcal{H}_i)$ is large and if $\sum_{j=0}^{b-1} |W(u_j)|$ is large, then $|A_i|$ is small.

Case 1: $|M_\ell^i(\mathcal{H}_i)| \geq \frac{1}{2(i+1)\Delta_\ell^{i+1}} \cdot e(\mathcal{H}_{i+1})$ for some $\ell \in [i]$.

Since $\deg_{\mathcal{H}_i}(T) \geq \Delta_\ell^i/2$ for every $T \in M_\ell^i(\mathcal{H}_i)$, it follows by the handshaking lemma that

$$e(\mathcal{H}_i) = \binom{i}{\ell}^{-1} \sum_{T \in \binom{V(\mathcal{H}_i)}{\ell}} \deg_{\mathcal{H}_i}(T) \geq \frac{|M_\ell^i(\mathcal{H}_i)| \cdot \Delta_\ell^i}{2 \binom{i}{\ell}}. \quad (13)$$

Recalling that $\Delta_\ell^i \geq p\Delta_\ell^{i+1}$, see (7), we have

$$e(\mathcal{H}_i) \geq \frac{e(\mathcal{H}_{i+1})}{4(i+1)\binom{i}{\ell}} \cdot \frac{\Delta_\ell^i}{\Delta_\ell^{i+1}} \geq \frac{p}{2^{i+2}(i+1)} \cdot e(\mathcal{H}_{i+1}) \geq \frac{p}{2^{k+1}k} \cdot e(\mathcal{H}_{i+1}),$$

as required.

Case 2: $\sum_{j=0}^{b-1} |W(u_j)| \geq \frac{1}{4\Delta_1^{i+1}} \cdot e(\mathcal{H}_{i+1})$.

We claim that in this case, $|A_i| \leq (1 - c_i)v(\mathcal{H})$. Indeed, we have

$$v(\mathcal{H}) - |A_i| = v(\mathcal{A}_{i+1}^{(0)}) - v(\mathcal{A}_{i+1}^{(b)}) = \sum_{j=0}^{b-1} |W(u_j)| \geq \frac{e(\mathcal{H}_{i+1})}{4\Delta_1^{i+1}}.$$

Recall that $\Delta_1^{i+1} \leq c2^k p^{k-i-1} \frac{e(\mathcal{H})}{v(\mathcal{H})}$ by Lemma 3.8. Thus,

$$v(\mathcal{H}) - |A_i| \geq \frac{p^{i+1-k}}{c2^{k+2}} \cdot \frac{v(\mathcal{H})}{e(\mathcal{H})} \cdot e(\mathcal{H}_{i+1}) \geq c_i v(\mathcal{H}),$$

since $e(\mathcal{H}_{i+1}) \geq c_{i+1} p^{k-(i+1)} e(\mathcal{H})$ and $c_{i+1}/(c2^{k+2}) \geq c_i$. \square

3.2. The proof of Proposition 3.1 and Theorem 2.2.

Proof of Proposition 3.1. Let k be an integer and let c be a positive constant. Furthermore, let $p \in (0, 1)$ and let \mathcal{H} be a k -uniform hypergraph that satisfy the assumptions of Proposition 3.1. Let $\delta = (ck2^{k+1})^{-k}$ and $b = pv(\mathcal{H})$. We will use the Scythe Algorithm, described in Section 3.1, to construct a family \mathcal{S} and functions f_0 and g_0 as in the statement of Proposition 3.1. We obtain them by running the following algorithm (with $\mathcal{H}_k = \mathcal{H}$) on every independent set $I \in \mathcal{I}(\mathcal{H})$. We shall define f_0 somewhat implicitly by defining a function $f_0^*: \mathcal{I}(\mathcal{H}) \rightarrow \mathcal{P}(V(\mathcal{H}))$ that is constant on the set $g_0^{-1}(S)$ for every $S \in \mathcal{S}$.

Constructing g_0 and f_0^* . Given an $I \in \mathcal{I}(\mathcal{H})$, set $i = k - 1$ and repeat the following:

- (1) Apply the Scythe Algorithm to \mathcal{H}_{i+1} and I . Suppose that it outputs \mathcal{H}_i , A_i and B_i .
- (2) If $|A_i| \leq (1 - \delta)v(\mathcal{H})$, then set $q = i$, $r = i + 1$ and STOP.
- (3) If $i > 1$, then set $i = i - 1$. Otherwise, set $q = r = 1$ and STOP.

Let I be an independent set and let us execute the above procedure (with $\mathcal{H}_k = \mathcal{H}$) on I . We claim that for every $i \in \{r, \dots, k\}$, the hypergraph \mathcal{H}_i satisfies properties (P1)–(P4) defined in Section 3.1. This follows by induction on $k - i$. The base of the induction, the case $i = k$, follows vacuously from the definitions of c_k and Δ_ℓ^k for $\ell \in [k]$. The inductive step follows from Lemmas 3.5, 3.7, and 3.9. To see this, note that since $|A_i| > (1 - \delta)v(\mathcal{H}) \geq (1 - c_i)v(\mathcal{H})$ for all $i \in \{r, \dots, k - 1\}$, then (9) in Lemma 3.9 always holds.

Now, let us define $g_0(I)$ and $f_0^*(I)$. Suppose first that $r > 1$ and note that in this case, the algorithm stopped in step (2), which means that $|A_q| \leq (1 - \delta)v(\mathcal{H})$; we set

$$g_0(I) = B_{k-1} \cup \dots \cup B_q \quad \text{and} \quad f_0^*(I) = A_q.$$

On the other hand, if $r = 1$, then we set

$$g_0(I) = B_{k-1} \cup \dots \cup B_1 \quad \text{and} \quad f_0^*(I) = \{v \in V(\mathcal{H}_1) : \{v\} \notin \mathcal{H}_1\}.$$

Finally, we let

$$\mathcal{S} = \{g_0(I) : I \in \mathcal{I}(\mathcal{H})\}.$$

We will define f_0 by letting $f_0(S) = f_0^*(I)$ for some $I \in g_0^{-1}(S)$. We first show that this definition will not depend on the choice of I . In fact, we shall prove a slightly stronger statement, which also establishes the consistency property of g_0 stated in the final line of Proposition 3.1.

Claim. Suppose that for some $I, I' \in \mathcal{I}(\mathcal{H})$, $g_0(I) \subseteq I'$ and $g_0(I') \subseteq I$. Then $g_0(I) = g_0(I')$ and $f_0^*(I) = f_0^*(I')$.

Proof of claim. Suppose that while running the algorithm on some I , we obtain a sequence (B_{k-1}, \dots, B_q) . Since $g_0(I)$ depends solely on (B_{k-1}, \dots, B_q) and, by Lemma 3.5, for each i , the hypergraph \mathcal{H}_i and the set A_i depend only on (B_{k-1}, \dots, B_i) , then also $f_0^*(I)$ depends solely on (B_{k-1}, \dots, B_q) . Hence, it suffices to show that if, while running the algorithm on some I' with $B_{k-1} \cup \dots \cup B_q \subseteq I'$, we obtain a sequence $(B'_{k-1}, \dots, B'_{q'})$ with $B'_{k-1} \cup \dots \cup B'_{q'} \subseteq I$, then $(B'_{k-1}, \dots, B'_{q'}) = (B_{k-1}, \dots, B_q)$. To this end, let us first observe that, under the above assumptions, for every $i \in [k-1]$, if $\mathcal{H}_{i+1} = \mathcal{H}'_{i+1}$, then $B_i = B'_i$. Indeed, note that B_i and B'_i are the outputs of the Scythe Algorithm executed on the inputs (\mathcal{H}_{i+1}, I) and (\mathcal{H}'_{i+1}, I') , respectively. Hence, if $\mathcal{H}_{i+1} = \mathcal{H}'_{i+1}$, then since

$$B_i \subseteq B_{k-1} \cup \dots \cup B_q \subseteq I' \quad \text{and} \quad B'_i \subseteq B'_{k-1} \cup \dots \cup B'_{q'} \subseteq I,$$

then Lemma 3.6 implies that $B_i = B'_i$. Since clearly $\mathcal{H}_k = \mathcal{H}'_k = \mathcal{H}$ and, as noted before, for each i , \mathcal{H}_{i+1} depends only on $(B_{k-1}, \dots, B_{i+1})$, it follows that $B_i = B'_i$ for all i , as required. \square

By the above claim, we can define f_0 by letting, for every $S \in \mathcal{S}$, $f_0(S) = f_0^*(I)$ for any $I \in g_0^{-1}(S)$. Finally, let us show that the \mathcal{S} , g_0 , and f_0 , which we have just defined, satisfy the required conditions, that is, for all $I, I' \in \mathcal{I}(\mathcal{H})$,

- (i) $|S| \leq (k-1)pv(\mathcal{H})$ for every $S \in \mathcal{S}$,
- (ii) $g_0(I) \subseteq I \subseteq f_0(g_0(I)) \cup g_0(I)$,
- (iii) $|f_0(g_0(I))| \leq (1-\delta)v(\mathcal{H})$,
- (iv) $g_0(I) \subseteq I'$ and $g_0(I') \subseteq I$ imply that $g_0(I) = g_0(I')$.

To see (i), simply recall that $|B_i| \leq pv(\mathcal{H})$ for every $i \in [k-1]$. To see (ii), note that $B_i \subseteq I \subseteq A_i \cup B_i$ for every $i \in \{q, \dots, k-1\}$, by Lemma 3.5, that I is an independent set in \mathcal{H}_1 (if $r = 1$) and, crucially, that $f_0(g_0(I)) = f_0^*(I)$. To see (iii), note that if $r > 1$, then $|A_q| \leq (1-\delta)v(\mathcal{H})$, see step (2) of the algorithm; if $r = 1$, then since $\Delta_1(\mathcal{H}_1) \leq c \cdot 2^k p^{k-1} e(\mathcal{H})/v(\mathcal{H})$, by Lemma 3.8 and property (P3), we have

$$|\{v \in V(\mathcal{H}_1) : \{v\} \in \mathcal{H}_1\}| \geq \frac{e(\mathcal{H}_1)}{\Delta_1(\mathcal{H}_1)} \geq \frac{c_1 p^{k-1} e(\mathcal{H})}{\Delta_1(\mathcal{H}_1)} \geq \delta v(\mathcal{H}),$$

since $\delta \leq c_1/(c \cdot 2^k)$ and \mathcal{H}_1 satisfies property (P4), so $e(\mathcal{H}_1) \geq c_1 p^{k-1} e(\mathcal{H})$. Finally, (iv) follows directly from the claim. \square

Proof of Theorem 2.2. The theorem follows by applying Proposition 3.1 a bounded number of times. Given an integer k and positive reals c and ε , let $\delta = \delta_{3.1}(c/\varepsilon)$ and let

$$C = (k-1) \cdot \left(\frac{1}{\delta} \log \frac{1}{\varepsilon} + 1 \right).$$

Let V be a finite set and let \mathcal{F} be an increasing family of subsets of V such that $|A| \geq \varepsilon|V|$ for every $A \in \mathcal{F}$. Let $p \in (0, 1)$ and suppose that \mathcal{H} is a k -uniform hypergraph on the vertex set V that is $(\mathcal{F}, \varepsilon)$ -dense and satisfies the assumptions of the theorem, that is,

$$\Delta_\ell(\mathcal{H}) \leq cp^{\ell-1} \frac{e(\mathcal{H})}{v(\mathcal{H})}$$

for every $\ell \in [k]$. We now show how to construct a family $\mathcal{S} \subseteq \binom{V(\mathcal{H})}{\leq Cp v(\mathcal{H})}$ and functions $f: \mathcal{S} \rightarrow \overline{\mathcal{F}}$ and $g: \mathcal{I}(\mathcal{H}) \rightarrow \mathcal{S}$ such that

$$g(I) \subseteq I \quad \text{and} \quad I \setminus g(I) \subseteq f(g(I)) \tag{14}$$

for every $I \in \mathcal{I}(\mathcal{H})$. Similarly as in the proof of Proposition 3.1, we shall define f via a function $f^*: \mathcal{I}(\mathcal{H}) \rightarrow \mathcal{P}(V)$ that is constant on each set $g^{-1}(S)$ with $S \in \mathcal{S}$.

Fix some $I \in \mathcal{I}(\mathcal{H})$. Using Proposition 3.1, we shall construct (for some $J \leq \frac{1}{\delta} \log \frac{1}{\varepsilon} + 1$) a sequence $(A_j, S_j)_{j=1}^J$ of pairs of subsets of V such that for each $j \in [J]$,

$$S_1 \cup \dots \cup S_j \subseteq I \subseteq A_j \cup S_1 \cup \dots \cup S_j.$$

Moreover, $A_j \in \overline{\mathcal{F}}$ while $|S_1 \cup \dots \cup S_j| \leq Cp v(\mathcal{H})$. Crucially, the set A_j will depend solely on $S_1 \cup \dots \cup S_j$. We will let $g(I) = S_1 \cup \dots \cup S_J$ and $f^*(I) = A_J$.

Construction. Let $S_0 = \emptyset$ and let $A_0 = V$. For $j = 0, 1, \dots$, do the following:

- (1) If $A_j \in \mathcal{F}$, then let $I_j = I \cap A_j$ and apply Proposition 3.1 with $c_{3.1} = c/\varepsilon$ and $p_{3.1} = p$ to the hypergraph $\mathcal{H}[A_j]$ and the set I_j to obtain sets $g_0(I_j)$ and $f_0(g_0(I_j))$ such that $g_0(I_j) \subseteq I_j$ and $I_j \setminus g_0(I_j) \subseteq f_0(g_0(I_j))$. Otherwise, if $A_j \in \overline{\mathcal{F}}$, then STOP.
- (2) Let $S_{j+1} = g_0(I_j)$ and let $A_{j+1} = f_0(g_0(I_j))$.

Let us first show that the above procedure is well-defined, that is, that the assumptions of Proposition 3.1 are satisfied each time we are in (1). To this end, fix some $A \subseteq V$ and note that if $A \in \mathcal{F}$, then, since \mathcal{H} is $(\mathcal{F}, \varepsilon)$ -dense,

$$\Delta_\ell(\mathcal{H}[A]) \leq \Delta_\ell(\mathcal{H}) \leq cp^{\ell-1} \frac{e(\mathcal{H})}{v(\mathcal{H})} \leq \frac{c}{\varepsilon} \cdot p^{\ell-1} \frac{e(\mathcal{H}[A])}{v(\mathcal{H}[A])},$$

where the last step follows since $e(\mathcal{H}[A]) \geq \varepsilon \cdot e(\mathcal{H})$ and $v(\mathcal{H}[A]) \leq v(\mathcal{H})$.

Next, let us show that the above procedure terminates, therefore producing a finite sequence (A_j, S_j) with $j \in [J]$. To this end, let us simply note that by Proposition 3.1, $|A_{j+1}| \leq (1 - \delta)|A_j|$ for all j , $A_0 = V$ and $|A| \geq \varepsilon|V|$ for every $A \in \mathcal{F}$. Moreover, since $A_{J-1} \in \mathcal{F}$, then

$$\varepsilon|V| \leq |A_{J-1}| \leq (1 - \delta)^{J-1} |A_0| = \exp(-(J-1)\delta)|V|$$

and hence $J \leq \frac{1}{\delta} \log \frac{1}{\varepsilon} + 1$. It immediately follows that

$$|g(I)| \leq \sum_{j=1}^J |S_j| \leq \sum_{j=1}^J (k-1)pv(\mathcal{H}[A_j]) \leq J(k-1)pv(\mathcal{H}) \leq Cpv(\mathcal{H}).$$

Finally, let $\mathcal{S} = \{g(I) : I \in \mathcal{I}(\mathcal{H})\}$. It remains to show that for every $S \in \mathcal{S}$, f^* is constant on $g^{-1}(S)$. Similarly as in the proof of Proposition 3.1, we shall prove a somewhat stronger statement.

Claim. Suppose that for some $I, I' \in \mathcal{I}(\mathcal{H})$, $g(I) \subseteq I'$ and $g(I') \subseteq I$. Then $g(I) = g(I')$ and $f^*(I) = f^*(I')$.

Proof of claim. Suppose that while running the above procedure on some I , we generate a sequence $(A_j, S_j)_{j=1}^J$. Since for each j , A_{j+1} depends solely on A_j and S_{j+1} , where $A_0 = V$, then both $g(I)$ and $f^*(I)$ depend solely on (S_1, \dots, S_J) . Hence, it suffices to show that if, while running the above procedure on some I' with $S_1 \cup \dots \cup S_J \subseteq I'$, we generate a sequence $(A'_j, S'_j)_{j=1}^{J'}$ with $S'_1 \cup \dots \cup S'_{J'} \subseteq I$, then $(S_1, \dots, S_J) = (S'_1, \dots, S'_{J'})$. To this end, it suffices to note that if $A_j = A'_j$, then, since

$$S_{j+1} \subseteq S_1 \cup \dots \cup S_J \subseteq I' \quad \text{and} \quad S'_{j+1} \subseteq S'_1 \cup \dots \cup S'_{J'} \subseteq I,$$

by the consistency property of g_0 stated in the final line of Proposition 3.1, $S_{j+1} = S'_{j+1}$. Since $A_0 = A'_0 = V$ and for each j , A_j depends only on (S_1, \dots, S_j) , it follows that $S_j = S'_j$ for all j , as required. \square

Finally, for every $S \in \mathcal{S}$, we let $f(S) = f^*(I)$ for some $I \in g^{-1}(S)$. This completes the proof of Theorem 2.2. \square

4. SZEMERÉDI'S THEOREM FOR SPARSE SETS

In this section, we prove Theorem 1.1 and derive from it Corollary 1.2. Before we get to the proofs, let us first remark that Theorem 1.1 and Corollary 1.2 are both sharp up to the value of the constant C in the lower bounds for p and m . More precisely, let us make the following two observations.

- (1) For every $\beta \in (0, 1)$, there is a positive c such that if $m \leq cn^{1-1/(k-1)}$, then the number of m -subsets of $[n]$ that contain no k -term AP is at least $(1 - \beta)^m \binom{n}{m}$. To see this, let $\varepsilon = \beta^2$ and observe that if c is sufficiently small and $m \leq cn^{1-1/(k-1)}$, then the expected number of k -term APs in a random $(1 + \varepsilon)m$ -subset of $[n]$ is smaller than $\varepsilon m/2$ and hence by Markov's inequality, at least half of all $(1 + \varepsilon)m$ -subsets of $[n]$ contain a subset of size m with no k -term AP. Hence⁸

$$\#\{m\text{-subsets of } [n] \text{ with no } k\text{-term AP}\} \geq \frac{\binom{n}{(1+\varepsilon)m}}{2\binom{n}{\varepsilon m}} \geq (1 - \sqrt{\varepsilon})^m \binom{n}{m},$$

where the final inequality holds since $\binom{n}{(1+\varepsilon)m} \geq \left(\frac{n}{2m}\right)^{\varepsilon m} \binom{n}{m}$ and $\binom{n}{\varepsilon m} \leq \left(\frac{en}{\varepsilon m}\right)^{\varepsilon m}$.

⁸We assume here, without loss of generality, that β (and hence also ε) is sufficiently small.

(2) There is a positive constant c such that if $p_n \leq cn^{-1/(k-1)}$, then

$$\mathbb{P}([n]_{p_n} \text{ is } (\delta, k)\text{-Szemerédi}) \rightarrow 0 \text{ as } n \rightarrow \infty.$$

For a (simple) proof of this statement, we refer the reader to [58].

We shall in fact prove the following somewhat stronger version of Corollary 1.2, originally proved by Schacht [58] (the approach of Conlon and Gowers [14] yields a somewhat weaker probability estimate).

Corollary 4.1. *For every $k \in \mathbb{N}$ and every $\delta \in (0, 1)$, there exists a constant C such that for all sufficiently large n , if $p \geq Cn^{-1/(k-1)}$, then*

$$\mathbb{P}([n]_p \text{ is } (\delta, k)\text{-Szemerédi}) \geq 1 - 2 \exp(-pn/8).$$

In the proofs of Theorem 1.1 and Corollary 4.1, and frequently in later sections, we shall need various estimates on binomial coefficients, which we list here for future reference. Let a , b , and c be integers satisfying $a \geq b \geq c \geq 0$. Then the following inequalities hold:

$$\binom{a}{b} \leq \left(\frac{ea}{b}\right)^b, \quad (15) \quad \binom{b}{c} \leq \left(\frac{b}{a}\right)^c \binom{a}{c}, \quad (16)$$

$$\binom{a}{b-c} \leq \left(\frac{b}{a-b}\right)^c \binom{a}{b}, \quad (17) \quad \binom{a}{c} \leq \left(\frac{a-c}{b-c}\right)^c \binom{b}{c}. \quad (18)$$

We remark that each inequality above follows easily from the definition of $\binom{a}{b}$.

Proof of Corollary 4.1. Fix $k \in \mathbb{N}$ and $\delta \in (0, 1)$, let $\beta = \delta/(2e) \cdot e^{-1/\delta}$, and set $C = 2C_{1.1}(\beta, k)/\delta$. Assume that $p \geq Cn^{-1/(k-1)}$, let $m = \delta pn/2$, and let X_m denote the number of m -subsets of $[n]_p$ that contain no k -term AP. By Theorem 1.1 and (15), we have

$$\mathbb{P}(X_m > 0) \leq \mathbb{E}[X_m] \leq \binom{\beta n}{m} p^m \leq \left(\frac{\beta e p n}{m}\right)^m = \left(\frac{2\beta e}{\delta}\right)^m = e^{-m/\delta}. \quad (19)$$

Let \mathcal{A} denote the event that $[n]_p$ is *not* (δ, k) -Szemerédi, i.e., that $[n]_p$ contains a subset with $\delta|[n]_p|$ elements and no k -term AP. By (19) and Chernoff's inequality (see, e.g., [3, Appendix A]), it follows that

$$\mathbb{P}(\mathcal{A}) \leq \mathbb{P}\left(\mathcal{A} \wedge |[n]_p| \geq \frac{pn}{2}\right) + \mathbb{P}\left(|[n]_p| < \frac{pn}{2}\right) \leq \mathbb{P}(X_m > 0) + e^{-pn/8} \leq 2e^{-pn/8},$$

as required. \square

Finally, let us show how to deduce Theorem 1.1 from Theorem 2.2. Our proof will use the following robust version of Szemerédi's theorem, which can be proved by a simple averaging argument, originally observed by Varnavides [65].

Lemma 4.2. *For every positive δ and $k \in [n]$, there exists a positive ε such that the following holds for all sufficiently large n . Every subset of $[n]$ with at least δn elements contains at least εn^2 k -term APs.*

Proof of Theorem 1.1. Given $k \in \mathbb{N}$ and positive β , let $\delta = \min\{\beta/2, 1/10\}$ and let $n \in \mathbb{N}$ be sufficiently large. Let \mathcal{H} be the k -uniform hypergraph of k -term APs in $[n]$, i.e., the hypergraph on the vertex set $[n]$ whose edges are all k -term APs in $[n]$, let \mathcal{F} denote the

family of subsets of $[n]$ with at least δn elements, and let $\varepsilon = \varepsilon_{4.2}(\delta, k)$. By Lemma 4.2, the hypergraph \mathcal{H} is $(\mathcal{F}, \varepsilon)$ -dense, provided that n is sufficiently large. Let $p = n^{-1/(k-1)}$ and let $c = 2k^2$. Since $e(\mathcal{H}) \geq n^2/k^2 \geq 2n^2/c$, it follows that

$$\Delta_1(\mathcal{H}) \leq k \cdot \frac{n}{k-1} \leq 2n \leq c \cdot p^{1-1} \frac{e(\mathcal{H})}{v(\mathcal{H})},$$

for every $\ell \in \{2, \dots, k-1\}$,

$$\Delta_\ell(\mathcal{H}) \leq \Delta_2(\mathcal{H}) \leq \binom{k}{2} \leq 2n^{1/(k-1)} \leq c \cdot p^{\ell-1} \frac{e(\mathcal{H})}{v(\mathcal{H})},$$

and $\Delta_k(\mathcal{H}) = 1 \leq c \cdot p^{k-1} \frac{e(\mathcal{H})}{v(\mathcal{H})}$.

Let $C' = C_{2.2}(k, \varepsilon, c)$, let $C = C'/\delta$, and assume that $m \geq Cn^{1-1/(k-1)} = Cpn$. Note that if $m > \delta n/2$, then $\mathcal{I}(\mathcal{H}, m) = 0$ by Szemerédi's theorem, so we may assume that $m \leq \delta n/2$. Since $C'pn \leq \delta m$, then by Theorem 2.2, there exists a family $\mathcal{S} \subseteq \binom{[n]}{\leq C'pn} \subseteq \binom{[n]}{\leq \delta m}$ and functions $f: \mathcal{S} \rightarrow \overline{\mathcal{F}}$ and $g: \mathcal{I}(\mathcal{H}) \rightarrow \mathcal{S}$, such that for every $I \in \mathcal{I}(\mathcal{H})$,

$$g(I) \subseteq I \quad \text{and} \quad I \setminus g(I) \subseteq f(g(I)).$$

Therefore, using (15) and (17), the number of independent sets of size m in \mathcal{H} can be estimated as follows:

$$\begin{aligned} |\mathcal{I}(\mathcal{H}, m)| &= \sum_{S \in \mathcal{S}} |\{I \in \mathcal{I}(\mathcal{H}, m) : g(I) = S\}| \leq \sum_{S \in \mathcal{S}} \binom{|f(S)|}{m - |S|} \\ &\leq \sum_{k \leq \delta m} \binom{n}{k} \binom{\delta n}{m-k} \leq \sum_{k \leq \delta m} \left(\frac{en}{k}\right)^k \left(\frac{m}{\delta n - m}\right)^k \binom{\delta n}{m}. \end{aligned}$$

Since $m \leq \delta n/2$ and the function $x \mapsto (y/x)^x$ is increasing on $(0, y/e)$, it follows that

$$|\mathcal{I}(\mathcal{H}, m)| \leq \sum_{k \leq \delta m} \left(\frac{2em}{\delta k}\right)^k \binom{\delta n}{m} \leq m \left(\frac{2e}{\delta^2}\right)^{\delta m} \binom{\delta n}{m} \leq \binom{\beta n}{m},$$

where the final inequality follows since $\binom{\delta n}{m} \leq 2^{-m} \binom{2\delta n}{m}$, by (16), and since $2^{1/\delta} > 2e/\delta^2$ if $\delta \leq 1/10$. This proves Theorem 1.1. \square

The same proof, combined with an analogue of Lemma 4.2 due to Furstenberg and Katznelson [27], yields the following generalization of Theorem 1.1, which strengthens both [14, Theorem 10.4] and [58, Theorem 2.3]. Given a set $F \subseteq \mathbb{N}^\ell$, we call a set of the form $a + bF = \{a + bx : x \in F\}$, with $a \in \mathbb{N}^\ell$ and $b \in \mathbb{Z} \setminus \{0\}$, a *homothetic copy* of F .

Theorem 4.3. *For every positive β , every $\ell \in \mathbb{N}$, and every finite configuration $F \subseteq \mathbb{N}^\ell$, there exist constants C and n_0 such that the following holds. For every $n \in \mathbb{N}$ with $n \geq n_0$, if $m \geq Cn^{\ell-1/(|F|-1)}$, then there are at most*

$$\binom{\beta n^\ell}{m}$$

m -subsets of $[n]^\ell$ that contain no homothetic copy of F .

Finally, using the famous polynomial Szemerédi theorem of Bergelson and Leibman [13], the same argument gives a counting version of [14, Theorem 10.7].

Theorem 4.4. *For every positive β and integers k and r , there exist constants C and n_0 such that the following holds. For every $n \in \mathbb{N}$ with $n \geq n_0$, if $m \geq Cn^{1-1/kr}$, then there are at most*

$$\binom{\beta n}{m}$$

m -subsets of $[n]$ that contain no set of the form $\{a, a + d^r, \dots, a + kd^r\}$.

5. EXTREMAL RESULTS FOR SPARSE SETS

In this section, we shall deduce from Theorem 2.2 two versions of the general transference theorem of Schacht [58, Theorem 3.3]. We remind the reader that a statement very similar to Schacht’s theorem was proved independently by Conlon and Gowers [14]. For the benefit of the readers who are familiar with [58], we shall state it using the terminology used there.

Definition 5.1. Let $\mathcal{H} = (\mathcal{H}_n)_{n \in \mathbb{N}}$ be a sequence of k -uniform hypergraphs and let $\alpha \in [0, 1)$. We say that \mathcal{H} is α -dense if the following is true: For every positive δ , there exist positive ε and n_0 such that for every n with $n \geq n_0$ and every $U \subseteq V(\mathcal{H}_n)$ with $|U| \geq (\alpha + \delta)v(\mathcal{H}_n)$, we have

$$e(\mathcal{H}_n[U]) \geq \varepsilon e(\mathcal{H}_n).$$

Let us remark here that Definition 2.1 is a generalization of Definition 5.1. Indeed, if \mathcal{F}_δ denotes the collection of all subsets of $V(\mathcal{H}_n)$ with at least $(\alpha + \delta)v(\mathcal{H}_n)$ elements, then a sequence \mathcal{H} of hypergraphs is α -dense if and only if for every positive δ , there exists a positive ε such that for all sufficiently large n , the hypergraph \mathcal{H}_n is $(\mathcal{F}_\delta, \varepsilon)$ -dense.

We start with the ‘random’ version of our extremal result, which was originally proved by Schacht [58, Theorem 3.3].

Theorem 5.2. *Let \mathcal{H} be a sequence of k -uniform hypergraphs, let $\alpha \in [0, 1)$, and let c be a positive constant. Suppose that $\mathbf{p} \in [0, 1]^\mathbb{N}$ is a sequence of probabilities such that for all sufficiently large $n \in \mathbb{N}$, and for every $\ell \in [k]$, we have*

$$\Delta_\ell(\mathcal{H}_n) \leq c \cdot p_n^{\ell-1} \frac{e(\mathcal{H}_n)}{v(\mathcal{H}_n)}. \tag{20}$$

If \mathcal{H} is α -dense, then the following holds. For every positive δ , there exists a constant C such that if $q_n \geq Cp_n$ and $q_n v(\mathcal{H}_n) \rightarrow \infty$ as $n \rightarrow \infty$, then a.a.s.

$$\alpha(\mathcal{H}_n[V(\mathcal{H}_n)_{q_n}]) \leq (\alpha + \delta)q_n v(\mathcal{H}_n).$$

We note that the probability bounds implicit in the ‘asymptotically almost surely’ statement that we obtain are, as in [58], optimal, that is, they decay exponentially in $p_n v(\mathcal{H}_n)$.

Remark 5.3. We remark that the only difference between Theorem 5.2 and [58, Theorem 3.3] are the assumptions on the hypergraph sequence \mathcal{H} . It turns out that this difference is only superficial, since condition (20) is essentially equivalent to the condition that \mathcal{H} is (K, \mathbf{p}) -bounded (see [58]). One easily checks that if \mathcal{H}_n satisfies (20) for sufficiently large n ,

then \mathcal{H} is (K, \mathbf{p}) -bounded for some constant K that depends only on c and k . Conversely, if \mathcal{H} is (K, \mathbf{p}) -bounded, then for all sufficiently large n , there is an $\mathcal{H}'_n \subseteq \mathcal{H}_n$ with at least $(1 - \varepsilon)e(\mathcal{H}_n)$ edges that satisfies (20) for some constant c that depends only on ε , k , and K . One obtains such \mathcal{H}'_n by repeatedly deleting from \mathcal{H}_n edges that contain an ℓ -set T with $\deg_{\mathcal{H}}(T) > c \cdot p_n^{\ell-1} e(\mathcal{H}_n) / v(\mathcal{H}_n)$. Finally, note that, trivially, if \mathcal{H}_n is $(\mathcal{F}, 2\varepsilon)$ -dense for some family $\mathcal{F} \subseteq \mathcal{P}(V(\mathcal{H}_n))$, then every \mathcal{H}'_n with $e(\mathcal{H}'_n) \geq (1 - \varepsilon)e(\mathcal{H}_n)$ is $(\mathcal{F}, \varepsilon)$ -dense.

Our methods also yield the following ‘counting’ analogue of Theorem 5.2. This generalizes Theorem 1.1, and does not follow from the methods of [14] or [58]. In the case $\alpha = 0$, it can be thought of as a strengthening of Theorem 5.2, see Corollary 1.2.

Theorem 5.4. *Let \mathcal{H} be a sequence of k -uniform hypergraphs, let $\alpha \in [0, 1)$, and let c be a positive constant. Suppose that $\mathbf{p} \in [0, 1]^{\mathbb{N}}$ is a sequence of probabilities such that for all sufficiently large $n \in \mathbb{N}$, and for every $\ell \in [k]$, we have*

$$\Delta_{\ell}(\mathcal{H}_n) \leq c \cdot p_n^{\ell-1} \frac{e(\mathcal{H}_n)}{v(\mathcal{H}_n)}.$$

If \mathcal{H} is α -dense, then the following holds. For every positive δ , there exists a constant C such that for all sufficiently large n , if $m \geq Cp_n v(\mathcal{H}_n)$, then

$$|\mathcal{I}(\mathcal{H}_n, m)| \leq \binom{(\alpha + \delta)v(\mathcal{H}_n)}{m}.$$

Proof of Theorem 5.2. Let $\alpha \in [0, 1)$, let $k \in \mathbb{N}$, let $\mathbf{p} \in [0, 1]^{\mathbb{N}}$, let $c \in (0, \infty)$, and let \mathcal{H} be a sequence of k -uniform hypergraphs as in the statement of Theorem 5.2. Furthermore, suppose that \mathcal{H} is α -dense and fix some positive δ ; without loss of generality, we may assume that δ is sufficiently small. Let $n \in \mathbb{N}$ be sufficiently large, let $\delta' = \delta/3$, and let \mathcal{F} denote the family of all subsets of $V(\mathcal{H}_n)$ with at least $(\alpha + \delta')v(\mathcal{H}_n)$ elements. Since \mathcal{H}_n is α -dense, it follows that \mathcal{H}_n is $(\mathcal{F}, \varepsilon)$ -dense for some small positive ε that does not depend on n . Let $C' = C_{2.2}(k, \varepsilon, c)$. By Theorem 2.2, there exist a family $\mathcal{S} \subseteq \binom{V(\mathcal{H}_n)}{\leq C' p_n v(\mathcal{H}_n)}$ and functions $f: \mathcal{S} \rightarrow \overline{\mathcal{F}}$ and $g: \mathcal{I}(\mathcal{H}_n) \rightarrow \mathcal{S}$ such that

$$g(I) \subseteq I \quad \text{and} \quad I \setminus g(I) \subseteq f(g(I))$$

for every $I \in \mathcal{I}(\mathcal{H}_n)$. Let $C = C' / \delta^3$ and assume that $q_n \geq Cp_n$. Let $m = (\alpha + \delta)q_n v(\mathcal{H}_n)$ and, for the sake of brevity, let us write $V = V(\mathcal{H}_n)$ and $q = q_n$. Observe that

$$\begin{aligned} \mathbb{P}\left(\alpha(\mathcal{H}_n[V_q]) \geq m\right) &= \mathbb{P}\left(I \subseteq V_q \text{ for some } I \in \mathcal{I}(\mathcal{H}_n, m)\right) \\ &\leq \sum_{S \in \mathcal{S}} \mathbb{P}\left(I \subseteq V_q \text{ for some } I \in \mathcal{I}(\mathcal{H}_n, m) \text{ such that } g(I) = S\right). \end{aligned} \tag{21}$$

Fix an $S \in \mathcal{S}$ and let $\mathcal{I}'_S = \{I \in \mathcal{I}(\mathcal{H}_n, m): g(I) = S\}$. We estimate the summand in the right-hand side of (21) as follows:

$$\mathbb{P}\left(I \subseteq V_q \text{ for some } I \in \mathcal{I}'_S\right) \leq \mathbb{P}(S \subseteq V_q) \cdot \mathbb{P}\left(|V_q \cap f(S)| \geq m - |S|\right). \tag{22}$$

To see the above inequality, simply note that for every $I \in \mathcal{I}'_S$, we have $I \setminus S \subseteq f(S)$.

Now, since $m = (\alpha + 3\delta')q_nv(\mathcal{H}_n)$ and $S \in \mathcal{S}$, then

$$|S| \leq C'p_nv(\mathcal{H}_n) \leq \delta^3qv(\mathcal{H}_n) \leq \delta'qv(\mathcal{H}_n)$$

and hence $m - |S| \geq (\alpha + 2\delta')qv(\mathcal{H}_n)$. On the other hand, since $|f(S)| \leq (\alpha + \delta')v(\mathcal{H}_n)$ by the definition of \mathcal{F} , then

$$\mathbb{E}[|V_q \cap f(S)|] \leq (\alpha + \delta')qv(\mathcal{H}_n).$$

Hence, by Chernoff's inequality, we have

$$\mathbb{P}(|V_q \cap f(S)| \geq m - |S|) \leq \exp\left(-\frac{(\delta')^2qv(\mathcal{H}_n)}{4}\right) = \exp\left(-\frac{\delta^2qv(\mathcal{H}_n)}{36}\right). \quad (23)$$

Finally, note that since $|S| \leq \delta^3qv(\mathcal{H}_n)$ for every $S \in \mathcal{S}$, and using (15),

$$\sum_{S \in \mathcal{S}} \mathbb{P}(S \subseteq V_q) \leq \sum_{s=0}^{\delta^3qv(\mathcal{H}_n)} \binom{v(\mathcal{H}_n)}{s} q^s \leq v(\mathcal{H}_n) \cdot \left(\frac{e}{\delta^3}\right)^{\delta^3qv(\mathcal{H}_n)}. \quad (24)$$

Putting (21), (22), (23), and (24) together, we obtain

$$\mathbb{P}(\alpha(\mathcal{H}_n[V_q]) \geq m) \leq \sum_{S \in \mathcal{S}} \mathbb{P}(S \subseteq V_q) \exp\left(-\frac{\delta^2qv(\mathcal{H}_n)}{36}\right) \leq \exp(-\delta^3qv(\mathcal{H}_n)),$$

as required. \square

Proof of Theorem 5.4. Let $\alpha \in [0, 1)$, let $k \in \mathbb{N}$, let $\mathbf{p} \in [0, 1]^{\mathbb{N}}$, let $c \in (0, \infty)$, and let \mathcal{H} be a sequence of k -uniform hypergraphs as in the statement of Theorem 5.2. Furthermore, suppose that \mathcal{H} is α -dense and fix some positive δ . Let n be sufficiently large, let $\delta' = \delta/2$, and let \mathcal{F} denote the family of all subsets of $V(\mathcal{H}_n)$ with at least $(\alpha + \delta')v(\mathcal{H}_n)$ elements. Since \mathcal{H}_n is α -dense, it follows that \mathcal{H}_n is $(\mathcal{F}, \varepsilon)$ -dense for some small positive ε that does not depend on n . Let $C' = C_{2.2}(k, \varepsilon, c)$. By Theorem 2.2, there exist a family $\mathcal{S} \subseteq \binom{V(\mathcal{H}_n)}{\leq C'p_nv(\mathcal{H}_n)}$ and functions $f: \mathcal{S} \rightarrow \overline{\mathcal{F}}$ and $g: \mathcal{I}(\mathcal{H}_n) \rightarrow \mathcal{S}$ such that

$$g(I) \subseteq I \quad \text{and} \quad I \setminus g(I) \subseteq f(g(I))$$

for every $I \in \mathcal{I}(\mathcal{H}_n)$. Let $C = C'/\delta^2$ and assume that $m \geq Cp_nv(\mathcal{H}_n)$. Fix an $S \in \mathcal{S}$, let $\mathcal{I}_S = \{I \in \mathcal{I}(\mathcal{H}_n, m): g(I) = S\}$, and note for future reference that

$$|S| \leq C'p_nv(\mathcal{H}_n) \leq \delta^2m. \quad (25)$$

Since $f(S) \in \overline{\mathcal{F}}$, we have $|f(S)| < (\alpha + \delta')v(\mathcal{H}_n)$. Therefore,

$$|\mathcal{I}_S| \leq \binom{|f(S)|}{m - |S|} \leq \binom{(\alpha + \delta')v(\mathcal{H}_n)}{m - |S|}.$$

To see the above inequality, simply note that for every $I \in \mathcal{I}_S$, we have $I \setminus S \subseteq f(S)$.

It follows, using (16) and (17), that

$$|\mathcal{I}_S| \leq \binom{(\alpha + \delta')v(\mathcal{H}_n)}{m - |S|} \leq \left(\frac{\alpha + \delta'}{\alpha + \delta}\right)^{m - |S|} \left(\frac{m}{(\alpha + \delta)v(\mathcal{H}_n) - m}\right)^{|S|} \binom{(\alpha + \delta)v(\mathcal{H}_n)}{m}. \quad (26)$$

Now, if $m \geq (\alpha + \delta')v(\mathcal{H}_n)$, then every m -subset of $V(\mathcal{H}_n)$ belongs to \mathcal{F} and hence there is no independent set of size m . We may therefore assume that $m < (\alpha + \delta')v(\mathcal{H}_n) = (\alpha + \delta/2)v(\mathcal{H}_n)$. Setting $s = |S|$, we obtain

$$\binom{v(\mathcal{H}_n)}{s} \cdot |\mathcal{I}_S| \leq \left(\frac{\alpha + \delta'}{\alpha + \delta}\right)^{m/2} \left(\frac{ev(\mathcal{H}_n)}{s} \cdot \frac{2m}{\delta v(\mathcal{H}_n)}\right)^s \binom{(\alpha + \delta)v(\mathcal{H}_n)}{m} \leq e^{-\delta^2 m} \binom{(\alpha + \delta)v(\mathcal{H}_n)}{m},$$

since $s \leq \delta^2 m$, by (25), and provided that δ is sufficiently small. It follows that

$$|\mathcal{I}(\mathcal{H}_n, m)| = \sum_{S \in \mathcal{S}} |\mathcal{I}_S| \leq \sum_{s=0}^{\delta^2 m} \binom{v(\mathcal{H}_n)}{s} \max\{|\mathcal{I}_S| : |S| = s\} \leq \binom{(\alpha + \delta)v(\mathcal{H}_n)}{m},$$

as claimed. \square

6. STABILITY RESULTS FOR SPARSE SETS

In this section, we shall deduce from Theorem 2.2 two versions of the general transference theorem for stability results proved by Conlon and Gowers [14]. Similarly as in Section 5, we shall state our results using the terminology used by Schacht [58]. We remark here that in parallel to this work, Schacht's method was adapted to yield sparse random analogues of stability statements by Samotij [56]. The main result of this section is most easily compared with [56, Theorem 3.4]. We begin by recalling the following definition from [1].

Definition 6.1. Let \mathcal{H} be a sequence of k -uniform hypergraphs, let α be a positive real, and let \mathcal{B} be a sequence of sets with $\mathcal{B}_n \subseteq \mathcal{P}(V(\mathcal{H}_n))$. We say that \mathcal{H} is (α, \mathcal{B}) -stable if for every positive δ , there exist positive ε and n_0 such that the following holds. For every n with $n \geq n_0$ and every $U \subseteq V(\mathcal{H}_n)$ with $|U| \geq (\alpha - \varepsilon)v(\mathcal{H}_n)$, we have either $e(\mathcal{H}_n[U]) \geq \varepsilon e(\mathcal{H}_n)$ or $|U \setminus B| \leq \delta v(\mathcal{H}_n)$ for some $B \in \mathcal{B}_n$.

Roughly speaking, a sequence \mathcal{H} of hypergraphs is (α, \mathcal{B}) -stable if for every $A \subseteq V(\mathcal{H}_n)$ that is almost as large as $\alpha v(\mathcal{H}_n)$, the set A is either very 'close' to some extremal set $B \in \mathcal{B}_n$ or it contains 'many' (a positive fraction of all) edges of \mathcal{H}_n . Note that in many natural settings, such a property does hold, for example, as a consequence of the Erdős-Simonovits stability theorem [16, 60] and the removal lemma for graphs.

We again start with the 'random' version of our stability result, which was originally proved in [56].

Theorem 6.2. *Let \mathcal{H} be a sequence of k -uniform hypergraphs, let $\alpha \in (0, 1)$, and let c be a positive constant. Let \mathbf{p} be a sequence of probabilities such that, for every $\ell \in [k]$,*

$$\Delta_\ell(\mathcal{H}_n) \leq c \cdot p_n^{\ell-1} \frac{e(\mathcal{H}_n)}{v(\mathcal{H}_n)}$$

and let \mathcal{B} be a sequence of sets with $\mathcal{B}_n \subseteq \mathcal{P}(V(\mathcal{H}_n))$.

If \mathcal{H} is (α, \mathcal{B}) -stable, then the following holds. For every positive δ , there exist ε and C such that if $q_n \geq Cp_n$ and $q_n v(\mathcal{H}_n) \rightarrow \infty$ as $n \rightarrow \infty$, then a.a.s. every independent set $I \subseteq V(\mathcal{H}_n)_{q_n}$ with $|I| \geq (\alpha - \varepsilon)q_n v(\mathcal{H}_n)$ satisfies $|I \setminus B| < \delta q_n v(\mathcal{H}_n)$ for some $B \in \mathcal{B}_n$.

The following theorem, a ‘counting’ analogue of Theorem 6.2, is our main stability result. A simple version of it, applicable to 3-uniform hypergraphs with $\Delta_2(\mathcal{H}_n) = O(1)$, was proved in [1] and used in [1, 2] to count sum-free subsets in Abelian groups and in the set $[n]$.

Theorem 6.3. *Let \mathcal{H} be a sequence of k -uniform hypergraphs, let $\alpha \in (0, 1)$, and let c be a positive constant. Let \mathbf{p} be a sequence of probabilities such that, for every $\ell \in [k]$,*

$$\Delta_\ell(\mathcal{H}_n) \leq c \cdot p_n^{\ell-1} \frac{e(\mathcal{H}_n)}{v(\mathcal{H}_n)}$$

and let \mathcal{B} be a sequence of sets with $\mathcal{B}_n \subseteq \mathcal{P}(V(\mathcal{H}_n))$.

If \mathcal{H} is (α, \mathcal{B}) -stable, then the following holds. For every positive δ , there exist ε and C such that if $m \geq Cp_nv(\mathcal{H}_n)$, then there are at most

$$(1 - \varepsilon)^m \binom{\alpha v(\mathcal{H}_n)}{m}$$

independent sets $I \in \mathcal{I}(\mathcal{H}_n, m)$ such that $|I \setminus B| \geq \delta m$ for every $B \in \mathcal{B}_n$.

Proof of Theorem 6.2. The proof is similar to the proof of Theorem 5.2. Let $k \in \mathbb{N}$, $\alpha \in (0, 1)$, $\mathbf{p} \in [0, 1]^\mathbb{N}$, $c \in (0, \infty)$, and \mathcal{H} and \mathcal{B} be as in the statement of Theorem 6.2. Furthermore, suppose that \mathcal{H} is (α, \mathcal{B}) -stable and fix some small positive δ . Let ε be a small positive constant, let n be sufficiently large, let $\delta' = \delta/3$ and $\varepsilon' = 3\varepsilon$, and set

$$\mathcal{F} = \{A \subseteq V(\mathcal{H}_n) : |A| \geq (\alpha - \varepsilon')v(\mathcal{H}_n) \text{ and } |A \setminus B| \geq \delta'v(\mathcal{H}_n) \text{ for every } B \in \mathcal{B}_n\}.$$

Since \mathcal{H} is (α, \mathcal{B}) -stable, it follows that \mathcal{H}_n is $(\mathcal{F}, \varepsilon)$ -dense, provided that ε is sufficiently small. Let $C' = C_{2.2}(k, \varepsilon, c)$. By Theorem 2.2, there exist a family $\mathcal{S} \subseteq \binom{V(\mathcal{H}_n)}{\leq C'p_nv(\mathcal{H}_n)}$ and functions $f: \mathcal{S} \rightarrow \overline{\mathcal{F}}$ and $g: \mathcal{I}(\mathcal{H}_n) \rightarrow \mathcal{S}$ such that

$$g(I) \subseteq I \quad \text{and} \quad I \setminus g(I) \subseteq f(g(I))$$

for every $I \in \mathcal{I}(\mathcal{H}_n)$. Let $C = C'/\varepsilon^3$ and assume that $q_n \geq Cp_n$. Let $m = (\alpha - \varepsilon)q_nv(\mathcal{H}_n)$ and, for the sake of brevity, let us write $V = V(\mathcal{H}_n)$ and $q = q_n$. Let

$$\mathcal{I}' = \{I \in \mathcal{I}(\mathcal{H}_n) : |I| \geq m \text{ and } |I \setminus B| \geq \delta qv(\mathcal{H}_n) \text{ for every } B \in \mathcal{B}_n\}$$

and let \mathcal{A} denote the event that $\mathcal{H}_n[V_q]$ contains an independent set $I \in \mathcal{I}'$. We are required to prove that $\mathbb{P}(\mathcal{A})$ tends to 0 as $n \rightarrow \infty$.

Observe first that

$$\mathbb{P}(\mathcal{A}) \leq \sum_{S \in \mathcal{S}} \mathbb{P}(I \subseteq V_q \text{ for some } I \in \mathcal{I}' \text{ such that } g(I) = S). \quad (27)$$

Fix an $S \in \mathcal{S}$, let $\mathcal{I}'_S = \{I \in \mathcal{I}' : g(I) = S\}$, and note for future reference that

$$|S| \leq C'p_nv(\mathcal{H}_n) \leq \varepsilon^3 qv(\mathcal{H}_n). \quad (28)$$

We claim that

$$\mathbb{P}(I \subseteq V_q \text{ for some } I \in \mathcal{I}'_S) \leq \mathbb{P}(S \subseteq V_q) \cdot \exp\left(-\frac{\varepsilon^2 qv(\mathcal{H}_n)}{4}\right). \quad (29)$$

In order to prove (29), recall that since $f(S) \in \overline{\mathcal{F}}$, we either have $|f(S)| < (\alpha - \varepsilon')v(\mathcal{H}_n)$ or $|f(S) \setminus B| < \delta'v(\mathcal{H}_n)$ for some $B \in \mathcal{B}_n$. We therefore consider two cases.

Case 1: $|f(S)| < (\alpha - \varepsilon')v(\mathcal{H}_n)$.

We bound the left-hand side of (29) as follows:

$$\mathbb{P}\left(I \subseteq V_q \text{ for some } I \in \mathcal{I}'_S\right) \leq \mathbb{P}(S \subseteq V_q) \cdot \mathbb{P}\left(|V_q \cap f(S)| \geq m - |S|\right). \quad (30)$$

In order to justify the above inequality, note that for every $I \in \mathcal{I}'_S$, we have $I \setminus S \subseteq f(S)$. Recall that $\varepsilon' = 3\varepsilon$. Since $m - |S| \geq (\alpha - 2\varepsilon)qv(\mathcal{H}_n)$, by (28), and

$$\mathbb{E}[|V_q \cap f(S)|] \leq (\alpha - \varepsilon')qv(\mathcal{H}_n) = (\alpha - 3\varepsilon)qv(\mathcal{H}_n),$$

then by Chernoff's inequality we have

$$\mathbb{P}\left(|V_q \cap f(S)| \geq m - |S|\right) \leq \exp\left(-\frac{\varepsilon^2qv(\mathcal{H}_n)}{4}\right). \quad (31)$$

Combining (30) and (31), we obtain (29), as required.

Case 2: $|f(S) \setminus B| < \delta'v(\mathcal{H}_n)$ for some $B \in \mathcal{B}_n$.

We estimate the left-hand side of (29) as follows:

$$\mathbb{P}\left(I \subseteq V_q \text{ for some } I \in \mathcal{I}'_S\right) \leq \mathbb{P}(S \subseteq V_q) \cdot \mathbb{P}\left(|V_q \cap (f(S) \setminus B)| \geq \delta'qv(\mathcal{H}_n) - |S|\right).$$

This follows from the definition of \mathcal{I}' and the fact that $I \setminus S \subseteq f(S)$ for every $I \in \mathcal{I}'_S$. Since $|f(S) \setminus B| < \delta'v(\mathcal{H}_n)$, we have

$$\mathbb{E}[|V_q \cap (f(S) \setminus B)|] < \delta'qv(\mathcal{H}_n),$$

whereas $\delta'qv(\mathcal{H}_n) - |S| \geq 2\delta'qv(\mathcal{H}_n)$ by (28) and since $\delta = 3\delta'$. By Chernoff's inequality, it follows that

$$\mathbb{P}\left(|V_q \cap (f(S) \setminus B)| \geq 3\delta'qv(\mathcal{H}_n) - |S|\right) \leq \exp\left(-\frac{(\delta')^2qv(\mathcal{H}_n)}{4}\right) \leq \exp\left(-\frac{\varepsilon^2qv(\mathcal{H}_n)}{4}\right)$$

since ε was chosen sufficiently small. Thus (29) follows in this case as well.

Finally, note that, since $|S| \leq \varepsilon^3qv(\mathcal{H}_n)$ for every $S \in \mathcal{S}$, as in (24), we have

$$\sum_{S \in \mathcal{S}} \mathbb{P}(S \subseteq V_q) \leq \sum_{s=0}^{\varepsilon^3qv(\mathcal{H}_n)} \binom{v(\mathcal{H}_n)}{s} q^s \leq v(\mathcal{H}_n) \cdot \left(\frac{e}{\varepsilon^3}\right)^{\varepsilon^3qv(\mathcal{H}_n)}. \quad (32)$$

Putting (27), (29), and (32) together, we obtain

$$\mathbb{P}(\mathcal{A}) \leq \sum_{S \in \mathcal{S}} \mathbb{P}(S \subseteq V_q) \exp\left(-\frac{\varepsilon^2qv(\mathcal{H}_n)}{4}\right) \leq \exp(-\varepsilon^3qv(\mathcal{H}_n)),$$

as required. \square

Proof of Theorem 6.3. Let $k \in \mathbb{N}$, $\alpha \in (0, 1)$, $\mathbf{p} \in [0, 1]^{\mathbb{N}}$, $c \in (0, \infty)$, and \mathcal{H} and \mathcal{B} be as in the statement of Theorem 6.3. Furthermore, suppose that \mathcal{H} is (α, \mathcal{B}) -stable and fix some positive δ . Let δ' be a sufficiently small positive constant (depending only on α and δ), let ε be a small positive constant, and let n be sufficiently large. Let $\varepsilon' = 2\varepsilon$, and set

$$\mathcal{F} = \{A \subseteq V(\mathcal{H}_n) : |A| \geq (\alpha - \varepsilon')v(\mathcal{H}_n) \text{ and } |A \setminus B| \geq \delta'v(\mathcal{H}_n) \text{ for every } B \in \mathcal{B}_n\}.$$

Since \mathcal{H} is (α, \mathcal{B}) -stable, it follows that \mathcal{H}_n is $(\mathcal{F}, \varepsilon)$ -dense, provided that ε is sufficiently small (as a function of δ'). Let $C' = C_{2,2}(k, \varepsilon, c)$. By Theorem 2.2, there exist a family $\mathcal{S} \subseteq \binom{V(\mathcal{H}_n)}{\leq C' p_n v(\mathcal{H}_n)}$ and functions $f: \mathcal{S} \rightarrow \overline{\mathcal{F}}$ and $g: \mathcal{I}(\mathcal{H}_n) \rightarrow \mathcal{S}$ such that

$$g(I) \subseteq I \quad \text{and} \quad I \setminus g(I) \subseteq f(g(I))$$

for every $I \in \mathcal{I}(\mathcal{H}_n)$. Let $C = C'/\varepsilon^2$, assume that $m \geq C p_n v(\mathcal{H}_n)$, and set

$$\mathcal{I}' = \{I \in \mathcal{I}(\mathcal{H}_n, m) : |I \setminus B| \geq \delta m \text{ for every } B \in \mathcal{B}_n\}.$$

Our task is to bound the size of \mathcal{I}' from above. To this end, fix an $S \in \mathcal{S}$ and let $\mathcal{I}'_S = \{I \in \mathcal{I}' : g(I) = S\}$. Note for future reference that

$$|S| \leq C' p_n v(\mathcal{H}_n) \leq \varepsilon^2 m. \quad (33)$$

Since $f(S) \in \overline{\mathcal{F}}$, we either have $|f(S)| < (\alpha - \varepsilon')v(\mathcal{H}_n)$ or $|f(S) \setminus B| < \delta'v(\mathcal{H}_n)$ for some $B \in \mathcal{B}_n$. We therefore consider two cases.

Case 1: $|f(S)| < (\alpha - \varepsilon')v(\mathcal{H}_n)$.

We claim that in this case

$$\binom{v(\mathcal{H}_n)}{|S|} \cdot |\mathcal{I}'_S| \leq \frac{(1 - \varepsilon)^m}{2m} \binom{\alpha v(\mathcal{H}_n)}{m}. \quad (34)$$

To prove (34), we first estimate the size of \mathcal{I}'_S as follows:

$$|\mathcal{I}'_S| \leq \binom{|f(S)|}{m - |S|} \leq \binom{(\alpha - \varepsilon')v(\mathcal{H}_n)}{m - |S|}.$$

The above inequality follows since $I \setminus S \subseteq f(S)$ for every $I \in \mathcal{I}'_S$.

It follows, using (16) and (17), as in (26), that

$$|\mathcal{I}'_S| \leq \binom{(\alpha - \varepsilon')v(\mathcal{H}_n)}{m - |S|} \leq \left(\frac{\alpha - \varepsilon'}{\alpha}\right)^{m - |S|} \left(\frac{m}{\alpha v(\mathcal{H}_n) - m}\right)^{|S|} \binom{\alpha v(\mathcal{H}_n)}{m}.$$

Now, if $m \geq (\alpha - \varepsilon')v(\mathcal{H}_n)$, then $\mathcal{I}' \subseteq \mathcal{F}$ and hence $\mathcal{I}' = \emptyset$, since \mathcal{H}_n is $(\mathcal{F}, \varepsilon)$ -dense. We may therefore assume that $m < (\alpha - \varepsilon')v(\mathcal{H}_n) = (\alpha - 2\varepsilon)v(\mathcal{H}_n)$. We obtain

$$\binom{v(\mathcal{H}_n)}{|S|} |\mathcal{I}'_S| \leq \left(\frac{\alpha - \varepsilon'}{\alpha}\right)^{m/2} \left(\frac{ev(\mathcal{H}_n)}{|S|} \cdot \frac{m}{2\varepsilon v(\mathcal{H}_n)}\right)^{|S|} \binom{\alpha v(\mathcal{H}_n)}{m} \leq \frac{(1 - \varepsilon)^m}{2m} \binom{\alpha v(\mathcal{H}_n)}{m},$$

since $|S| \leq \varepsilon^2 m$ and $\varepsilon' = 2\varepsilon$, as claimed.

Case 2: $|f(S) \setminus B| < \delta'v(\mathcal{H}_n)$ for some $B \in \mathcal{B}_n$.

We claim that in this case

$$\binom{v(\mathcal{H}_n)}{|S|} \cdot |\mathcal{I}'_S| \leq \delta^m \binom{\alpha v(\mathcal{H}_n)}{m}. \quad (35)$$

To prove (35), we first estimate the size of \mathcal{I}'_S as follows:

$$|\mathcal{I}'_S| \leq \binom{|f(S) \setminus B|}{\delta m - |S|} \binom{|f(S)|}{m - \delta m} \leq \binom{\delta' v(\mathcal{H}_n)}{\delta m - |S|} \binom{v(\mathcal{H}_n)}{m - \delta m}. \quad (36)$$

To see the first inequality, recall that every $I \in \mathcal{I}'_S$ contains at least $\delta m - |S|$ elements of $f(S) \setminus B$ for every $B \in \mathcal{B}_n$. Recall that $|S| \leq \varepsilon^2 m$ and note that therefore, if $m \geq (\alpha/2)v(\mathcal{H}_n)$, then $\delta m - |S| \geq \delta' v(\mathcal{H}_n)$ and hence $\mathcal{I}'_S = \emptyset$. Thus, we may assume that $m < (\alpha/2)v(\mathcal{H}_n)$. It follows, using (17) and (18), that

$$\binom{v(\mathcal{H}_n)}{m - \delta m} \leq \left(\frac{m}{v(\mathcal{H}_n) - m} \right)^{\delta m} \binom{v(\mathcal{H}_n)}{m} \leq \left(\frac{2m}{v(\mathcal{H}_n)} \right)^{\delta m} \left(\frac{2}{\alpha} \right)^m \binom{\alpha v(\mathcal{H}_n)}{m}. \quad (37)$$

Hence, by (33), (36), and (37), using (15), we have

$$\begin{aligned} \binom{v(\mathcal{H}_n)}{|S|} |\mathcal{I}'_S| &\leq \left(\frac{ev(\mathcal{H}_n)}{|S|} \right)^{|S|} \left(\frac{2e\delta' v(\mathcal{H}_n)}{\delta m} \right)^{\delta m - |S|} \left(\frac{2m}{v(\mathcal{H}_n)} \right)^{\delta m} \left(\frac{2}{\alpha} \right)^m \binom{\alpha v(\mathcal{H}_n)}{m} \\ &\leq \left(\frac{1}{|S|} \cdot \frac{\delta m}{2\delta'} \right)^{|S|} \left(\frac{4e\delta'}{\delta} \right)^{\delta m} \left(\frac{2}{\alpha} \right)^m \binom{\alpha v(\mathcal{H}_n)}{m} \leq \delta^m \binom{\alpha v(\mathcal{H}_n)}{m}, \end{aligned}$$

as claimed, since $|S| \leq \varepsilon^2 m$ and δ' and ε were chosen to be sufficiently small. Indeed, note that (for this calculation, and assuming that δ is sufficiently small) $\delta' = \delta^3 \cdot (\delta\alpha/2e)^{1/\delta}$ and $\varepsilon < \delta'$ suffice.

Finally, by (34) and (35), we obtain

$$|\mathcal{I}'| = \sum_{S \in \mathcal{S}} |\mathcal{I}'_S| \leq \sum_{s=0}^{\varepsilon^2 m} \binom{v(\mathcal{H}_n)}{s} \max \{ |\mathcal{I}'_S| : |S| = s \} \leq (1 - \varepsilon)^m \binom{\alpha v(\mathcal{H}_n)}{m},$$

as claimed. □

7. TURÁN'S PROBLEM IN RANDOM GRAPHS

In this section, we shall deduce from Theorems 5.2 and 6.2 the sparse random analogues of the classical theorems of Erdős and Stone [20] and Turán [64] and of Erdős and Simonovits [16, 60], Theorems 1.3 and 1.5. In fact, we will prove a natural generalization of Theorem 1.3 to t -uniform hypergraphs, Theorem 7.2 below, which was already proved by Conlon and Gowers [14] and Schacht [58]. We first recall the following generalization of the notion of 2-density of a graph to t -uniform hypergraphs.

Definition 7.1. Let H be a t -uniform hypergraph with at least $t + 1$ vertices. We define the t -density of H , denoted by $m_t(H)$, by

$$m_t(H) = \max \left\{ \frac{e(H') - 1}{v(H') - t} : H' \subseteq H \text{ with } v(H') \geq t + 1 \right\}.$$

We also recall that the *Turán density* of a t -uniform hypergraph H , denoted $\pi(H)$, is defined by

$$\pi(H) = \lim_{n \rightarrow \infty} \frac{\text{ex}(K_n^t, H)}{\binom{n}{t}}, \quad (38)$$

where, as usual, $\text{ex}(K_n^t, H)$ is the Turán number for H , that is, the maximum number of edges in an H -free t -uniform hypergraph with n vertices.

Theorem 7.2. *For every t -uniform hypergraph H with $\Delta(H) \geq 2$ and every positive δ , there exists a positive constant C such that if $q_n \geq Cn^{-1/m_t(H)}$, then*

$$\mathbb{P} \left(\text{ex}(G^t(n, q_n), H) \leq (\pi(H) + \delta)q_n \binom{n}{t} \right) \rightarrow 1$$

as $n \rightarrow \infty$.

Once again, we emphasize that we actually obtain essentially optimal bounds on the probability in the above statement, i.e., bounds of the form $1 - \exp(-bq_n n^t)$ for some positive constant b that depends only on H and δ .

Theorems 7.2 and 1.5, and hence also Theorem 1.3, will follow easily from our general transference results, Theorems 5.2 and 6.2, the classical supersaturation results of Erdős and Simonovits [19] (for Theorem 7.2), and the stability theorem of Erdős and Simonovits [16, 60] together with the so-called graph removal lemma (for Theorem 1.5). We only need to check that the hypergraph of copies of H in the complete hypergraph K_n^t , to which we would like to apply our transference theorems, satisfies the assumptions of Theorems 5.2 and 6.2. Since we are going to use this fact several times in this and later sections, we state it as a separate proposition.

Let H be an arbitrary t -uniform hypergraph. The *hypergraph of copies of H in K_n^t* is the $e(H)$ -uniform hypergraph on the vertex set $E(K_n^t)$ whose edges are the edge sets of all copies of H in K_n^t .

Proposition 7.3. *Let n and t be integers with $t \geq 2$ and let H be a t -uniform hypergraph. Set $k = e(H)$ and let \mathcal{H} be the k -uniform hypergraph of copies of H in K_n^t . There exists a positive constant c such that, letting $p = n^{-1/m_t(H)}$,*

$$\Delta_\ell(\mathcal{H}) \leq c \cdot p^{\ell-1} \frac{e(\mathcal{H})}{v(\mathcal{H})} \quad (39)$$

for every $\ell \in [k]$.

Proof. Note that $v(\mathcal{H}) = \binom{n}{t} = \Theta(n^t)$ and that $e(\mathcal{H}) = \frac{(v(H))!}{|\text{Aut}(H)|} \cdot \binom{n}{v(H)} = \Theta(n^{v(H)})$. By the definition of p and $m_t(H)$, we have

$$p^{e(H')-1} n^{v(H')-t} \geq 1 \quad (40)$$

for every $H' \subseteq H$. Now, for each $\ell \in [k]$,

$$\Delta_\ell(\mathcal{H}) \leq c' \cdot \max \left\{ n^{v(H)-v(H')} : H' \subseteq H \text{ with } e(H') = \ell \right\}$$

for some positive constant c' . Since $e(\mathcal{H})/v(\mathcal{H}) \geq c'' \cdot n^{v(H)-t}$ for some constant c'' , it follows that

$$\begin{aligned} \Delta_\ell(\mathcal{H}) \cdot \left(p^{\ell-1} \frac{e(\mathcal{H})}{v(\mathcal{H})} \right)^{-1} &\leq c' \cdot \frac{v(\mathcal{H})}{e(\mathcal{H})} \cdot \max_{H' \subseteq H: e(H')=\ell} \left(\frac{n^{v(H)}}{p^{e(H')-1} n^{v(H')}} \right) \\ &\leq \frac{c'}{c''} \cdot \max_{H' \subseteq H: e(H')=\ell} \left(\frac{1}{p^{e(H')-1}} n^{v(H')-t} \right) \leq \frac{c'}{c''}, \end{aligned}$$

where the last inequality follows by (40). \square

Proof of Theorem 7.2. Let H be a t -uniform hypergraph, let $k = e(H)$, and let $(\mathcal{H}_n)_{n \in \mathbb{N}}$ be the sequence of k -uniform hypergraphs of copies of H in K_n^t . Let $\alpha = \pi(H)$, let δ be a positive constant, and let $p_n = n^{-1/m_t(H)}$. It follows easily from the supersaturation theorem of Erdős and Simonovits [19] that \mathcal{H} is α -dense, see [58]. Let $C = C_{5,2}(\mathcal{H}, \delta)$ and assume that $q_n \geq Cp_n = Cn^{-1/m_t(H)}$. Note that the assumption that H contains a vertex of degree at least 2 implies that $m_t(H) > 1/t$ and hence $q_n v(\mathcal{H}_n) \rightarrow \infty$ as $n \rightarrow \infty$. Together with Proposition 7.3, this implies that \mathcal{H} satisfies the assumptions of Theorem 5.2 and hence with probability tending to 1 as $n \rightarrow \infty$,

$$\text{ex}(G^t(n, q_n), H) = \alpha (\mathcal{H}_n[E(G^t(n, q_n))]) \leq (\pi(H) + \delta) q_n \binom{n}{t},$$

as required. \square

In the proof of Theorems 1.5 and 1.7, we shall need the following proposition, which is a fairly straightforward consequence of the Erdős-Simonovits stability theorem [16, 60] and the graph removal lemma [17]. A proof of this statement can be found in [56]. We remark that a new proof of the graph removal lemma, which avoids the use of the Szemerédi regularity lemma, was given recently by Fox [21].

Proposition 7.4. *For every graph H and every positive δ , there exists a positive ε such that the following holds for every $n \in \mathbb{N}$. If G is an n -vertex graph with*

$$e(G) \geq \left(1 - \frac{1}{\chi(H) - 1} - \varepsilon \right) \binom{n}{2},$$

then either G may be made $(\chi(H) - 1)$ -partite by removing from it at most δn^2 edges or G contains at least $\varepsilon n^{v(H)}$ copies of H .

Proof of Theorem 1.5. Let H be a graph, let $k = e(H)$, and let $(\mathcal{H}_n)_{n \in \mathbb{N}}$ be the sequence of k -uniform hypergraphs of copies of H in K_n . Let $\alpha = \pi(H) = \left(1 - \frac{1}{\chi(H)-1} \right)$, let δ be a positive constant, and let $p_n = n^{-1/m_2(H)}$. Moreover, let \mathcal{B}_n be the family of all complete $(\chi(H) - 1)$ -partite subgraphs of K_n . By Proposition 7.4, \mathcal{H} is (α, \mathcal{B}) -stable. Let $C = C_{6,2}(\mathcal{H}, \delta)$, let $\varepsilon = \varepsilon_{6,2}(\mathcal{H}, \delta)$, and assume that $q_n \geq Cp_n = Cn^{-1/m_2(H)}$. Note that the assumption that H contains a vertex of degree at least 2 implies that $m_2(H) > 1/2$ and hence $q_n v(\mathcal{H}_n) \rightarrow \infty$ as $n \rightarrow \infty$. Together with Proposition 7.3, the discussion above implies that \mathcal{H} satisfies the assumptions of Theorem 6.2 and hence with probability tending to 1 as $n \rightarrow \infty$, every independent set $G' \subseteq G(n, q_n)$ with $|G'| \geq (\alpha - \varepsilon) q_n v(\mathcal{H}_n)$ satisfies $|G' \setminus B| \leq \delta q_n v(\mathcal{H}_n)$

for some $B \in \mathcal{B}_n$. In other words, with probability tending to 1 as $n \rightarrow \infty$, every H -free subgraph of $G(n, q_n)$ with at least $\left(1 - \frac{1}{\chi(H)-1} - \varepsilon\right) \binom{n}{2} q_n$ edges can be made $(\chi(H) - 1)$ -partite by removing from it at most $\delta q_n \binom{n}{2}$ edges, as required. \square

8. THE TYPICAL STRUCTURE OF H -FREE GRAPHS

In this section, we shall deduce from Theorems 5.4 and 6.3 the sparse analogue of the theorem of Erdős, Frankl, and Rödl [17], Theorem 1.6, and an approximate sparse analogue of the result of Erdős, Kleitman, and Rothschild [18], Theorem 1.7. We stress once again that neither proof employs Szemerédi's regularity lemma. In order to prove Theorem 1.6, we are actually going to prove the following natural generalization of it to t -uniform hypergraphs. Generalizing the definition stated in Section 1.3, given integers n and m with $0 \leq m \leq \binom{n}{t}$ and a t -uniform hypergraph H , let us denote by $f_{n,m}(H)$ the number of H -free t -uniform hypergraphs on the vertex set $[n]$ that have exactly m edges.

Theorem 8.1. *For every t -uniform hypergraph H and every positive δ , there exists a positive constant C such that the following holds. For every $n \in \mathbb{N}$, if $m \geq Cn^{t-1/m_t(H)}$, then*

$$\binom{\text{ex}(n, H)}{m} \leq f_{n,m}(H) \leq \binom{\text{ex}(n, H) + \delta n^t}{m}.$$

We remark that Theorem 8.1 refines a result of Nagle, Rödl, and Schacht [48], who, using the hypergraph regularity lemma, generalized (4) to t -uniform hypergraphs.

Proof of Theorem 8.1. Let H be a t -uniform hypergraph, let $k = e(H)$, and let $(\mathcal{H}_n)_{n \in \mathbb{N}}$ be the sequence of k -uniform hypergraphs of copies of H in K_n^t . Let $\alpha = \pi(H)$, see (38), let δ be a positive constant, and let $p_n = n^{-1/m_t(H)}$. It follows easily from the supersaturation theorem of Erdős and Simonovits [19] that \mathcal{H} is α -dense, see [58]. Let $C = C_{5.4}(\mathcal{H}, \delta)$ and assume that $m \geq Cn^{t-1/m_t(H)} \geq Cp_n v(\mathcal{H}_n)$. Note that Proposition 7.3 implies that \mathcal{H} satisfies the assumptions of Theorem 5.4 and hence

$$f_{n,m}(H) = |I(\mathcal{H}_n, m)| \leq \binom{(\pi(H) + \delta) \binom{n}{t}}{m} \leq \binom{\text{ex}(n, H) + \delta n^t}{m},$$

as required. The claimed lower bound on $f_{n,m}(H)$ is trivial. \square

Proof of Theorem 1.7. Let H be a graph, let $k = e(H)$, and let $(\mathcal{H}_n)_{n \in \mathbb{N}}$ be the sequence of k -uniform hypergraphs of copies of H in K_n . Let $\alpha = \pi(H) = 1 - \frac{1}{\chi(H)-1}$, let δ be a positive constant, and let $p_n = n^{-1/m_2(H)}$. Moreover, let \mathcal{B}_n be the family of all complete $(\chi(H) - 1)$ -partite subgraphs of K_n . By Proposition 7.4, \mathcal{H} is (α, \mathcal{B}) -stable. Let $C = C_{6.3}(\mathcal{H}, \delta)$, let $\varepsilon = \varepsilon_{6.3}(\mathcal{H}, \delta)$, and assume that $m \geq Cn^{2-1/m_2(H)} \geq Cp_n v(\mathcal{H}_n)$. Together with Proposition 7.3, this implies that \mathcal{H} satisfies the assumptions of Theorem 6.3 and hence, letting $f_{n,m}^\delta(H)$ denote the number of H -free graphs on the vertex set $[n]$ that have exactly m edges and that are not $(\delta, \chi(H) - 1)$ -partite,

$$f_{n,m}^\delta(H) \leq (1 - \varepsilon)^m \binom{\pi(H) \binom{n}{2}}{m}.$$

Finally, note that (trivially),

$$f_{n,m}(H) \geq \binom{\pi(H) \binom{n}{2}}{m}$$

and hence $f_{n,m}^\delta(H) = o(f_{n,m}(H))$, as claimed. \square

For t -uniform hypergraphs, there is no general stability theorem known; however, such results have been proved for a few specific hypergraphs (see [25, 26, 36, 37]), and in each case we obtain a corresponding result for sparse hypergraphs. For example, following [9], let F_5 denote the ‘3-uniform triangle’, i.e., the hypergraph with edge set isomorphic to $\{123, 124, 345\}$, and say that a 3-uniform hypergraph is *triangle-free* if it contains no copy of F_5 . The following theorem follows easily, as above, from Theorem 6.3 combined with the hypergraph removal lemma of Gowers [33] and Rödl and Skokan [54] and the stability theorem for 3-uniform triangle-free hypergraphs, which was proved by Keevash and Mubayi [36].

Theorem 8.2. *For every positive δ , there exists a constant C such that the following holds. If $m \geq Cn^2$, then almost every triangle-free 3-uniform hypergraph with n vertices and m edges can be made tripartite by removing from it at most δm edges.*

Proof. Let $(\mathcal{H}_n)_{n \in \mathbb{N}}$ be the sequence of 3-uniform hypergraphs of copies of F_5 in K_n^3 , set $\alpha = 2/9$, and let \mathcal{B}_n denote the collection of all complete tripartite subhypergraphs of K_n^3 . By the hypergraph removal lemma [54, Theorem 1.3], combined with the stability theorem for triangle-free 3-uniform hypergraphs [36, Theorem 1.6], it follows that \mathcal{H} is (α, \mathcal{B}) -dense.

It follows by Proposition 7.3 that \mathcal{H} satisfies the conditions of Theorem 6.3 with $p_n = n^{-1}$. Hence the number of triangle-free 3-uniform hypergraphs with n vertices and m edges that cannot be made tripartite by removing at most δm edges is at most

$$(1 - \varepsilon)^m \binom{\pi(F_5) \binom{n}{3}}{m},$$

which easily implies the theorem. \square

Finally, we remark that Theorem 8.2 can be seen as an approximate sparse analogue of a result of Balogh and Mubayi [9], who used the hypergraph regularity lemma and [36, Theorem 1.6] to show that almost all triangle-free 3-uniform hypergraphs are tripartite. For similar results for other forbidden hypergraphs, see [8] and [50].

9. THE KLR CONJECTURE

In this section, we shall deduce from Theorem 2.2 the KLR conjecture, Theorem 1.9. As in the preceding sections, the proof will be a fairly straightforward application of Theorem 2.2 to an appropriately defined hypergraph \mathcal{H} and family $\mathcal{F} \subseteq \mathcal{P}(V(\mathcal{H}))$. Let H be an arbitrary graph and let \mathcal{H} be the $e(H)$ -uniform hypergraph of canonical copies of H in the complete blow-up of H . Defining an appropriate family \mathcal{F} and showing that \mathcal{H} is $(\mathcal{F}, \varepsilon)$ -dense will require some work.

Given a graph H and integers $n_1, \dots, n_{v(H)}$, let us denote by $\mathcal{G}(H; n_1, \dots, n_{v(H)})$ the collection of all graphs G constructed in the following way. The vertex set of G is a disjoint union $V_1 \cup \dots \cup V_{v(H)}$ of sets of sizes $n_1, \dots, n_{v(H)}$, respectively, one for each vertex of H . The only edges of G lie between those pairs of sets (V_i, V_j) such that $\{i, j\}$ is an edge of H . Recall the definition of $\mathcal{G}(H, n, m, p, \varepsilon)$ from Section 1.4 and observe that $\mathcal{G}(H, n, m, p, \varepsilon) \subseteq \mathcal{G}(H; n, \dots, n)$ for all m, p , and ε .

The following lemma, which is a robust version of the embedding lemma, stated in Section 1.4, suggests the right choice of \mathcal{F} . The lemma is well-known, and so we omit the (standard) proof.

Lemma 9.1. *Let H be a graph and let $\delta: (0, 1] \rightarrow (0, 1)$ be an arbitrary function. There exist positive constants α_0, ξ , and N such that for every collection of integers $n_1, \dots, n_{v(H)}$ satisfying $n_1, \dots, n_{v(H)} \geq N$ and every graph $G \in \mathcal{G}(H; n_1, \dots, n_{v(H)})$, one of the following holds:*

- (a) G contains at least $\xi n_1 \dots n_{v(H)}$ canonical copies of H .
- (b) There exist a positive constant α with $\alpha \geq \alpha_0$, an edge $\{i, j\} \in E(H)$, and sets $A_i \subseteq V_i, A_j \subseteq V_j$ such that $|A_i| \geq \alpha n_i, |A_j| \geq \alpha n_j$, and $d_G(A_i, A_j) < \delta(\alpha)$.

Our next lemma is also straightforward. It allows us to count (ε, p) -regular subgraphs of a graph that has a ‘hole’, as in Lemma 9.1(b). Recall that $\mathcal{G}(K_2, n, m, p, \varepsilon)$ denotes the collection of all (ε, p) -regular bipartite graphs with m edges and n vertices in each part. Given such G , let $V_1(G)$ and $V_2(G)$ denote the two parts. For each $\beta \in (0, 1)$, define a function $\delta: (0, 1] \rightarrow (0, 1)$ by setting

$$\delta(x) = \frac{1}{4e} \left(\frac{\beta}{2} \right)^{2/x^2} \tag{41}$$

for each $x \in (0, 1]$. The following lemma says that a graph \tilde{G} that has a hole of size αn and density at most $\delta(\alpha)$ has very few subgraphs in $\mathcal{G}(K_2, n, m, m/n^2, \varepsilon)$.

Lemma 9.2. *For every positive α_0 and β , there exists a positive constant ε such that the following holds. Let $\tilde{G} \subseteq K_{n,n}$ be such that there exist subsets $A \subseteq V_1(\tilde{G})$ and $B \subseteq V_2(\tilde{G})$ with*

$$\min\{|A|, |B|\} \geq \alpha n \quad \text{and} \quad d_G(A, B) < \delta(\alpha)$$

for some $\alpha \in [\alpha_0, 1]$, and let $S \subseteq \tilde{G}$. Then, for every m with $|S|/\varepsilon \leq m \leq n^2$, there are at most

$$\beta^m \binom{n^2}{m - |S|}$$

subgraphs of \tilde{G} that belong to $\mathcal{G}(K_2, n, m, m/n^2, \varepsilon)$ and contain S .

Proof. We begin by noting that, by choosing random subsets of A and B if necessary, we may assume that $|A| = |B| = \alpha n$. Set $\varepsilon = \min\{\alpha_0^2/4, 1/4\}$, write \mathcal{G}^* for the family of all subgraphs of \tilde{G} that belong to $\mathcal{G}(K_2, n, m, m/n^2, \varepsilon)$ and contain S , and let $G \in \mathcal{G}^*$. In particular, G is (ε, p) -regular, where $p = m/n^2$ and since $\varepsilon \leq \alpha$, it follows that the pair (A, B) must have density at least $(1 - \varepsilon)p$ in G , and hence must contain at least $(1 - \varepsilon - \varepsilon/\alpha^2)p|A||B|$ edges of $E(G) \setminus S$, since $|S| \leq \varepsilon m$. Set $m' = m - |S|$ and $\varepsilon' = \varepsilon(1 + 1/\alpha^2) \leq 1/2$, and let

we write $e_{\tilde{G}}(A, B)$ for the number of edges of \tilde{G} that lie between the sets A and B . Since $d_{\tilde{G}}(A, B) < \delta(\alpha)$, then the number of choices for G can be estimated as follows:

$$|\mathcal{G}^*| \leq \sum_{\ell \geq (1-\varepsilon')p|A||B|} \binom{e_{\tilde{G}}(A, B)}{\ell} \binom{e(\tilde{G}) - e_{\tilde{G}}(A, B)}{m' - \ell} \leq \sum_{\ell \geq \alpha^2 m/2} \binom{\delta(\alpha)\alpha^2 n^2}{\ell} \binom{n^2}{m' - \ell}. \quad (42)$$

Note that the right-hand side of (42) is zero if $m > 2\delta(\alpha)n^2$, so we may assume that $m' \leq m \leq 2\delta(\alpha)n^2 \leq n^2/2$. Thus, using (15) and (17), (42) implies that

$$|\mathcal{G}^*| \leq \sum_{\ell \geq \alpha^2 m/2} \left(\frac{e\delta(\alpha)\alpha^2 n^2}{\ell} \right)^\ell \left(\frac{m'}{n^2 - m'} \right)^\ell \binom{n^2}{m'} \leq \sum_{\ell \geq \alpha^2 m/2} \left(\frac{2e\delta(\alpha)\alpha^2 m}{\ell} \right)^\ell \binom{n^2}{m'}. \quad (43)$$

Since $\delta(\alpha) < 1/4e$, the summand in the right-hand side of (43) is decreasing in ℓ on $(\alpha^2 m/2, \infty)$ and hence

$$|\mathcal{G}^*| \leq m(4e\delta(\alpha))^{\alpha^2 m/2} \binom{n^2}{m'} \leq \beta^m \binom{n^2}{m'},$$

as required, since $(4e\delta(\alpha))^{\alpha^2/2} = \beta/2$. \square

We can now easily deduce Theorem 1.9 from Theorem 2.2.

Proof of Theorem 1.9. Let H be a fixed graph, let $n \in \mathbb{N}$, and let $H(n)$ be the largest graph in the family $\mathcal{G}(H; n, \dots, n)$, i.e., the complete blow-up of H , where each vertex of H is replaced by an independent set of size n and each edge of H is replaced by the complete bipartite graph $K_{n,n}$. Let \mathcal{H} be the $e(H)$ -uniform hypergraph on the vertex set $E(H(n))$ whose edges are all $n^{v(H)}$ canonical copies of H in $H(n)$.

Fix an arbitrary positive constant β , let $\delta: (0, 1] \rightarrow (0, 1)$ be the function defined in (41) with β replaced by $\beta/2$, i.e., set

$$\delta(x) = \frac{1}{4e} \left(\frac{\beta}{4} \right)^{2/x^2}$$

for each $x \in (0, 1]$, and let $\alpha_0 = (\alpha_0)_{9.1}(H, \delta)$, $\xi = \xi_{9.1}(H, \delta)$, and $N = N_{9.1}(H, \delta)$. Let \mathcal{F} be the family of all subgraphs of $H(n)$, i.e., graphs in $\mathcal{G}(H; n, \dots, n)$, for which (b) in Lemma 9.1 is *not* satisfied. Clearly \mathcal{F} is an upset, and so, by Lemma 9.1, \mathcal{H} is (\mathcal{F}, ξ) -dense provided that $n \geq N$.

Now, since \mathcal{H} is contained in the hypergraph of all copies of H in the complete graph on $v(H)n$ vertices and contains a positive proportion of those copies, it follows from Proposition 7.3 that \mathcal{H} satisfies the assumptions of Theorem 2.2 with $p = n^{2-1/m_2(H)}$ and $\varepsilon = \xi$, for some constant c depending only on H . Therefore, there is a constant C' , a family $\mathcal{S} \subseteq \binom{E(H(n))}{\leq C'n^{2-1/m_2(H)}}$, and functions $f: \mathcal{S} \rightarrow \overline{\mathcal{F}}$ and $g: \mathcal{I}(\mathcal{H}) \rightarrow \mathcal{S}$ such that

$$g(I) \subseteq I \quad \text{and} \quad I \setminus g(I) \subseteq f(g(I))$$

for every $I \in \mathcal{I}(\mathcal{H})$.

Let ε be a sufficiently small positive constant such that, in particular, $\varepsilon \leq \varepsilon_{9.2}(\alpha_0, \beta/2)$, let $C = C'/\varepsilon$, and suppose that $m \geq Cn^{2-1/m_2(H)}$. Let $\mathcal{G}^* = \mathcal{G}^*(H, n, m, m/n^2, \varepsilon)$ and note that $\mathcal{G}^* \subseteq \mathcal{I}(\mathcal{H})$. We are required to bound from above the number of graphs in \mathcal{G}^* .

To this end, fix an $S \in \mathcal{S}$, let

$$\mathcal{G}_S^* = \{G \in \mathcal{G}^* : g(G) = S\},$$

and let $G_S = f(S)$. For each $\{i, j\} \in E(H)$, let $s(i, j) = e_S(V_i, V_j)$ and note that $\sum_{ij \in E(H)} s(i, j) = |S|$. Since

$$|S| \leq C'n^{2-1/m_2(H)} \leq \varepsilon \cdot Cn^{2-1/m_2(H)} \leq \varepsilon m,$$

then $s(i, j) \leq \varepsilon m$ for every $\{i, j\} \in E(H)$.

Now, since $G_S \in \overline{\mathcal{F}}$, it follows that there exist an $\alpha \in [\alpha_0, 1]$, an edge $\{i, j\} \in E(H)$, and sets $A_i \subseteq V_i$, $A_j \subseteq V_j$ such that $|A_i|, |A_j| \geq \alpha n$ and $d_{G_S}(A_i, A_j) < \delta(\alpha)$. By Lemma 9.2, it follows that there are at most

$$\left(\frac{\beta}{2}\right)^m \binom{n^2}{m - s(i, j)}$$

choices for the edges between V_i and V_j such that $G[V_i, V_j] \in \mathcal{G}(K_2, n, m, m/n^2, \varepsilon)$ and $S[V_i, V_j] \subseteq G[V_i, V_j] \subseteq S \cup G_S[V_i, V_j]$. It follows immediately that

$$|\mathcal{G}_S^*| \leq \left(\frac{\beta}{2}\right)^m \prod_{ij \in E(H)} \binom{n^2}{m - s(i, j)}.$$

Summing over sets $S \in \mathcal{S}$, and using (15) and (17), we obtain

$$\begin{aligned} |\mathcal{G}^*| &\leq \sum_{S \in \mathcal{S}} \left(\frac{\beta}{2}\right)^m \prod_{ij \in E(H)} \left(\frac{m}{n^2 - m}\right)^{s(i, j)} \binom{n^2}{m} = \left(\frac{\beta}{2}\right)^m \binom{n^2}{m}^{e(H)} \sum_{S \in \mathcal{S}} \left(\frac{m}{n^2 - m}\right)^{|S|} \\ &\leq \left(\frac{\beta}{2}\right)^m \binom{n^2}{m}^{e(H)} \sum_{s \leq \varepsilon m} \binom{e(H)n^2}{s} \left(\frac{2m}{n^2}\right)^s \leq \left(\frac{\beta}{2}\right)^m \binom{n^2}{m}^{e(H)} \sum_{s \leq \varepsilon m} \left(\frac{2e \cdot e(H)m}{s}\right)^s. \end{aligned}$$

Now, since ε was chosen to be sufficiently small, it follows that the summand above is increasing in s on $(0, \varepsilon m]$ and hence

$$|\mathcal{G}^*| \leq \left(\frac{\beta}{2}\right)^m \binom{n^2}{m}^{e(H)} m \left(\frac{2e \cdot e(H)}{\varepsilon}\right)^{\varepsilon m} \leq \beta^m \binom{n^2}{m}^{e(H)},$$

as required. \square

Acknowledgement. The third author would like to thank Noga Alon and David Conlon for stimulating discussions. The authors would also like to thank David Conlon and Yoshiharu Kohayakawa for helpful comments on the manuscript, and David Saxton for pointing out the usefulness of allowing multiple edges. Finally, we would like to thank the anonymous referee for a very careful reading of the proof, and a plenitude of helpful suggestions.

REFERENCES

1. N. Alon, J. Balogh, R. Morris, and W. Samotij, *Counting sum-free sets in Abelian groups*, to appear in Israel J. Math.
2. ———, *A refinement of the Cameron-Erdős Conjecture*, Proc. London Math. Soc. **108** (2014), 44–72.
3. N. Alon and J.H. Spencer, *The probabilistic method*, third ed., Wiley-Interscience Series in Discrete Math. and Optimization, John Wiley & Sons Inc., Hoboken, NJ, 2008, With an appendix on the life and work of Paul Erdős.
4. L. Babai, M. Simonovits, and J. Spencer, *Extremal subgraphs of random graphs*, J. Graph Theory **14** (1990), 599–622.
5. J. Balogh, B. Bollobás, and M. Simonovits, *The number of graphs without forbidden subgraphs*, J. Combin. Theory Ser. B **91** (2004), 1–24.
6. ———, *The typical structure of graphs without given excluded subgraphs*, Random Structures Algorithms **34** (2009), 305–318.
7. J. Balogh, R. Morris, W. Samotij, and L. Warnke, *The typical structure of sparse K_{r+1} -free graphs*, submitted.
8. J. Balogh and D. Mubayi, *Almost all triple systems with independent neighborhoods are semi-bipartite*, J. Combin. Theory Ser. A **118** (2011), 1494–1518.
9. ———, *Almost all triangle-free triple systems are tripartite*, Combinatorica **32** (2012), 143–169.
10. J. Balogh and W. Samotij, *The number of $K_{m,m}$ -free graphs*, Combinatorica **31** (2011), 131–150.
11. ———, *The number of $K_{s,t}$ -free graphs*, J. Lond. Math. Soc. **83** (2011), 368–388.
12. M. Behrisch, *Random graphs without a short cycle*, Master’s thesis, Humboldt-Universität zu Berlin, 2002.
13. V. Bergelson and A. Leibman, *Polynomial extensions of van der Waerden’s and Szemerédi’s theorems*, J. Amer. Math. Soc. **9** (1996), 725–753.
14. D. Conlon and W.T. Gowers, *Combinatorial theorems in sparse random sets*, submitted.
15. D. Conlon, W.T. Gowers, W. Samotij, and M. Schacht, *On the KLR conjecture in random graphs*, to appear in Israel J. Math.
16. P. Erdős, *Some recent results on extremal problems in graph theory. Results*, Theory of Graphs (Internat. Sympos., Rome, 1966), Gordon and Breach, New York, 1967, pp. 117–123 (English); pp. 124–130 (French).
17. P. Erdős, P. Frankl, and V. Rödl, *The asymptotic number of graphs not containing a fixed subgraph and a problem for hypergraphs having no exponent*, Graphs Combin. **2** (1986), 113–121.
18. P. Erdős, D. J. Kleitman, and B. L. Rothschild, *Asymptotic enumeration of K_n -free graphs*, Colloquio Internazionale sulle Teorie Combinatorie (Rome, 1973), Tomo II, Accad. Naz. Lincei, Rome, 1976, pp. 19–27. Atti dei Convegni Lincei, No. 17.
19. P. Erdős and M. Simonovits, *Supersaturated graphs and hypergraphs*, Combinatorica **3** (1983), 181–192.
20. P. Erdős and A.H. Stone, *On the structure of linear graphs*, Bull. Amer. Math. Soc. **52** (1946), 1087–1091.
21. J. Fox, *A new proof of the graph removal lemma*, Ann. Math. **174** (2011), 561–579.
22. P. Frankl and V. Rödl, *Large triangle-free subgraphs in graphs without K_4* , Graphs Combin. **2** (1986), 135–144.
23. E. Friedgut, V. Rödl, and M. Schacht, *Ramsey properties of random discrete structures*, Random Structures Algorithms **37** (2010), 407–436.
24. Z. Füredi, *Random Ramsey graphs for the four-cycle*, Discrete Math. **126** (1994), 407–410.
25. Z. Füredi, O. Pikhurko, and M. Simonovits, *On triple systems with independent neighbourhoods*, Combin. Probab. Comput. **14** (2005), 795–813.
26. Z. Füredi and M. Simonovits, *Triple systems not containing a Fano configuration*, Combin. Probab. Comput. **14** (2005), 467–484.
27. H. Furstenberg and Y. Katznelson, *An ergodic Szemerédi theorem for commuting transformations*, J. Analyse Math. **34** (1979), 275–291.

28. S. Gerke, *Random graphs with constraints*, 2005, Habilitationsschrift, Institut für Informatik, Technische Universität München.
29. S. Gerke, Y. Kohayakawa, V. Rödl, and A. Steger, *Small subsets inherit sparse ϵ -regularity*, J. Combin. Theory Ser. B **97** (2007), 34–56.
30. S. Gerke, H.J. Prömel, T. Schickinger, A. Steger, and A. Taraz, *K_4 -free subgraphs of random graphs revisited*, Combinatorica **27** (2007), 329–365.
31. S. Gerke, T. Schickinger, and A. Steger, *K_5 -free subgraphs of random graphs*, Random Structures Algorithms **24** (2004), 194–232.
32. S. Gerke and A. Steger, *The sparse regularity lemma and its applications*, Surveys in combinatorics 2005, London Math. Soc. Lecture Note Ser., vol. 327, Cambridge Univ. Press, Cambridge, 2005, pp. 227–258.
33. W.T. Gowers, *Hypergraph regularity and the multidimensional Szemerédi theorem*, Ann. Math. **166** (2007), 897–946.
34. P.E. Haxell, Y. Kohayakawa, and T. Łuczak, *Turán’s extremal problem in random graphs: forbidding even cycles*, J. Combin. Theory Ser. B **64** (1995), 273–287.
35. ———, *Turán’s extremal problem in random graphs: forbidding odd cycles*, Combinatorica **16** (1996), 107–122.
36. P. Keevash and D. Mubayi, *Stability theorems for cancellative hypergraphs*, J. Combin. Theory Ser. B **92** (2004), 163–175.
37. P. Keevash and B. Sudakov, *The Turán number of the Fano plane*, Combinatorica **25** (2005), 561–574.
38. D.J. Kleitman and K.J. Winston, *On the number of graphs without 4-cycles*, Discrete Math. **41** (1982), 167–172.
39. Y. Kohayakawa, *Szemerédi’s regularity lemma for sparse graphs*, Foundations of computational mathematics (Rio de Janeiro, 1997), Springer, Berlin, 1997, pp. 216–230.
40. Y. Kohayakawa and B. Kreuter, *Threshold functions for asymmetric Ramsey properties involving cycles*, Random Structures Algorithms **11** (1997), 245–276.
41. Y. Kohayakawa, T. Łuczak, and V. Rödl, *Arithmetic progressions of length three in subsets of a random set*, Acta Arith. **75** (1996), 133–163.
42. ———, *On K^4 -free subgraphs of random graphs*, Combinatorica **17** (1997), 173–213.
43. Y. Kohayakawa and V. Rödl, *Regular pairs in sparse random graphs. I*, Random Structures Algorithms **22** (2003), 359–434.
44. Y. Kohayakawa, V. Rödl, and M. Schacht, *The Turán theorem for random graphs*, Combin. Probab. Comput. **13** (2004), 61–91.
45. Y. Kohayakawa, M. Schacht, and R. Spöhel, *Upper bounds on probability thresholds for asymmetric Ramsey properties*, Random Structures Algorithms **44** (2014), 1–28.
46. T. Łuczak, *On triangle-free random graphs*, Random Structures Algorithms **16** (2000), 260–276.
47. M. Marciniszyn, J. Skokan, R. Spöhel, and A. Steger, *Asymmetric Ramsey properties of random graphs involving cliques*, Random Structures Algorithms **34** (2009), 419–453.
48. B. Nagle, V. Rödl, and M. Schacht, *Extremal hypergraph problems and the regularity method*, Topics in discrete mathematics, Algorithms Combin., vol. 26, Springer, Berlin, 2006, pp. 247–278.
49. D. Osthus, H.J. Prömel, and A. Taraz, *For which densities are random triangle-free graphs almost surely bipartite?*, Combinatorica **23** (2003), 105–150, Paul Erdős and his mathematics (Budapest, 1999).
50. Y. Person and M. Schacht, *Almost all hypergraphs without Fano planes are bipartite*, Proceedings of the Twentieth Annual ACM-SIAM Symposium on Discrete Algorithms (Philadelphia, PA), SIAM, 2009, pp. 217–226.
51. F.P. Ramsey, *On a problem of formal logic*, Proc. London Math. Soc. **30** (1930), 264–286.
52. V. Rödl and A. Ruciński, *Lower bounds on probability thresholds for Ramsey properties*, Combinatorics, Paul Erdős is eighty, Vol. 1, Bolyai Soc. Math. Stud., János Bolyai Math. Soc., Budapest, 1993, pp. 317–346.
53. ———, *Threshold functions for Ramsey properties*, J. Amer. Math. Soc. **8** (1995), 917–942.

54. V. Rödl and J. Skokan, *Applications of the regularity lemma for uniform hypergraphs*, Random Structures Algorithms **28** (2006), 180–194.
55. I.Z. Ruzsa and E. Szemerédi, *Triple systems with no six points carrying three triangles*, Combinatorics (Proc. Fifth Hungarian Colloq., Keszthely, 1976), Vol. II, Colloq. Math. Soc. János Bolyai, vol. 18, North-Holland, Amsterdam, 1978, pp. 939–945.
56. W. Samotij, *Stability results for random discrete structures*, to appear in Random Structures Algorithms.
57. D. Saxton and A. Thomason, *Hypergraph containers*, submitted.
58. M. Schacht, *Extremal results for random discrete structures*, submitted.
59. A. Scott, *Szemerédi’s regularity lemma for matrices and sparse graphs*, Combin. Probab. Comput. **20** (2011), 455–466.
60. M. Simonovits, *A method for solving extremal problems in graph theory, stability problems*, Theory of Graphs (Proc. Colloq., Tihany, 1966), Academic Press, New York, 1968, pp. 279–319.
61. T. Szabó and V. H. Vu, *Turán’s theorem in sparse random graphs*, Random Structures Algorithms **23** (2003), 225–234.
62. E. Szemerédi, *On sets of integers containing no k elements in arithmetic progression*, Acta Arith. **27** (1975), 199–245.
63. ———, *Regular partitions of graphs*, Problèmes combinatoires et théorie des graphes (Colloq. Internat. CNRS, Univ. Orsay, Orsay, 1976), Colloq. Internat. CNRS, vol. 260, CNRS, Paris, 1978, pp. 399–401.
64. P. Turán, *Eine Extremalaufgabe aus der Graphentheorie*, Mat. Fiz. Lapok **48** (1941), 436–452.
65. P. Varnavides, *On certain sets of positive density*, J. London Math. Soc. **34** (1959), 358–360.

DEPARTMENT OF MATHEMATICS, UNIVERSITY OF ILLINOIS, 1409 W. GREEN STREET, URBANA, IL 61801

E-mail address: jobal@math.uiuc.edu

IMPA, ESTRADA DONA CASTORINA 110, JARDIM BOTÂNICO, RIO DE JANEIRO, RJ, BRASIL

E-mail address: rob@impa.br

SCHOOL OF MATHEMATICAL SCIENCES, TEL AVIV UNIVERSITY, TEL AVIV 69978, ISRAEL; AND TRINITY COLLEGE, CAMBRIDGE CB2 1TQ, UK

E-mail address: ws299@cam.ac.uk