A Randomized Version of Ramsey's Theorem

Luca Gugelmann* Yury Person† Angelika Steger* Henning Thomas*,1

* Institute of Theoretical Computer Science
ETH Zurich, 8092 Zurich, Switzerland
{lgugelmann|steger|hthomas}@inf.ethz.ch

† Institute of Mathematics
FU Berlin, 14195 Berlin, Germany
person@math.fu-berlin.de

Abstract. The standard randomization of Ramsey's theorem [11] asks for a fixed graph F and a fixed number r of colors: for what densities $p = p(n)$ can we asymptotically almost surely color the edges of the random graph $G(n, p)$ with r colors without creating a monochromatic copy of F. This question was solved in full generality by Rödl and Ruciński [12, 14]. In this paper we consider a different randomization that was recently suggested by Allen et al. [1]. Let $\mathcal{R}_F(n, q)$ be a random subset of all copies of F on a vertex set V_n of size n, in which every copy is present independently with probability q. For which functions $q = q(n)$ can we color the edges of the complete graph on V_n with r colors such that no monochromatic copy of F is contained in $\mathcal{R}_F(n, q)$? We answer this question for strictly 2-balanced graphs F. Moreover, we combine both randomizations and prove a threshold result for the property that there exists an r-edge-coloring of $G(n, p)$ such that no monochromatic copy of F is contained in $\mathcal{R}_F(n, q)$.

1. Introduction

Ramsey theory dates back to a seminal paper by Frank P. Ramsey [11] from 1930. In the context of graph theory his fundamental theorem reads as follows. For all integers $k \geq 2$ there exists an integer $R(k)$ such that no matter how one colors the edges of the complete graph $K_{R(k)}$ with two colors, there will always be a monochromatic copy of K_k. We use the standard arrow notation for this property: for two graphs F and G we have $G \to (F)^e_r$ if every edge-coloring of G with r colors contains a monochromatic copy of F. (For vertex-colorings we write v instead of e.) Ramsey's theorem was generalized in various ways, see [7] for an overview. Loosely speaking one can say that the core of Ramsey theory are statements of the form: every sufficiently large structure contains certain substructures so robustly that they show up monochromatically in every coloring; or formulated intuitively: complete disorder is impossible. Random graphs on the other hand describe structures that are, at least intuitively, as unordered as they can possibly get. Studying Ramsey type phenomena in random graphs is thus a natural and, as it turned out, a challenging task.

This study was initiated by Łuczak, Ruciński, and Voigt [9] who for a given graph F and an integer $r \geq 2$ found the threshold for the vertex case, i.e., the threshold for the property $G(n, p) \to (F)^v_r$. Here $G(n, p)$ denotes the usual binomial random graph $G(n, p)$ with edge probability p. They also established the threshold for the property $G(n, p) \to (K_3)^e_2$. Thereupon, in a series of papers Rödl and

[1]The author was supported by the Swiss National Science Foundation, grant 200021-120284.

Ruciński showed lower bounds for the density threshold of $G(n,p) \to (F)_r^e$ [12], extended the known result about triangles to an arbitrary number of colors [13] and solved the problem in full generality [14]. Formally, their result reads as follows. For every graph G we use $V(G)$ and $E(G)$ to denote its vertex and edge set and by $v(G)$ and $e(G)$ their sizes. For every graph G on at least 3 vertices we set $d_2(G) = (e(G)-1)/(v(G)-2)$ and $d_2(K_1) = 0$, $d_2(2K_1) = 0$ (where $2K_1$ denotes the empty graph on two vertices), and $d_2(K_2) = 1/2$. We write $m_2(G)$ for the so-called *2-density*, formally defined as

$$m_2(G) = \max_{H \subseteq G} d_2(H) \ .$$

Theorem 1 ([12, 14, 8])**.** *For all integers $r \geq 2$ and for every non-empty graph F which is not a forest of stars and paths of length 3 there exist constants $c > 0$ and $C > 0$ such that*

$$\lim_{n\to\infty} \Pr\left[G(n,p) \to (F)_r^e\right] = \begin{cases} 0 & \text{if } p < cn^{-1/m_2(F)} \quad , \\ 1 & \text{if } p > Cn^{-1/m_2(F)} \quad . \end{cases}$$

Note that $p = n^{-1/m_2(F)}$ is the density where we expect that every edge is contained in roughly a constant number of copies of F. This observation can be used to provide an intuitive understanding of the bounds of Theorem 1. If c is very small, then the number of copies of F is a.a.s. (*asymptotically almost surely*, i.e., with probability $1 - o(1)$ if n tends to infinity) small enough that they are so scattered that a coloring without a monochromatic copy of F can be found. If on the other hand C is big then these copies a.a.s. overlap so heavily that every coloring has to induce at least one monochromatic copy of F. Extending this work, Friedgut and Krivelevich [4] proved the thresholds of Theorem 1 to be sharp for most trees, and later Friedgut, Rödl, Ruciński and Tetali [5] showed that this is also the case for F being a triangle.

In this paper we follow up on a different way to introduce randomness into Ramsey theory that was suggested in a recent paper by Allen, Böttcher, Hladký, and Piguet [1]. Note that the setup from Theorem 1 essentially studies the question of how many edges we need to remove from the complete graph (in a random fashion) such that the remaining graph can be colored without a monochromatic clique of size k. (For simplicity we here just consider the case $F = K_k$.) Allen et al. suggest to study the question for the case that we do not care about all cliques K_k, but only want to avoid certain cliques. More formally, assume we have a k-uniform hypergraph $H = (V_n, E(H))$ and we ask: does every coloring of the edges of the complete graph K_n on vertex set V_n with r colors induce a monochromatic k-clique *that forms a hyperedge in H*? We use the notation $K_n \xrightarrow{H} (K_k)_r^e$ to denote this property. The question asked in [1] is the following. Assume $H^{(k)}(n,q)$ is a binomial k-uniform hypergraph with edge probability q. What is the threshold $q = q(n)$ for the property $K_n \xrightarrow{H^{(k)}(n,q)} (K_k)_r^e$. In [1] the authors study the corresponding question for Turán's Theorem [17]? Our first main result not only solves this problem, but – similar as in [1] – also combines it with the classical probabilistic approach by considering a random graph $G(n,p)$ instead of K_n. More precisely, we show:

Theorem 2 (Cliques)**.** *Let $k \geq 3$ and $r \geq 2$ be fixed integers. There exist constants $c = c(k,r) > 0$ and $C = C(k,r) > 0$ such that*

$$\lim_{n\to\infty} \Pr[G(n,p) \xrightarrow{H^{(k)}(n,q)} (K_k)_r^e] = \begin{cases} 0, & \text{if } n^k p^{\binom{k}{2}} q \leq cn^2 p \\ 1, & \text{if } n^k p^{\binom{k}{2}} q \geq Cn^2 p \end{cases} \ .$$

We also study the generalization of Theorem 2 to strictly 2-balanced graphs F instead of cliques. A graph F is *2-balanced* if for every subgraph $J \subseteq F$ we have

$d_2(J) \leq d_2(F)$, and *strictly 2-balanced* if the inequality is strict for every $J \neq F$. Note that in order to prove the 0-statement for the classical probabilistic Ramsey theorem, i.e., Theorem 1, we only need to consider strictly 2-balanced graphs: if we know that it holds for all strictly 2-balanced graphs then it easily follows for all other graphs F by considering an appropriate strictly 2-balanced subgraph. In the context of Theorem 2 such an argument is no longer trivially true. Intuitively the threshold for the $G(n,p)$ part of the theorem is determined by a strictly 2-balanced subgraph which might be distinct from F itself. The hypergraph on the other hand forbids entire copies of F and not just the 2-balanced subgraph. We thus restrict our considerations to strictly 2-balanced graphs F and leave possible further generalizations to future work.

In the following let F denote a strictly 2-balanced graph. Note that for graphs F different from cliques we will in general have more than one possible copy of F on a given vertex set of size $v(F)$. We thus need to specify which of these subgraphs are forbidden to appear monochromatically. To define this formally, we call a subset $\mathcal{R}$ of all unlabeled copies of F in K_n a *restriction set*. Then, for every graph G on n vertices we have $G \xrightarrow{\mathcal{R}} (F)^e_r$ if every r-edge-coloring of G contains a monochromatic copy of F that is contained in $\mathcal{R}$. Moreover, let $\mathcal{R}_F(n,q)$ denote a random subset of all unlabeled copies of F in K_n in which every copy is present independently of the others with probability q.

Theorem 3 (Main Result)**.** *Let F be a strictly 2-balanced graph with at least 3 edges. Let $r \geq 2$ be a fixed integer. There exist constants $c = c(F,r) > 0$ and $C = C(F,r) > 0$ such that*

$$\lim_{n\to\infty} \Pr[G(n,p) \xrightarrow{\mathcal{R}_F(n,q)} (F)^e_r] = \begin{cases} 0, & \text{if } n^{v(F)}p^{e(F)}q \leq cn^2p \\ 1, & \text{if } n^{v(F)}p^{e(F)}q \geq Cn^2p \end{cases}.$$

Note that the statement is not true for strictly 2-balanced graphs with less than 3 edges. For $F = K_{1,2}$ and $q = 1$ the threshold is given by the appearance of the star $K_{1,r+1}$ in $G(n,p)$ which is known to be Poisson-distributed in the regime where p is of order $n^{1-1/(r+1)}$ (cf. [3]).

Moreover, note that Theorem 2 is a special case of Theorem 3. Observe also that for $q = 1$ the threshold matches with the one from Theorem 1, that is, our result can be viewed as a natural generalization of Theorem 1. The proof of the 1-statement of our results builds upon the proof of the 1-statement of Theorem 1. In fact, this proof applies to all graphs F that have $m_2(F) \geq 1$. Our proof of the 0-statement of Theorem 3 on the other hand follows a new approach by further developing algorithmic ideas from [10] to obtain a suitable coloring of $G(n,p)$ (see section 3 in [10]) and combining them with a general theorem from [12] about the global density of graphs which are Ramsey with respect to a given graph, i.e., when $H \nrightarrow (F)^e_2$, see Theorem 4.

2. Proof of Theorem 3

The intuition behind the threshold can be stated similarly to the intuition of Theorem 1 as follows. Call a copy of F in a graph G *bad* with respect to a restriction set $\mathcal{R}$ if it is contained in $\mathcal{R}$. For a fixed edge in $G(n,p)$ we expect order of $n^{v(F)-2}p^{e(F)-1}q$ bad copies of F in $G(n,p)$ with respect to $\mathcal{R}_F(n,q)$ that contain this edge. Hence, if c is small and p and q satisfy the inequality of the 0-statement we expect so few bad copies of F on a fixed edge that these copies are indeed so scattered that we can find an edge-coloring without a monochromatic copy of F. However, if C is large and p and q are as in the 1-statement, then the bad copies of F overlap so heavily that no such coloring can be found.

We now prove Theorem 3. We first address the easier 1-statement. Throughout the remainder of this paper, we consider F to be a fixed strictly 2-balanced graph with at least 3 edges and $r \geq 2$ to be fixed. We write $\mathcal{R}(n,q)$ instead of $\mathcal{R}_F(n,q)$.

2.1. **The 1-statement.** We have to show that for a large enough C and $n^{v(F)}p^{e(F)}q \geq Cn^2p$ the random graph $G(n,p)$ and the random restriction set $\mathcal{R}(n,q)$ a.a.s. have the property that every r-edge-coloring contains a monochromatic bad copy of F. Observe that $n^{v(F)}p^{e(F)}q \geq Cn^2p$ in particular implies $p \geq C'n^{-1/m_2(F)}$. We now use Theorem 3 from [14] which states that there exist constants $a, b > 0$ such that for C' large enough $p \geq C'n^{-1/m_2(F)}$ implies that every r-edge-coloring of $G(n,p)$ contains at least $an^{v(F)}p^{e(F)}$ monochromatic copies of F with probability at least $1 - 2^{-bn^2p}$. Also note that by Chernoff bounds (see e.g. Chapter 2, Theorem 2.1 in [8]) it holds that

$$\Pr\left[e(G(n,p)) \geq 2\binom{n}{2}p\right] \leq \mathrm{e}^{-\Theta(n^2p)}$$

and thus there are a.a.s. at most $r^{2\binom{n}{2}p}$ r-edge-colorings of $G(n,p)$. Let $\mathcal{Q}$ denote the set of all graphs for which every r-edge-coloring contains at least $an^{v(F)}p^{e(F)}$ monochromatic copies of F and which have at most $2\binom{n}{2}p$ edges. Then,

$$\Pr[G(n,p) \notin \mathcal{Q}] \leq 2^{-bn^2p} + \mathrm{e}^{-\Theta(n^2p)} = o(1) \ . \tag{1}$$

Now, the probability that $G(n,p) \overset{\mathcal{R}(n,q)}{\not\longrightarrow} (F)^e_r$ conditioned on that $G(n,p)$ satisfies $\mathcal{Q}$ can be bounded from above with a union bound over all r-edge-colorings by

$$r^{2\binom{n}{2}p}(1-q)^{an^{v(F)}p^{e(F)}} \leq \mathrm{e}^{\ln(r)n^2p - an^{v(F)}p^{e(F)}q} \leq \mathrm{e}^{(\ln r - aC)n^2p} \ , \tag{2}$$

where we used $n^{v(F)}p^{e(F)}q \geq Cn^2p$ in the last step. We now have

$$\begin{aligned}
&\Pr\left[G(n,p) \overset{\mathcal{R}(n,q)}{\not\longrightarrow} (F)^e_r\right] \\
&\leq \Pr\left[G(n,p) \overset{\mathcal{R}(n,q)}{\not\longrightarrow} (F)^e_r \mid G(n,p) \in \mathcal{Q}\right] \cdot \Pr[G(n,p) \in \mathcal{Q}] + \Pr[G(n,p) \notin \mathcal{Q}] \\
&\overset{(1),(2)}{\leq} \mathrm{e}^{(\ln r - aC)n^2p} + o(1) \ .
\end{aligned}$$

Clearly, for large enough C this probability tends to 0. This finishes the proof of the 1-statement.

2.2. **The 0-statement.** Recall that we need to show that for $n^{v(F)}p^{e(F)}q \leq cn^2p$ the random graph $G(n,p)$ and the random restriction set $\mathcal{R}(n,q)$ a.a.s. have the property that there *exists* an r-edge-coloring of $G(n,p)$ that does not contain any monochromatic bad copy of F. We call such a coloring *valid*. In the remainder we describe an algorithm that finds such a coloring and show that it succeeds with high probability.

In a first step we identify a set of edges that we can color easily. Let G be a graph and $\mathcal{R}$ be a restriction set. Let $e \in E(G)$ be an edge which is contained in at most one bad copy of F with respect to $\mathcal{R}$. Then clearly, if there exists a valid coloring for $G - e$ (the graph obtained from G by removing e), we can extend it to one for G since we can assign at least $r - 1 \geq 1$ colors to e without creating a monochromatic bad copy of F. We call such edges *open* with respect to $\mathcal{R}$ and all other edges *closed* (with respect to $\mathcal{R}$). It is easy to see that successively removing open edges yields the unique maximum subgraph of G in which every edge is contained in at least two bad copies of F, where maximum is with respect to subgraph inclusion. We call this subgraph the *F-core* of G (with respect to $\mathcal{R}$). By the above argument, it suffices to find a valid coloring for the F-core of G.

We say that a subgraph H of the F-core of G is *F-closed* with respect to $\mathcal{R}$ if every bad copy of F from the F-core of G is either contained in H or edge-disjoint with H. It is easy to see that the edges of the F-core can be partitioned into minimal F-closed subgraphs where minimal is with respect to subgraph inclusion. Furthermore, each such subgraph can be colored separately in order to find a valid coloring of the F-core. The key property of F-closed subgraphs H follows from the definition of the F-core. For every edge $e \in E(H)$ there are at least two bad copies $F_1, F_2 \subseteq H$ which contain e.

Grow Sequences. In the following we describe a procedure that yields a sequence of bad copies of F which construct an F-closed subgraph. Let F_1 be a bad copy of F from the F-core. Now, for every $\ell \geq 1$, we let $F_{\ell+1}$ be a bad copy of F from the F-core such that $F_{\ell+1}$ intersects $\bigcup_{i=1}^{\ell} F_i$ in at least one edge and for which $F_{\ell+1} \neq F_i$ for every $1 \leq i \leq \ell$. If no such copy exists in the F-core of G the sequence ends after the ℓ-th step and we set $S := (F_1, F_2, \ldots, F_\ell)$ and $G(S) := \bigcup_{i=1}^{\ell} F_i$ as *the graph of S*. We call such a sequence a *grow sequence*, and say that S is *contained* in or *appears* in G with respect to $\mathcal{R}$, meaning that $G(S)$ is a subgraph of G and that every F_i is contained in $\mathcal{R}$. Observe that for every minimal F-closed subgraph H there exists a grow sequence S such that $G(S) = H$. It thus suffices to show that we can find a valid coloring for the graph of every grow sequence that is contained in $G(n,p)$ with respect to $\mathcal{R}(n,q)$.

The following theorem by Rödl and Ruciński states that there exists a valid coloring for a graph whenever it is sparse enough. For a graph H let $d(H) = e(H)/v(H)$ denote the density of H and let $m(H) = \max_{J\subseteq H} d(J)$ denote the maximum density of H.

Theorem 4 ([12]). *Let G and H be two graphs. If $m(H) \leq m_2(G)$ and $m_2(G) > 1$ then $H \nrightarrow (G)^e_2$.*

Note that since F is strictly 2-balanced and has at least 3 edges it satisfies $m_2(F) > 1$. Hence, we can find a valid coloring for every grow sequence S which satisfies $m(G(S)) \leq m_2(F)$. It remains to deal with grow sequences that encode denser subgraphs. The most important structural property of a grow sequence is that every edge of its graph is contained in at least two copies of F. Intuitively speaking, this forces its graph to be dense and to contain many copies of F. On the other hand, the conditions on p and q imply that $G(n,p)$ and $\mathcal{R}(n,q)$ have the property that a.a.s. every subgraph has few edges in $G(n,p)$ or few copies in $\mathcal{R}(n,q)$. In the remainder we use these two density restrictions in the following way. For the *length* of a grow sequence S we write $\ell(S)$. We show that there exists a constant L such that asymptotically almost surely

(i) there are no grow sequences in $G(n,p)$ with respect to $\mathcal{R}(n,q)$ of length more than L, and
(ii) for every grow sequence S of length at most L that appears in $G(n,p)$ with respect to $\mathcal{R}(n,q)$ we have that $G(S)$ is sparse enough so that Theorem 4 applies.

These two conditions are given formally in Lemma 5 and 6 below.

Lemma 5. *For every strictly 2-balanced graph F with at least 3 edges and every integer $r \geq 2$ there exist constants $c = c(F,r) > 0$ and $L = L(F,r) > 0$ such that if $p \leq cn^{-1/m_2(F)}q^{-1/(e(F)-1)}$ then $G(n,p)$ and $\mathcal{R}(n,q)$ a.a.s. satisfy that every grow sequence in $G(n,p)$ (with respect to $\mathcal{R}(n,q)$) has length at most L.*

Note that our assumption $n^{v(F)}p^{e(F)}q \leq cn^2p$ is equivalent to the property $p \leq c'n^{-1/m_2(F)}q^{-1/(e(F)-1)}$ for $c' = c^{1/(e(F)-1)}$.

Lemma 6. *For every strictly 2-balanced graph F with at least 3 edges and every integer $r \geq 2$, and for all constants $c, L > 0$ we have that if $p \leq cn^{-1/m_2(F)}q^{-1/(e(F)-1)}$ then $G(n,p)$ and $\mathcal{R}(n,q)$ a.a.s. satisfy that every grow sequence S in $G(n,p)$ (with respect to $\mathcal{R}(n,q)$) of length at most L satisfies $m(G(S)) \leq m_2(F)$.*

With these two ingredients the proof of the 0-statement from Theorem 3 is straightforward.

Proof (of Theorem 3, 0-statement). Choose $c = c(F,r)$ and $L = L(F,r)$ according to Lemma 5. Then, by Lemmas 5 and 6, $G(n,p)$ and $\mathcal{R}(n,q)$ a.a.s. satisfy that every grow sequence S in $G(n,p)$ with respect to $\mathcal{R}(n,q)$ has length at most L and satisfies $m(G(S)) \leq m_2(F)$. Hence, a.a.s. every F-closed subgraph H in $G(n,p)$ with respect to $\mathcal{R}(n,q)$ satisfies $m(H) \leq m_2(F)$. Conditioning on this property of $G(n,p)$ and $\mathcal{R}(n,q)$ Theorem 4 guarantees that we can find a valid r-edge-coloring for every such subgraph. Moreover, the union of the valid colorings of all F-closed subgraphs yields a valid coloring of the F-core of $G(n,p)$ which can be extended to a valid coloring of $G(n,p)$. □

3. Proof of Lemma 6

We will on several occasions use the following proposition.

Proposition 7. *Let $a, c, \beta \geq 0$ and $b, d > 0$. Then we have*

$$\frac{a}{b} > \beta \quad \text{and} \quad \frac{c}{d} \geq \beta \quad \Longrightarrow \quad \frac{a+c}{b+d} > \beta \ , \text{ and}$$
$$\frac{a}{b} < \beta \quad \text{and} \quad \frac{c}{d} \leq \beta \quad \Longrightarrow \quad \frac{a+c}{b+d} < \beta \ .$$

Proof (of Lemma 6). As a first step of the proof we show that every grow sequence $S = (F_1, F_2, \ldots, F_\ell)$ of length at most L that satisfies $m(G(S)) > m_2(F)$ in fact satisfies $d(G(S)) > m_2(F)$. We show this by inductively constructing a finite sequence of graphs $H_0, H_1, \ldots$ all of which have density at least $m_2(F)$ and which end in the graph $G(S)$. Let $H_0 \subseteq G(S)$ be an arbitrary densest subgraph of $G(S)$, i.e., $d(H_0) = m(G(S)) > m_2(F)$. If $H_i = H$, we are done. Otherwise there exists a copy F_j in S that is neither entirely contained in H_i nor edge-disjoint from H_i. Hence, F_j overlaps with H_i in a subgraph J that contains at least one edge, that is, $e(J) \geq 1$ and $v(J) \geq 2$. We set $H_{i+1} = H_i \cup F_j$. Denote by e_i and v_i the number of new edges and new vertices, i.e.,

$$e_i = e(F) - e(J) \quad \text{and} \quad v_i = v(F) - v(J) \ . \tag{3}$$

If there are no new vertices ($v_i = 0$), then clearly $d(H_{i+1}) \geq d(H_i) > m_2(F)$. Hence, we may assume from now on that there is at least one new vertex ($v_i \geq 1$). We show $e_i/v_i \geq m_2(F)$ with a case distinction.

Case 1: $e(J) = 1$. In this case,

$$\frac{e_i}{v_i} \overset{(3)}{=} \frac{e(F) - e(J)}{v(F) - v(J)} \overset{v(J)\geq 2}{\geq} \frac{e(F) - 1}{v(F) - 2} = m_2(F) \ .$$

Case 2: $e(J) \geq 2$. Clearly, this implies $v(J) \geq 3$. Recall that F is strictly 2-balanced and that $J \neq F$ since there is at least one new vertex. Thus, $e(J) - 1 < m_2(F)(v(J) - 2)$, and we can settle this case with

$$\begin{aligned} \frac{e_i}{v_i} \overset{(3)}{=} \frac{e(F) - e(J)}{v(F) - v(J)} &= \frac{e(F) - 1 - (e(J) - 1)}{v(F) - 2 - (v(J) - 2)} \\ &> \frac{m_2(F)(v(F) - 2) - m_2(F)(v(J) - 2)}{v(F) - 2 - (v(J) - 2)} = m_2(F) \ . \end{aligned} \tag{4}$$

Hence, we have

$$d(H_{i+1}) = \frac{e(H_{i+1})}{v(H_{i+1})} = \frac{e(H_i) + e_i}{v(H_i) + v_i} > m_2(F) \ ,$$

where in the last step we used Proposition 7 together with the assumption $d(H_i) = e(H_i)/v(H_i) > m_2(F)$ and $e_i/v_i \geq m_2(F)$. Clearly, after a constant number of iterations we arrive at $H_i = G(S)$.

We now show that the probability that there exists a grow sequence S of length at most L which satisfies $m(G(S)) > m_2(F)$ and which appears in $G(n,p)$ with respect to $\mathcal{R}(n,q)$ is $o(1)$. Observe that a grow sequence of length at most L can involve at most $v_L := 2 + L \cdot (v(F) - 2)$ vertices. Now fix $k \leq v_L$ and let U be a fixed vertex set of size k. We show that the probability that there is a grow sequence S of length at most L with $m(G(S)) > m_2(F)$ and $v(G(S)) = k$ that appears on U is $o(n^{-k})$. With a union bound over all k and all vertex sets U of size k this then concludes the proof of the lemma.

Observe that there are at most $\mathcal{O}(1)$ possible grow sequences that can be accommodated on the vertex set U. Consider an arbitrary but fixed possible grow sequence S of length at most L that is contained in U and that satisfies $v(G(S)) = k$ and $m(G(S)) > m_2(F)$. First observe that since $H := G(S)$ is F-closed we have that for every edge there must be two copies of F in S that contain it. Since every such copy contains $e(F)$ edges S must have length at least $2e(H)/e(F)$. Hence, using $p \leq cn^{-1/m_2(F)}q^{-1/(e(F)-1)}$ the probability that S appears on U in $G(n,p)$ with respect to $\mathcal{R}(n,q)$ can be bounded by

$$\begin{aligned} p^{e(H)}q^{2e(H)/e(F)} &\leq c^{e(H)}n^{-e(H)/m_2(F)} \underbrace{q^{2e(H)/e(F)-e(H)/(e(F)-1)}}_{\leq 1, \text{ since } e(F)\geq 3} \\ &\leq c^{e(H)}n^{-e(H)/m_2(F)} \overset{d(H)>m_2(F)}{=} o(n^{-v(H)}) = o(n^{-k}) \ . \end{aligned}$$

This concludes the proof of the lemma. □

4. Proof of Lemma 5

Canonical Grow Sequences. We first take a closer look at grow sequences and describe a unique canonical way to construct them. For this, let G be a graph, $\mathcal{R}$ be a restriction set, and fix an arbitrary ordering of the vertices of G. Note that this vertex ordering also naturally induces an ordering on the edges of G. Using these orderings we can say for two vertex or edge sets X and Y which of the two is *lexicographically smaller*. Now, let H be a minimal F-closed subgraph of G with respect to $\mathcal{R}$.

We can now construct a grow sequence S for H inductively in the following canonical way. Let $\mathcal{F}(H)$ denote the set of all bad copies of F in H. We set F_1 to be the lexicographically smallest copy of $\mathcal{F}(H)$. Now, for every $\ell \geq 1$, let $G(S,\ell) := \bigcup_{i=1}^{\ell} F_i$. If $G(S,\ell)$ contains an open edge with respect to the restriction set $\{F_1, F_2, \dots, F_\ell\}$, then we let e denote the lexicographically smallest open edge. Since H is F-closed, there must be at least one bad copy of F in $\mathcal{F}(H)$ that is not in $G(S,\ell)$ and contains e. We set $F_{\ell+1}$ to be the lexicographically smallest such copy. Otherwise, if $G(S,\ell)$ only contains closed edges, then we choose $F_{\ell+1}$ to be the lexicographically smallest copy in $\mathcal{F}(H)$ which is not edge-disjoint from $G(S,\ell)$ and not yet in the sequence $F_1, F_2, \dots, F_\ell$.

We call a grow sequence *canonical* if it follows the above procedure. Then, every F-closed subgraph has exactly one corresponding canonical grow sequence. Hence, in order to prove Lemma 5 it suffices to show that $G(n,p)$ a.a.s. does not contain a canonical grow sequence of size more than L with respect to $\mathcal{R}(n,q)$.

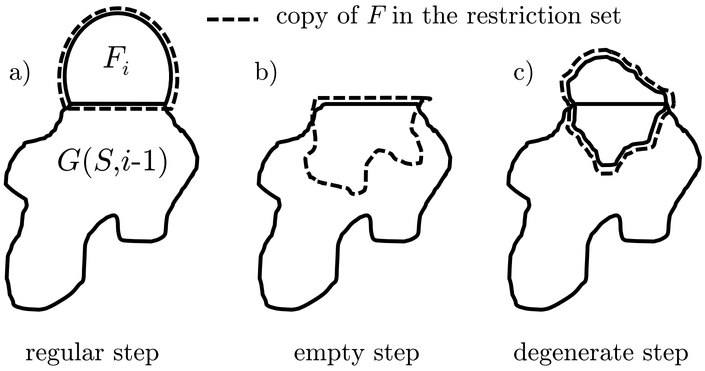

FIGURE 1. Illustration of the different step types of a grow sequence

We point out here that the crucial property of canonical grow sequences is the following. In case that $G(S,i)$ has open edges we know that the copy F_{i+1} contains the lexicographically smallest of them. Roughly speaking, this property will turn out to be important in a later first moment method argument: when counting the number of grow sequences there are no choices for these two vertices.

Step Types. Let $S = (F_1, F_2, \ldots, F_\ell)$ be a canonical grow sequence. We now view S from the perspective of building a graph step by step from $G(S,0)$ over $G(S,1)$ and so on up to $G(S)$. We split the steps of this process into different types, which are also illustrated in Figure 1. We call step one the *first step.* We call a step $i \geq 2$ *regular* if the intersection of F_i with $G(S,i-1)$ consists exactly of two vertices connected by an edge (Figure 1 *a*)), *empty* if the intersection is the whole copy F_i (Figure 1 *b*)) and *degenerate* otherwise (Figure 1 *c*)). Note that an empty step only imposes a copy of F in the restriction set. Besides categorizing the steps into regular, empty and degenerate ones we introduce another categorization. We call step i *open* if $G(S,i-1)$ contains open edges with respect to the restriction set $\{F_1, F_2, \ldots, F_{i-1}\}$ and *closed* otherwise. Note that since S is canonical the copy F_i of an open step i contains the lexicographically smallest open edge in $G(S,i-1)$.

Let v_i denote the number of new vertices added in step i, i.e., $v_i := v(G(S,i)) - v(G(S,i-1))$ and let similarly e_i denote the number of new edges. It is easy to see that for every regular and every empty step we have $v_i - e_i/m_2(F) = 0$, and that for every degenerate step $v_i - e_i/m_2(F) < 0$ since F is strictly 2-balanced (using a calculation similar to (4)). Furthermore, since the number of possible intersections of F_i with $G(S,i-1)$ is constant, there exists a constant $\delta = \delta(F) > 0$ such that every degenerate step satisfies

$$v_i - e_i/m_2(F) < -\delta \ . \tag{5}$$

Before we continue we give an intuitive reasoning of how we will use (4). We will use a first moment method to show that a.a.s. grow sequences of length more than L do not appear. In order to do so we need to *count* grow sequences. To do that we will distinguish grow sequences according to their number of degenerate steps. Observe that for a degenerate step we essentially have to choose v_i new vertices (at most n^{v_i} ways) and require the presence of e_i new edges and a new copy of F in the restriction set (which happens with probability $p^{e_i}q$). As $p \leq cn^{-1/m_2(F)}q^{-1/(e(F)-1)}$, we see that (5) implies that the probability for K degenerate steps decreases with $n^{-\delta K}$. That is, we should not expect that we have grow sequences with 'many' degenerate steps. In order to handle grow sequences with 'few' degenerate steps we will now prove some deterministic properties of canonical grow sequences.

We show that long sequences have to contain a certain number of degenerate steps, and that the number of closed regular steps and empty steps can be bounded in terms of the number of degenerate steps. For a grow sequence S we use $\mathrm{degen}(S), \mathrm{empty}(S), \mathrm{regular}^o(S), \mathrm{regular}^c(S)$ to denote the number of degenerate, empty, open regular and closed regular steps in S. Note that $\mathrm{degen}(S) + \mathrm{empty}(S) + \mathrm{regular}^o(S) + \mathrm{regular}^c(S) = \ell(S) - 1$, where the -1 accounts for the first step. Moreover, we write $X_F(G)$ for the number of (not necessarily bad) copies of F in G.

Claim 8. *Let $S = (F_1, F_2, \dots, F_\ell)$ be a grow sequence. If step i is empty or regular, then $X_F(G(S,i)) \le X_F(G(S,i-1)) + 1$.*

Proof. The claim is trivial if step i is an empty step. Otherwise step i is regular, and we can denote by $e = \{u,v\}$ the intersection of F_i with $G(S,i-1)$. Assume that in addition to the copy F_i there is another copy $\tilde{F}$ of F created in that step. Then clearly, $\tilde{F}$ contains at least one edge from the new edges of step i, i.e., from $E(F_i) \setminus \{e\}$, and at least one edge from the old edges $E(G(S,i-1)) \setminus \{e\}$. Hence, the graphs $\tilde{F}_{\mathrm{new}} = \tilde{F}[V(F_i)]$ and $\tilde{F}_{\mathrm{old}} = \tilde{F}[V(G(S,i-1))]$ are both non-empty. Since every strictly 2-balanced graph is 2-vertex-connected it follows that both u and v are contained in $V(\tilde{F})$, $V(\tilde{F}_{\mathrm{new}})$ and $V(\tilde{F}_{\mathrm{old}})$. Moreover, removing u and v disconnects $\tilde{F}$.
Case 1: $e \in E(\tilde{F})$. Then we have

$$m_2(F) = m_2(\tilde{F}) = \frac{e(\tilde{F}) - 1}{v(\tilde{F}) - 2} = \frac{e(\tilde{F}_{\mathrm{new}}) - 1 + e(\tilde{F}_{\mathrm{old}}) - 1}{v(\tilde{F}_{\mathrm{new}}) - 2 + v(\tilde{F}_{\mathrm{old}}) - 2} < m_2(F) \ ,$$

where the last step follows from Proposition 7 and the fact that F is strictly 2-balanced. Clearly, this is a contradiction.
Case 2: $e \notin E(\tilde{F})$. Then we have

$$m_2(F) = m_2(\tilde{F}) = \frac{e(\tilde{F}) - 1}{v(\tilde{F}) - 2} = \frac{e(\tilde{F}_{\mathrm{new}}) + e(\tilde{F}_{\mathrm{old}}) - 1}{v(\tilde{F}_{\mathrm{new}}) - 2 + v(\tilde{F}_{\mathrm{old}}) - 2} \ . \tag{6}$$

Since F is strictly 2-balanced we have $(e(\tilde{F}_{\mathrm{old}})-1)/(v(\tilde{F}_{\mathrm{old}})-2) < m_2(F)$. Hence, (6) together with Proposition 7 implies

$$\frac{e(\tilde{F}_{\mathrm{new}}) + 1 - 1}{v(\tilde{F}_{\mathrm{new}}) - 2} = \frac{e(\tilde{F}_{\mathrm{new}})}{v(\tilde{F}_{\mathrm{new}}) - 2} > m_2(F) \ .$$

Since $\tilde{F}_{\mathrm{new}} \cup \{e\}$ is a subgraph of F_i with $e(\tilde{F}_{\mathrm{new}}) + 1$ edges and $v(\tilde{F}_{\mathrm{new}})$ vertices, this contradicts the property that F is strictly 2-balanced. □

Claim 9. *For every canonical grow sequence $S = (F_1, \dots, F_\ell)$ we have* $\mathrm{regular}^c(S) \le \mathrm{degen}(S)$. *Moreover, for every prefix sequence $S_i = (F_1, \dots, F_i)$, where $i \le \ell$, we have* $\mathrm{regular}^c(S_i) \le \mathrm{degen}(S_i)$.

Proof. Let $S = (F_1, F_2, \dots, F_\ell)$ be a canonical grow sequence, and let $r_1, r_2, \dots$ denote the indices of its closed regular steps and set $r_0 = 1$ for convenience. We show that for every $i \ge 0$ there is at least one degenerate step between steps r_i and r_{i+1}. Clearly, $G(S, r_i)$ contains $e(F) - 1$ open edges with respect to the restriction set $\{F_1, F_2, \dots, F_{r_i}\}$ and thus, step $r_i + 1$ is open. Moreover, since r_i is a closed step all edges in $G(S, r_i - 1)$ are closed which together with Claim 8 implies that the *only* copy of F in $G(S, r_i)$ that contains open edges is F_{r_i}. Since S is canonical, step $r_i + 1$ cannot be empty (an empty step would require a copy of F in $G(S, r_i)$ that contains open edges and is not yet in S). Step $r_i + 1$ is thus either degenerate or an open regular step. In the former case we are done, and in the latter case we can apply the above argument to obtain that step $r_i + 2$ is either degenerate or open regular and so on. Eventually, there must be a degenerate step between steps r_i and r_{i+1} since $G(S, r_{i+1}-1)$ does not contain open edges with respect to $\{F_1, F_2, \ldots, F_{r_{i+1}-1}\}$. □

Claim 10. *There exist constants $c_1 = c_1(F)$ and $c_2 = c_2(F)$ such that for every canonical grow sequence S we have*

$$\text{empty}(S) \leq \big(c_1 \cdot \text{degen}(S)\big)^{c_2} \ .$$

Proof. Let $S = (F_1, F_2, \ldots, F_\ell)$ be a canonical grow sequence. Assume that the i-th step of the sequence S is an empty step. Then $F_i \subseteq G(S, i-1)$ and $F_i \neq F_j$ for every $1 \leq j \leq i-1$. Note that such a copy can only be created in a degenerate step since by Claim 8 the only copy of F created in a regular step is the copy of the step itself. Hence, F_i must contain an edge that was added in a degenerate step. We call such edges *degenerate*. Let $X_{F,e}(G)$ denote the number of copies of F in G that contain edge e. Then we have

$$\text{empty}(S) \leq \sum_{\substack{e \in E(G(S)) \\ e \text{ is degenerate}}} X_{F,e}(G(S)) \ . \tag{7}$$

Since there are at most $e(F)\text{degen}(S)$ degenerate edges in $G(S)$ it clearly suffices to bound $X_{F,e}(G(S))$ for every edge e by a bound similar to the one in the claim statement. Instead of bounding this quantity directly, let C_k^e denote for every edge $e \in E(G(S))$ the number of *induced* cycles of length k in $G(S)$ (where $3 \leq k \leq v(F)$) that contain e. Moreover, let C_k denote the maximum over all edges in $G(S)$. We will show that there exist constants $c_1' = c_1'(F)$ and $c_2' = c_2'(F)$ such that

$$C_k \leq (c_1' \cdot \text{degen}(S))^{c_2'} \ . \tag{8}$$

If we assume (8), then we can bound $X_{F,e}(G(S))$ for a fixed edge $e \in E(G(S))$ as follows. Let $\mathcal{F}$ denote the set of all supergraphs of F on $v(F)$ vertices. Clearly, for every copy of F in $G(S)$ that contains e, there is also an *induced* copy of some $J \in \mathcal{F}$ in $G(S)$ that contains e. Moreover, such an induced copy J can accommodate only a constant number of copies of F that contain e. Let $\bar{X}_{J,e}(G)$ denote the number of induced copies of J in G that contain edge e. As $|\mathcal{F}|$ is bounded by a constant (depending on F) it suffices to bound $\bar{X}_{J,e}(G)$ for every $J \in \mathcal{F}$ with a bound similar to the one in the claim statement.

Let $J \in \mathcal{F}$ be fixed. For every edge $e' \in E(J)$ fix an induced cycle $C(e')$ in J that contains e'. Note that this is possible since F is 2-connected. Clearly, the union of all these $e(J) \leq v(F)^2$ cycles covers $E(J)$, that is,

$$\bigcup_{e' \in E(J)} E(C(e')) = E(J) \ .$$

We can now count the number of induced copies of J in $G(S)$ that contain e as follows. First choose the role of e in the copy of J ($e(J)$ possibilites). Say e has the role of $e' \in E(J)$. Then we have to choose an induced copy of $C(e')$ in $G(S)$ that contains e (at most $(c_1' \cdot \text{degen}(S))^{c_2'}$ possibilities by (8)). Continue with some other edge of that cycle and choose its role in the copy of J, and then choose a copy of the corresponding induced cycle on it and so on. In this way we obtain

$$\begin{aligned} \bar{X}_{J,e}(G(S)) &\leq \big(e(J) \cdot (c_1' \cdot \text{degen}(S))^{c_2'}\big)^{e(J)} \\ &\overset{e(J) \leq v(F)^2}{\leq} (c_1'' \cdot \text{degen}(S))^{c_2''} \end{aligned} \tag{9}$$

for appropriately chosen constants $c_1'' = c_1''(F)$ and $c_2'' = c_2''(F)$. As explained above this concludes the proof.

It remains to show (8). For every $0 \leq i \leq \ell$ let $S_i = (F_1, F_2, \ldots, F_i)$ denote the prefix of S including the first i steps. For every k let $P^{k,i}_{\{u,v\}}$ denote the number of induced paths of length k with endpoints u and v in $G(S_i) - uv$ for $k \geq 2$ and in $G(S_i)$ for $k = 1$. (A path of length k is a path with k edges.) Denote with $P^{k,i}$ the maximum of $P^{k,i}_{\{u,v\}}$ over all vertex pairs $\{u, v\}$. We will show that there exists a constant $\beta = \beta(F) > 0$ such that for every $1 \leq k \leq v(F) - 1$ and every $1 \leq i \leq \ell$ we have

$$P^{k,i} \leq (\beta(\mathrm{degen}(S_i) + 1))^{k-1} \ . \tag{10}$$

As $C_k \leq P^{k-1,\ell}$ for all $3 \leq k \leq v(F)$ this establishes (8). It thus remains to show (10). First observe that (10) is trivially true for $k = 1$ as $P^{1,i} \leq 1$ for every i and we therefore assume $k \geq 2$ in the remainder. Let $\{u, v\}$ be a fixed vertex pair. We will show that if β is large enough, then we have for every $k \geq 2$ and every $1 \leq i \leq \ell$ that

$$P^{k,i}_{\{u,v\}} \leq \frac{\beta}{2}\big(\mathrm{regular}^{\{u,v\}}(S_i) + 1\big) + 6e(F) \sum_{j=1}^{\mathrm{degen}(S_i)+1} (\beta j)^{k-2} \ , \tag{11}$$

where $\mathrm{regular}^{\{u,v\}}(S_i)$ denotes the number of regular steps $j \leq i$ in which we attach a copy of F to the edge $\{u, v\}$, that is, for which $V(F_j) \cap V(G(S, j-1)) = \{u, v\}$. First we show that (11) implies (10). For this, observe that $V(F_j) \cap V(G(S, j-1)) = \{u, v\}$ is satisfied for at most one open regular step and at most $\mathrm{regular}^c(S_i)$ closed regular steps $j \leq i$, and hence by Claim 9, $\mathrm{regular}^{\{u,v\}}(S_i) \leq \mathrm{degen}(S_i) + 1$. Using $k \leq v(F)$ we thus easily conclude that for a sufficiently large β we have that (11) implies (10). It thus remains to show (11) for every $k \geq 2$ and every $1 \leq i \leq \ell$. For this, we do an induction on k and i. More precisely, we will use that for all smaller values of k we have (10) (the trivial case $k = 1$ serves as induction base) and for all smaller values of i we have (11). We make a case distinction according to the step type of step i.

Case 1. Step i is regular and $P^{k,i-1}_{\{u,v\}} = 0$, or $i = 1$. Using that $P^{k,i}_{\{u,v\}}$ counts *induced* paths, it is not hard to see that $P^{k,i}_{\{u,v\}}$ can only be non-zero if at least one of the two vertices $\{u, v\}$ lies within the new vertices added in step i, that is, if $\{u, v\} \cap (V(F_i) \setminus V(G(S, i-1))) \neq \emptyset$. Clearly, if both u and v lie within $V(F_i)$, there can be at most $v(F)^k \leq v(F)^{v(F)}$ induced paths of length k between u and v and thus (11) is satisfied if $\beta \geq 2v(F)^{v(F)}$. Otherwise u must lie within the old vertices, i.e., $u \in V(G(S, i-1)) \setminus V(F_i)$ and v in the new vertices, or vice versa. Then, we can count the number of induced paths of length k from u to v as follows (see also Figure 2 a)). Let $\{x, y\} = V(F_i) \cap V(G(S, i-1))$, and observe that in this case $\{u, v\} \cap \{x, y\} = \emptyset$. Now, every such path consists of an induced path of length s from u to x or y for some $1 \leq s \leq k-1$ and an induced path of length $k - s$ from x or y to v (clearly, there are at most $v(F)^{v(F)}$ such paths since they are contained in F_i). Note that paths containing both x and y can also be expressed this way by simply including the edge $\{x, y\}$ in one of the two subpaths of length s and $k - s$. Hence, we obtain

$$\begin{aligned} P^{k,i}_{\{u,v\}} &\leq \sum_{s=1}^{k-1} P^{s,i}_{\{u,x\}} P^{k-s,i}_{\{x,v\}} + \sum_{s=1}^{k-1} P^{s,i}_{\{u,y\}} P^{k-s,i}_{\{y,v\}} \leq 4P^{k-1,i} + 2v(F)^{v(F)} \sum_{s=2}^{k-2} P^{s,i} \\ &\overset{(10), k \leq v(F)}{\leq} 4(\beta(\mathrm{degen}(S_i) + 1))^{k-2} + 2v(F)^{v(F)+1}(\beta(\mathrm{degen}(S_i) + 1))^{k-3} \\ &\leq 6(\beta(\mathrm{degen}(S_i) + 1))^{k-2} \ , \end{aligned}$$

where the last step holds whenever β is sufficiently large. Clearly, this implies (11).

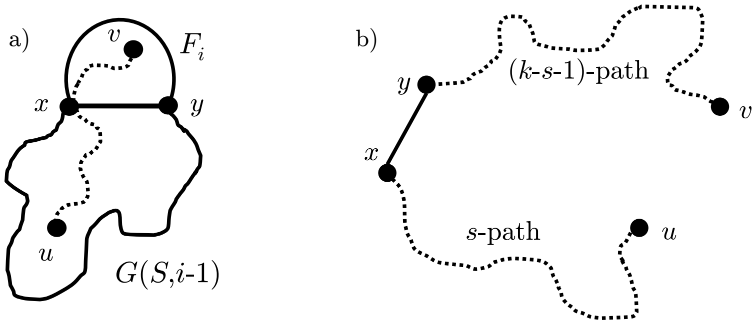

FIGURE 2. Creating a new induced path of length k from u to v

Case 2. Step i is regular and $P^{k,i-1}_{\{u,v\}} \neq 0$. It is easy to see that in this case we can only have $P^{k,i}_{\{u,v\}} > P^{k,i-1}_{\{u,v\}}$ if the copy F_i is attached to the edge $\{u,v\}$, that is, $V(F_i) \cap V(G(S, i-1)) = \{u,v\}$. Hence, $\text{regular}^{\{u,v\}}(S_i) = \text{regular}^{\{u,v\}}(S_{i-1}) + 1$ and step i can create at most $v(F)^k \leq v(F)^{v(F)}$ new induced paths of length k between u and v. Thus (11) certainly remains true if $\beta \geq 2v(F)^{v(F)}$.
Case 3. Step i is degenerate. In this case we insert the edges of F_i one by one. So assume a single edge $\{x,y\}$ is inserted. Clearly, this can only create a new induced path of length k from u to v if for some $0 \leq s \leq k-1$ we have that there is an induced path of length s from u to x and an induced path of length $k-s-1$ from v to y or similarly for paths from u to y and from v to x (see Figure 2 b)). Observe that the case $\{x,y\} = \{u,v\}$ does not create any new induced paths of length k since $k \geq 2$. Furthermore, the case $|\{x,y\} \cap \{u,v\}| = 1$ corresponds to the case $s = 0$ or $s = k-1$, and we set $P^{0,j} = 1$ for convenience. Then, the number of induced paths of length k from u to v that were created by the insertion of $\{x,y\}$ is at most

$$\begin{aligned}
\sum_{s=0}^{k-1} P^{s,j}_{\{u,x\}} P^{k-s-1,j}_{\{v,y\}} &+ \sum_{s=0}^{k-1} P^{s,j}_{\{u,y\}} P^{k-s-1,j}_{\{v,x\}} \\
&\leq 4P^{k-1,j} + 2\sum_{s=1}^{k-2} P^{s,j} P^{k-s-1,j} \\
&\leq 6(\beta(\text{degen}(S_i)+1))^{k-2} \ ,
\end{aligned}$$

cf. the calculation above. Since we insert at most $e(F)$ edges in step i we thus have

$$P^{k,i}_{\{u,v\}} \leq P^{k,i-1}_{\{u,v\}} + e(F)6(\beta(\text{degen}(S_i)+1))^{k-2} \ .$$

As the assumption of this case is that step i is degenerate this implies that (11) holds in this case as well. This thus concludes the proof of the claim □

Claim 11. *For every $d \geq 1$ there exists a constant $\ell_{\max}(d)$ such that every canonical grow sequence with at most d degenerate steps has length at most $\ell_{\max}(d)$.*

Proof. Let $d \geq 1$ be a constant and let S be a grow sequence of length ℓ. By Claim 10, there can be at most $C' := (c_1 d)^{c_2} + d + 1$ steps that are not regular. Hence, S contains at least $\ell - C'$ regular steps each of which increases the number of open edges by at least $e(F) - 2$. Moreover, every non-regular step decreases the number of open edges by at most $e(F)$. Hence, since $G(S)$ is F-closed we have $(\ell - C')(e(F) - 2) \leq C'e(F)$ and thus $\ell \leq \frac{e(F)}{e(F)-2}C' + C'$. □

We can now turn to the proof of Lemma 6.

Proof (of Lemma 5). Let $\mathcal{S}$ denote the set of all canonical grow sequences of length more than $L = L(F)$ (we will fix this constant later). Note that by a first moment argument it suffices to show that for an appropriate constant $c = c(F)$ we have that if $p \leq cn^{-1/m_2(F)}q^{-1/(e(F)-1)}$ the expected number of sequences from $\mathcal{S}$ contained in $G(n,p)$ with respect to $\mathcal{R}(n,q)$ is $o(1)$.

Note that if we can show for a sequence $S \in \mathcal{S}$ that a *prefix sequence* of it, i.e., a sequence obtained by considering only the first k steps for some $1 \leq k \leq \ell(S)$, is not contained in $G(n,p)$ with respect to $\mathcal{R}(n,q)$ then S as well does not appear. Hence, if we can find a set $\mathrm{Pre}(\mathcal{S})$ such that all sequences of $\mathcal{S}$ have a prefix sequence in $\mathrm{Pre}(\mathcal{S})$ and such that a.a.s. no sequence from $\mathrm{Pre}(\mathcal{S})$ appears in $G(n,p)$ with respect to $\mathcal{R}(n,q)$ then we are done.

With this in mind we define $\mathrm{Pre}(\mathcal{S})$ as the set of the following prefixes. Let $d_{\max} = d_{\max}(F)$ be a constant which we will determine later on. For each $S \in \mathcal{S}$ we include the prefix sequence of S containing either all steps up to (and including) the $d_{\max}$-th degenerate step or all steps up to the $\log n$-th step, if the index of the $d_{\max}$-th degenerate step is larger than $\log n$. Note that this is well defined as we can force any sequence $S \in \mathcal{S}$ to contain at least $d_{\max}$ many degenerate steps by choosing L large enough, cf. Claim 11.

The key intuition is that prefixes $S \in \mathrm{Pre}(\mathcal{S})$ containing 'many' degenerate steps give rise to a very dense graph $G(S)$, which correspondingly is unlikely to appear. On the other hand prefixes $S \in \mathrm{Pre}(\mathcal{S})$ with 'few' degenerate steps must contain many regular steps. The corresponding graph $G(S)$ will then be very large and also unlikely to appear.

As each prefix sequence in $\mathrm{Pre}(\mathcal{S})$ contains at most $d_{\max}$ degenerate steps we have by Claim 9 and Claim 10 that they also contain at most $d_{\max}$ closed regular and at most $(c_1 d_{\max})^{c_2} =: e_{\max}$ empty steps. Let $m = 2d_{\max} + e_{\max}$ denote the maximum number of steps that are not open regular in any prefix sequence. Note that m is a fixed constant depending only on F. Hence, we can choose L such that $L > m$ holds.

We define $\mathrm{Pre}^{d_{\max}}(\mathcal{S})$ as the set containing all prefix sequences from $\mathrm{Pre}(\mathcal{S})$ with exactly $d_{\max}$ degenerate steps and $\mathrm{Pre}^{\log n}(\mathcal{S})$ as the set of those with length exactly $\log n$. Clearly, every prefix sequence in $\mathrm{Pre}(\mathcal{S})$ is in at least one of the two subsets. We consider both subsets separately.

Sequences with $\log n$ steps. We start with prefix sequences in $\mathrm{Pre}^{\log n}(\mathcal{S})$. We can bound the number of elements in $\mathrm{Pre}^{\log n}(\mathcal{S})$ by counting all sequences of steps of length $\log n$ which contain at most m steps that are not open regular. To do so we first fix the number of steps that are not open regular, their types (i.e., open or closed, and regular, empty, degenerate) and their position in the sequence. For this we have at most

$$m(5 \log n)^m \tag{12}$$

choices. Then for any sequence with a fixed configuration of steps we have at most $n^{v(F)}$ choices for the copy of a step that is not open regular. As there are at most m of these we have in total at most

$$n^{v(F)m} \tag{13}$$

different choices for these steps. All remaining steps are open regular steps. Recall that for every open step two vertices of its copy of F are determined by all previous steps since all sequences in $\mathcal{S}$ are canonical. So for each of the at most $\log n$ open regular steps we only need to choose $v(F) - 2$ new vertices and the role of the two predetermined vertices in the copy of F. Thus, every open regular step gives at most

$$v(F)^2 n^{v(F)-2} \tag{14}$$

choices. The very first step is special and we model it by choosing two vertices as the starting edge (n^2 choices), and another $v(F)^2 n^{v(F)-2}$ choices as in an open regular step. Together with (12), (13) and (14) we therefore have that the number of elements in $\mathrm{Pre}^{\log n}(\mathcal{S})$ is at most

$$m\big(5\log(n)\big)^m \cdot n^{v(F)m} \cdot n^2 \cdot \big(v(F)^2 n^{v(F)-2}\big)^{\log n} \quad .$$

For a fixed sequence in $\mathrm{Pre}^{\log n}(S)$ the probability that it appears in $G(n,p)$ with respect to $\mathcal{R}(n,q)$ is bounded from above by the probability that the first $\log n - m$ open regular steps appear. Each such step requires $e(F)-1$ new edges to be present in $G(n,p)$ and one new copy of F to be present in $\mathcal{R}(n,q)$, so this probability is at most $(p^{e(F)-1}q)^{\log n - m}$. Using $n^{v(F)-2}p^{e(F)-1}q \leq c$ we can now deduce that the number $X_{\mathrm{Pre}^{\log n}(\mathcal{S})}$ of sequences from $\mathrm{Pre}^{\log n}(\mathcal{S})$ that appear in $G(n,p)$ with respect to $\mathcal{R}(n,q)$ satisfies

$$\begin{aligned}\mathbf{E}[X_{\mathrm{Pre}^{\log n}(\mathcal{S})}] &\leq m\big(5\log(n)\big)^m \cdot n^{v(F)m} \cdot n^2 \cdot \big(v(F)^2 n^{v(F)-2}\big)^{\log n} \cdot (p^{e(F)-1}q)^{\log n - m} \\ &\leq m\big(5\log(n)\big)^m \cdot n^{2v(F)m+2} \cdot v(F)^{2\log n} \cdot c^{\log n - m} \\ &= o(n) \cdot n^{2v(F)m+2+2\log(v(F))-\log(1/c)} \quad .\end{aligned}$$

As m is a constant depending only on F we can choose $c = c(F)$ such that the above expectation is $o(1)$.

Sequences with $d_{\max}$ degenerate steps. It remains to consider the prefix sequences in $\mathrm{Pre}^{d_{\max}}(\mathcal{S})$, i.e., those which contain exactly $d_{\max}$ degenerate steps. We partition these sequences further into sets $\mathrm{Pre}^{d_{\max}}(V,E,O,\ell)$ which contain all sequences from $\mathrm{Pre}^{d_{\max}}(\mathcal{S})$ of length exactly ℓ for which the total number of new vertices and new edges added in the $d_{\max}$ degenerate steps are exactly V and E respectively, and for which the number of empty steps is exactly O. Clearly, this set can only be non-empty if $1 \leq V \leq d_{\max}(v(F)-3)$, $1 \leq E \leq d_{\max}(e(F)-2)$, $0 \leq O \leq e_{\max}$ and $1 \leq \ell \leq \log n$. Hence, the total number of subsets that we need to consider is bounded by

$$d_{\max}^2(v(F)-3)(e(F)-2)e_{\max}\log n \leq \log(n)^2 \tag{15}$$

for n large enough.

Recall that for every degenerate step the numbers v_{new} of new vertices and e_{new} of new edges satisfy $v_{\mathrm{new}} - e_{\mathrm{new}}/m_2(F) < -\delta$ where $\delta = \delta(F) > 0$, cf. (5). Therefore the set $\mathrm{Pre}^{d_{\max}}(V,E,O,\ell)$ can only be non-empty if

$$V - E/m_2(F) < -\delta d_{\max} \quad . \tag{16}$$

We now derive a bound on the number of sequences contained in a set $\mathrm{Pre}^{d_{\max}}(V,E,O,\ell)$. Similar to (12) we have at most

$$m(5\ell)^m \tag{17}$$

choices for the step configuration, i.e., the number of steps that are not open regular, their types and positions. Moreover, we again model the first step by choosing a starting edge (n^2 choices) and by counting the remaining choices similar to an open regular step.

For each of the O empty steps in a sequence we need to choose a copy of F within the *old* vertices, i.e., the vertices which have appeared previously in the sequence. As the entire sequence contains at most $v(F)\ell$ vertices we have at most

$$(v(F)\ell)^{v(F)O} \tag{18}$$

choices in total.

Considering all degenerate steps at once we need to choose a total of $d_{\max}v(F)$ vertices for them, V of which are new vertices and not from the previously seen ones, resulting in at most n^V choices, and $d_{\max}v(F) - V \le d_{\max}v(F)$ of which are chosen from the previously seen ones, giving at most $(v(F)\ell)^{d_{\max}v(F)}$ choices. Having fixed the new vertices and old vertices it remains to choose for every degenerate step which vertices of the copy of F are from the old and which from the new vertices. This gives at most another $2^{v(F)d_{\max}} \le (v(F)\ell)^{v(F)d_{\max}}$ choices. In total the number of choices for degenerate steps is bounded by

$$n^V (v(F)\ell)^{2v(F)d_{\max}} \quad . \tag{19}$$

All other $\ell - d_{\max} - O$ steps are regular steps. Similar to (14) each open regular step gives rise to at most $\big(v(F)^2 n^{v(F)-2}\big)$ choices. For the at most $d_{\max}$ closed regular steps we additionally have to select the two vertices that in contrast to open regular steps are not predetermined. This accounts at most for another $(v(F)\ell)^{2d_{\max}}$ choices. In total all regular steps give rise to at most

$$(v(F)\ell)^{2d_{\max}} \big(v(F)^2 n^{v(F)-2}\big)^{\ell - d_{\max} - O} \tag{20}$$

choices. Combining (17), (18), (19), (20) and the n^2 choices for the starting edge, we get that the number of sequences in the set $\mathrm{Pre}^{d_{\max}}(V, E, O, \ell)$ is at most

$$\begin{aligned} & m(5\ell)^m \cdot n^2 \cdot (v(F)\ell)^{v(F)O} \cdot n^V (v(F)\ell)^{2v(F)d_{\max}} \\ & \qquad \cdot (v(F)\ell)^{2d_{\max}} \big(v(F)^2 n^{v(F)-2}\big)^{\ell - d_{\max} - O} \\ & \quad \le \underbrace{m(5v(F)\ell)^{m+v(F)O+2v(F)d_{\max}+2d_{\max}}}_{=o(n)} \cdot \underbrace{v(F)^{2\ell}}_{\le n^{2\log(v(F))}} \cdot n^{2+V+(v(F)-2)(\ell-d_{\max}-O)} \\ & \quad \le n^{3+2\log(v(F))+V+(v(F)-2)(\ell-d_{\max}-O)} \quad , \end{aligned} \tag{21}$$

where the last step holds for n large enough. It is easy to see that a prefix sequence from the set $\mathrm{Pre}^{d_{\max}}(V, E, O, \ell)$ appears in $G(n, p)$ with respect to $\mathcal{R}(n, q)$ with probability

$$q^\ell \big(p^{e(F)-1}\big)^{\ell - d_{\max} - O} p^E \quad . \tag{22}$$

Combining (21) and (22) with $n^{v(F)-2}p^{e(F)-1}q \le c$ and $p \le n^{-1/m_2(F)} q^{-1/(e(F)-1)}$ which holds if we choose $c \le 1$ we obtain that for every subset $\mathrm{Pre}^{d_{\max}}(V, E, O, \ell)$ the number $X_{\mathrm{Pre}^{d_{\max}}(V,E,O,\ell)}$ of sequences that are present in $G(n, p)$ with respect to $\mathcal{R}(n, q)$ satisfies

$$\begin{aligned} \mathbf{E}[X_{\mathrm{Pre}^{d_{\max}}(V,E,O,\ell)}] & \le n^{3+2\log(v(F))+V+(v(F)-2)(\ell-d_{\max}-O)} \cdot q^\ell \big(p^{e(F)-1}\big)^{\ell-d_{\max}-O} p^E \\ & \le n^{3+2\log(v(F))+V} p^E q^{d_{\max}+O} \underbrace{c^{\ell-d_{\max}-O}}_{\le 1} \\ & \le n^{3+2\log(v(F))+V-E/m_2(F)} \underbrace{q^{-(E/(e(F)-1))+d_{\max}+O}}_{\le 1 \text{ since } E \le d_{\max}(e(F)-2)} \\ & \overset{(16)}{\le} n^{3+2\log(v(F))-\delta d_{\max}} \quad . \end{aligned}$$

Since this bound is independent of V, E, O and ℓ, we obtain with (15) that we have for large enough n that

$$\mathbf{E}[X_{\mathrm{Pre}^{d_{\max}}(\mathcal{S})}] \le \log(n)^2 n^{3+2\log(v(F))-\delta d_{\max}} \le n^{4+2\log(v(F))-\delta d_{\max}} \quad . \tag{23}$$

Choosing $d_{\max} = d_{\max}(F)$ large enough and $L = L(F)$ such that every canonical grow sequence with at least L steps has at least $d_{\max}$ degenerate steps (cf. Claim 11) we have that the expectation in (23) is $o(1)$. □

5. Concluding remarks

It is natural to ask whether there exists a generalization of Theorem 1 for a random k-uniform hypergraph $H^{(k)}(n,q)$. Indeed the 1-statement of the corresponding generalization was shown for the hypergraph of the tetrahedron $K_4^{(3)}$ by Rödl and Ruciński [15], and for k-partite k-uniform hypergraphs by Rödl, Ruciński and Schacht [16]. Quite recently, it has been further generalized to more k-uniform hypergraphs by Friedgut, Rödl and Schacht [6] and by Conlon and Gowers [2](in the case of strictly k-balanced k-uniform hypergraphs) independently. The result from [2, 6] is the following:

Theorem 12. *Let $r,k \geq 2$ be natural numbers and F be a k-uniform hypergraph with $\Delta(F) \geq 2$, then there exists a constant $C > 0$ such that if $q \geq Cn^{-1/m_k(F)}$ the following holds:*

$$\lim_{n\to\infty} \Pr(H^{(k)}(n,q) \longrightarrow (F)^e_r) = 1.$$

Here m_k denotes the so-called k-density for a k-uniform hypergraph F, which generalizes m_2-density in an obvious way:

$$m_k(F) := \max_{F'\subseteq F, e(F')\geq 2} d_k(F') \quad \text{where} \quad d_k(F') = \begin{cases} \frac{e(F')-1}{v(F')-k}, & e(F') \geq 2, \\ \frac{1}{k}, & e(F') = 1, \\ 0, & \text{otherwise.} \end{cases}$$

With the methods of the present paper and some additional ideas we can show the (expected) corresponding 0-statement of Theorem 12 for complete k-uniform hypergraphs. We are currently working on generalizing the arguments to a wider class of hypergraphs. This will appear elsewhere.

Acknowledgments. We would like to thank the anonymous reviewers for their helpful comments.

References

[1] P. Allen, J. Böttcher, J. Hladký, and D. Piguet. Turánnical hypergraphs. *Random Structures & Algorithms, to appear.*
[2] D. Conlon and W.T. Gowers. Combinatorial theorems in sparse random sets. *arXiv:1011.4310v1*, 2011.
[3] P. Erdős and A. Rényi. On the evolution of random graphs. *Publ. Math. Inst. Hung. Acad. Sci*, 5:17–61, 1960.
[4] E. Friedgut and M. Krivelevich. Sharp thresholds for certain Ramsey properties of random graphs. *Random Structures & Algorithms*, 17(1):1–19, 2000.
[5] E. Friedgut, V. Rödl, A. Ruciński, and P. Tetali. A sharp threshold for random graphs with a monochromatic triangle in every edge coloring. *Mem. Amer. Math. Soc.*, 179(845):vi+66, 2006.
[6] E. Friedgut, V. Rödl, and M. Schacht. Ramsey properties of random discrete structures. *Random Structures & Algorithms*, 37(4):407–436, 2010.
[7] R.L. Graham, B.L. Rothschild, and J.H. Spencer. *Ramsey theory.* Wiley-Interscience, 1990.
[8] S. Janson, T. Łuczak, and A. Rucinski. *Random graphs.* Wiley-Interscience Series in Discrete Mathematics and Optimization. Wiley-Interscience, New York, 2000.
[9] T. Łuczak, A. Ruciński, and B. Voigt. Ramsey properties of random graphs. *Journal of Combinatorial Theory, Series B*, 56(1):55 – 68, 1992.
[10] M. Marciniszyn, J. Skokan, R. Spöhel, and A. Steger. Threshold functions for asymmetric ramsey properties involving cliques. *Approximation, Randomization, and Combinatorial Optimization. Algorithms and Techniques*, pages 462–474, 2006.
[11] F.P. Ramsey. On a problem in formal logic. *Proc. Lond. Math. Soc.*, 30:264–286, 1930.
[12] V. Rödl and A. Ruciński. Lower bounds on probability thresholds for Ramsey properties. *Combinatorics, Paul Erdős is eighty*, 1:317–346, 1993.
[13] V. Rödl and A. Ruciński. Random graphs with monochromatic triangles in every edge coloring. *Random Structures & Algorithms*, 5(2):253–270, 1994.

[14] V. Rödl and A. Ruciński. Threshold functions for Ramsey properties. *J. Amer. Math. Soc.*, 8(4):917–942, 1995.

[15] V. Rödl and A. Ruciński. Ramsey properties of random hypergraphs. *J. Comb. Theory, Ser. A*, 81(1):1–33, 1998.

[16] V. Rödl, A. Ruciński, and M. Schacht. Ramsey properties of random k-partite, k-uniform hypergraphs. *SIAM J. Discrete Math.*, 21(2):442–460, 2007.

[17] P. Turán. Eine Extremalaufgabe aus der Graphentheorie. *Mat. Fiz. Lapok*, 48:436–452, 1941.