

UPPER TAILS FOR SUBGRAPH COUNTS IN RANDOM GRAPHS

BY

SVANTE JANSON

*Department of Mathematics, Uppsala University
PO Box 480, SE-751 06 Uppsala, Sweden
e-mail: svante.janson@math.uu.se
URL: <http://www.math.uu.se/~svante/>*

AND

KRZYSZTOF OLESZKIEWICZ*

*Institute of Mathematics, Warsaw University
Banacha 2, 02-097 Warsaw, Poland
e-mail: koles@mimuw.edu.pl*

AND

ANDRZEJ RUCIŃSKI**

*Department of Discrete Mathematics, Adam Mickiewicz University
Poznań, Poland
e-mail: rucinski@amu.edu.pl
URL: <http://main.amu.edu.pl/~rucinski/>*

ABSTRACT

Let G be a fixed graph and let X_G be the number of copies of G contained in the random graph $G(n, p)$. We prove exponential bounds on the upper tail of X_G which are best possible up to a logarithmic factor in the exponent. Our argument relies on an extension of Alon's result about the maximum number of copies of G in a graph with a given number of edges. Similar bounds are proved for the random graph $G(n, M)$ too.

* Research of the second author supported by KBN grant 2 P03A 027 22.

** Research of the third author supported by KBN grant 2 P03A 15 23.

Received November 4, 2002

1. Introduction

Let $G = (V(G), E(G))$ be a fixed graph. By a “copy” of G in another graph F we mean any, not necessarily induced subgraph of F , isomorphic with G . We use the notation $v_G = |V(G)|$ and $e_G = |E(G)|$ for the numbers of vertices and edges. (For typographic reasons we sometimes write $e(G)$ instead.) We assume that $e_G > 0$.

As usual, $G(n, p)$ denotes the random graph with n labelled vertices and the edges added randomly, such that each of the $\binom{n}{2}$ possible edges exists with probability p , independently of the other edges; see e.g. [11]. Let X_G be the number of copies of G contained in the random graph $G(n, p)$.

The distribution of X_G has been studied extensively since the pioneering paper by Erdős and Rényi [6], where the first results were given. A general threshold for $\{X_G > 0\}$ was established by Bollobás [3] at $p = n^{-1/m_G}$, where $m_G := \max_{H \subseteq G} e_H/v_H$, with X_G asymptotically Poisson at the threshold for the so-called strictly balanced graphs. The asymptotic normality of X_G for a wide range of p (as long as $pn^{1/m_G} \rightarrow \infty$ and $n^2(1-p) \rightarrow \infty$) was proven in [20].

Next, it was shown that the lower tail of the distribution of X_G decays exponentially in the expectation of the least expected subgraph of G , see [10] ($\mathbf{P}(X_G = 0)$) and [9] (the general case). Namely, let $\Psi_H := n^{v_H} p^{e_H}$, which is roughly the expected number of copies of H in $G(n, p)$. Then, for all $\epsilon \in (0, 1]$, with $c_\epsilon > 0$ depending on G and ϵ ,

$$(1.1) \quad \mathbf{P}(X_G \leq (1 - \epsilon)\mathbf{E}X_G) \leq \exp(-c_\epsilon \min_{H \subseteq G, e_H > 0} \Psi_H).$$

This is best possible, provided p stays away from 1, as by the FKG inequality, $-\log \mathbf{P}(X_G = 0) \leq -\log \mathbf{P}(X_H = 0) = O(\Psi_H)$ for every $H \subseteq G$. It is easy to see that $\min_{H \subseteq G} \Psi_H$ is achieved by $H = G$ for small p , and by $H = K_2$ for large p . For some graphs these are the only two cases, but there are graphs G for which as many as $\frac{2}{5}v_G$ subgraphs give the minimum for different ranges of p [21]. For more on small subgraphs of random graphs see [11].

Estimate (1.1) has found numerous applications within the theory of random graphs (see [11]). In the nineties, in the course of study of Ramsey properties of random graphs ([18, 19]) a need for a similar upper tail bound was articulated (see Appendix B for a typical application). So far only some ad hoc techniques have been used to obtain partial results. Spencer [22] gave bounds on the upper tail not only for X_G but also for the number of copies of G extending a given set of vertices. For balanced graphs and with some restrictions on the range of

p , Vu [24] derived an upper bound of the form:

$$\mathbf{P}(X_G \geq (1 + \epsilon)\mathbf{E}X_G) \leq \exp(-c_\epsilon \Psi_G^{1/(v_G-1)}).$$

In the same paper Vu also gave a general lower bound:

$$\mathbf{P}(X_G \geq (1 + \epsilon)\mathbf{E}X_G) \geq \exp(-C_\epsilon \Psi_G^{1/\alpha_G^*} \log(1/p)),$$

where α_G^* is the fractional independence number of G (see Appendix A for definition), and noticed that for stars $G = K_{1,k}$ it essentially matches the upper bound. He also asked under what circumstances the exponent in the upper bound could be improved to $-c_\epsilon \Psi_G^{1/\alpha_G^*}$, essentially matching the lower bound. (Actually, Vu’s lower bound is not always correct, because the construction in [24] may require more than n vertices; cf. the complete result for stars in Corollary 1.8 below. It is correct at least when p is so small that $\Psi_G^{1/\alpha_G^*} = O(n)$, and (for all p) when $\alpha_G^* = v_G/2$, for example, for regular graphs.)

Several techniques to bound the upper tail are surveyed in [12], see also [11]. Most of the methods do not give optimal results, but for the graphs K_4 and C_4 and some ranges of p , upper bounds matching Vu’s lower bound (up to a factor $O(\log(1/p))$ in the exponent) were found in [13] and [12]. Almost optimal results have also been obtained by Kim and Vu for K_3 [15] and K_4 (personal communication), and by Panchenko [17] for cycles. Some related results on the upper tail of the number of triangles in different models of random graphs are given in [5].

In this paper, we prove upper and lower exponential bounds on the upper tail of X_G for all G and all p , which match up to a factor $O(\log(1/p))$ in the exponent. Our main probabilistic result can be stated as follows.

Let $N(F, H)$ be the number of copies of H in another graph F and let $N(n, m, H)$ be the maximum of $N(F, H)$ over all graphs F with $v_F \leq n$ and $e_F \leq m$. In other words, $N(n, m, H)$ is the largest number of copies of H that can be packed in n vertices and m edges. Recall that $\Psi_H := n^{v_H} p^{e_H}$, and let

$$M_G^* = M_G^*(n, p) := \begin{cases} \max\{m \leq \binom{n}{2} : \forall H \subseteq G N(n, m, H) \leq \Psi_H\}, & p \geq n^{-2}, \\ 1, & p < n^{-2}. \end{cases} \tag{1.2}$$

The reason for the special definition in the extreme case $p < n^{-2}$ (when $\Psi_{K_2} = n^2 p < 1$) is to prevent $M_G^* = 0$. Since $N(n, 1, H) = 0$ unless $e_H \leq 1$, it is easily checked that $1 \leq M_G^* \leq \binom{n}{2}$.

Remark 1.1: It suffices to consider in (1.2) *connected* subgraphs $H \subseteq G$, because if H is the disjoint union of two subgraphs H_1 and H_2 , then $\Psi_H = \Psi_{H_1} \Psi_{H_2}$

while $N(F, H) \leq N(F, H_1)N(F, H_2)$ for each F and consequently $N(n, m, H) \leq N(n, m, H_1)N(n, m, H_2)$.

THEOREM 1.2: *For every graph G and for every $t > 1$ there exist constants $c(t, G) > 0$ and $C(t, G) > 0$ such that for all $n \geq v_G$ and $p \in (0, 1)$*

$$\mathbf{P}(X_G \geq t\mathbf{E}X_G) \leq \exp\{-c(t, G)M_G^*(n, p)\},$$

and, provided $t\mathbf{E}X_G \leq N(K_n, G)$,

$$\mathbf{P}(X_G \geq t\mathbf{E}X_G) \geq p^{C(t, G)M_G^*(n, p)}.$$

If $t\mathbf{E}X_G > N(K_n, G)$, the probability is trivially 0. Note that the condition $t\mathbf{E}X_G \leq N(K_n, G)$ is equivalent to $tp^{e_G} \leq 1$, so, for fixed t , we only have to consider p bounded away from 1.

Our proof is not strong enough to yield sharp estimates of the dependence of $c(t, G)$ and $C(t, G)$ on t , see Remark 8.2. We are mainly interested in the case of constant t , for example $t = 2$, and the theorem then says that $\log \mathbf{P}(X_G \geq t\mathbf{E}X_G)$ equals $-M_G^*$ up to a factor that is at most $O(\log 1/p)$. Theorem 1.2 is proved in Sections 2 (upper bound) and 3 (lower bound). It turns out that the same result holds with no essential modifications for the random graph $G(n, M)$ too, see Section 4.

To apply Theorem 1.2, it is crucial to have a fair estimate of the extremal parameter $N(n, m, H)$ for every graph $H \subseteq G$ in order to estimate M_G^* . This leads to the combinatorial part of this paper.

Let α_H^* be the *fractional independence number* of H (see Appendix A for the definition and basic properties of α_H^*). Noga Alon proved in his first published paper [1] (see also [7]) that $N(m, H) = \Theta(m^{\alpha_H^*})$, where $N(m, H) = N(\infty, m, H)$ is the analogue of $N(n, m, H)$ without restricting the number of vertices. Here and below, $a = \Theta(b)$ means that $cb \leq a \leq Cb$ for some positive constants c and C that may depend on the considered fixed graph (G or H) but not on any other parameters such as m , n or p . The next theorem, which is our main graph theoretic result, extends Alon's estimate.

THEOREM 1.3: *For every graph H without isolated vertices, and for all $m \geq e_H$ and $n \geq v_H$, we have*

$$(1.3) \quad N(n, m, H) = \begin{cases} \Theta(m^{\alpha_H^*}) & \text{if } m \leq n, \\ \Theta(m^{v_H - \alpha_H^*} n^{2\alpha_H^* - v_H}) & \text{if } n \leq m \leq \binom{n}{2}, \\ \Theta(n^{v_H}) & \text{if } m \geq \binom{n}{2}. \end{cases}$$

The first case above ($m \leq n$) is equivalent to Alon’s result, while the last case is trivial and stated here just for completeness. Theorem 1.3 is proved in Section 5.

Remark 1.4: As an immediate consequence of (and a supplement to) Theorem 1.3, if H has v_0 isolated vertices, then (1.3) remains valid for $m \geq n$, while for $m \leq n$ we have $N(n, m, H) = \Theta(m^{\alpha_H^* - v_0} n^{v_0})$.

The obtained bounds on $N(n, m, H)$ translate to the parameter $M_G^*(n, p)$, so vital for Theorem 1.2. Recall that $m_G := \max_{H \subseteq G} e_H/v_H$, and note that $p < n^{-1/m_G}$ if and only if $\min_{H \subseteq G} \Psi_H < 1$. Further, let Δ_G be the maximum degree of G . It is easily seen that $m_G \leq \Delta_G/2 < \Delta_G$.

THEOREM 1.5: *For every graph G and $n \geq v_G$ we have*

$$(1.4) \quad M_G^*(n, p) = \begin{cases} \Theta(1) & \text{if } p \leq n^{-1/m_G}, \\ \Theta(\min_{H \subseteq G} \Psi_H^{1/\alpha_H^*}) & \text{if } n^{-1/m_G} \leq p \leq n^{-1/\Delta_G}, \\ \Theta(n^2 p^{\Delta_G}) & \text{if } p \geq n^{-1/\Delta_G}. \end{cases}$$

Remark 1.6: Note that $\min_{H \subseteq G} \Psi_H^{1/\alpha_H^*}$ is achieved by a connected subgraph of G . Indeed, if $H = H_1 \cup H_2 \subseteq G$ is a disjoint union of two subgraphs of G and $(n^{v_H} p^{e_H})^{\alpha_{H_i}^*} < (n^{v_{H_i}} p^{e_{H_i}})^{\alpha_{H_i}^*}$ for both $i = 1, 2$, then we get a contradiction by multiplying the two inequalities. Hence $\Psi_H^{1/\alpha_H^*} \geq \Psi_{H_i}^{1/\alpha_{H_i}^*}$ for $i = 1$ or 2 .

Moreover, for $p \leq n^{-1/\Delta_G}$ as in (1.4), it suffices to consider H with $e_H > 0$, since otherwise $\Psi_H^{1/\alpha_H^*} = n \geq n^{1+1/\Delta_G} p = \Psi_K^{1/\alpha_K^*}$, where $K = K_{1, \Delta_G}$ is a maximal star in G .

Combining Theorems Theorems 1.2 and 1.5, we find upper and lower bounds for $\mathbf{P}(X_G \geq t \mathbf{E}X_G)$ expressed in terms of n and p . Note that they are similar to, but clearly distinct from, the lower tail estimates in (1.1), see Remark 8.3. We state explicit results for some particular classes of graphs. We consider below only $p \geq n^{-1/m_G}$, since otherwise $M_G^* = \Theta(1)$ by Theorem 1.5. First, for k -regular graphs the formula for M_G^* is indeed very simple.

COROLLARY 1.7: *If G is a k -regular graph, then $M_G^* = \Theta(n^2 p^k)$ for all $p \geq n^{-1/m_G} = n^{-2/k}$. More generally, for a graph G we have $M_G^* = \Theta(n^2 p^k)$ for all $p \geq n^{-1/m_G}$ if and only if $\Delta_G = k$ and G has a k -regular subgraph. In particular, if G is connected, then $M_G^* = \Theta(n^2 p^k)$ for all $p \geq n^{-1/m_G}$ if and only if G is k -regular.*

However, for other graphs, we may have different expressions for M_G^* for different ranges of p . Below we give two examples of that.

COROLLARY 1.8: Let G be the k -armed star $K_{1,k}$, with $k \geq 1$, and assume $p \geq n^{-1/m_G} = n^{-1-1/k}$. Then

$$(1.5) \quad M_G^* = \begin{cases} \Theta(n^{1+1/k}p) & \text{if } p \leq n^{-1/k}, \\ \Theta(n^2p^k) & \text{if } p \geq n^{-1/k}. \end{cases}$$

COROLLARY 1.9: Let P_k be the path on k vertices, i.e. of length $k - 1$, and assume $p \geq n^{-1/m_{P_k}} = n^{-1-1/(k-1)}$. Then, if $k \geq 3$ is odd,

$$M_{P_k}^* = \begin{cases} \Theta(n^{2\frac{k}{k+1}}p^{2\frac{k-1}{k+1}}) & \text{if } p \leq n^{-1/2}, \\ \Theta(n^2p^2) & \text{if } p \geq n^{-1/2}, \end{cases}$$

and, if $k \geq 4$ is even,

$$M_{P_k}^* = \begin{cases} \Theta(n^2p^{2\frac{k-1}{k}}) & \text{if } p \leq n^{-1}, \\ \Theta(n^{2\frac{k-1}{k}}p^{2\frac{k-2}{k}}) & \text{if } n^{-1} \leq p \leq n^{-1/2}, \\ \Theta(n^2p^2) & \text{if } p \geq n^{-1/2}. \end{cases}$$

Note that $P_2 = K_2$ is covered by Corollary 1.7 or 1.8, and quite different. Theorem 1.5 and its corollaries are proved in Section 6.

It is easily seen that for any G we have a pattern of the same type as in the corollaries above; $M_G^*(n, p)$ is always given by a sequence of one or several expressions of the type $\Theta(n^a p^b)$, each expression valid for a certain range of p . This is further explained in Section 7.

ACKNOWLEDGEMENT: This research was initiated after the third author gave an invited talk at the VII-th Conference on Probability in Będlewo in May 2002 to an audience including the second author, and then carried on through the summer in pairwise collaborations made possible by the Pacific Institute of Mathematical Sciences, Vancouver (Thematic Programme on Asymptotic Geometric Analysis), and the Isaac Newton Institute for Mathematical Sciences in Cambridge, U.K. (Programme on Computing, Probability and Combinatorics). We would like to thank the organizers of these meetings and programs for inviting us there, as well as the sponsoring organizations for financial support.

We are also very grateful to Van Vu for sending us his most recent manuscripts and to Noga Alon, Laci Lovász and Bruce Reed for providing us with some relevant reference information.

Finally, we acknowledge that the construction in Example 7.8 was found after massive experimentation with Maple.

2. Proof of the upper bound

It will be convenient in the proof to count *ordered* copies of G , i.e. copies of G with the vertices ordered in some way, consistent with a fixed ordering of the vertices of G . Equivalently, an ordered copy of G in F can be regarded as an injective mapping of $V(G)$ into $V(F)$ that maps edges into edges. We let Y_G denote the total number of ordered copies of G in $G(n, p)$. Note that $Y_G = \text{aut}(G)X_G$, where $\text{aut}(G)$ is the number of automorphisms of G , and thus $\mathbf{P}(Y_G \geq t\mathbf{E}Y_G) = \mathbf{P}(X_G \geq t\mathbf{E}X_G)$.

For convenience, we also extend the definition of $N(n, m, H)$ to non-integer m ; clearly $N(n, m, H) = N(n, \lfloor m \rfloor, H)$. We will use a simple estimate.

LEMMA 2.1: *For every H with $e_H > 0$ there is a constant C_H such that if $n \geq v_H$ and $0 \leq m_1 < m_2 \leq \binom{n}{2}$, then*

$$N(n, m_1, H) \leq C_H \frac{m_1}{m_2} N(n, m_2, H).$$

Proof: This can be shown directly from the definition, but for us it is simpler to see that it is an immediate corollary of Theorem 1.3 and the obvious inequalities $1 \leq \alpha_H^* \leq v_H - 1$. (The case $m_1 < e_H$ is trivial, since then $N(n, m_1, H) = 0$.)

Let $G_1, \dots, G_f, f = \binom{n}{v_G}$, be all ordered copies of G in the complete graph K_n on the labelled vertex set $[n] = \{1, \dots, n\}$. Furthermore, for $i = 1, \dots, f$, set $Y_i = 1$ if $G_i \subset G(n, p)$ and $Y_i = 0$ otherwise. Thus $Y_G = \sum_i Y_i$.

Set $\mu_H := \mathbf{E}Y_H = \binom{n}{v_H} p^{e_H} = \Theta(\Psi_H)$. We will first show by induction that for all $m = 1, 2, \dots$

$$(2.1) \quad \mathbf{E}Y_G^m \leq \mu_G^m \left(1 + v_G! \sum_{H \subseteq G} \frac{N(n, (m-1)e_G, H)}{\mu_H} \right)^{m-1}.$$

The above sum extends over all subgraphs H of G with $e_H > 0$. It is trivially true for $m = 1$.

For $m \geq 2$, we set $F = F(i_1, \dots, i_{m-1}) = G_{i_1} \cup \dots \cup G_{i_{m-1}}$. The number of copies of G which are edge-disjoint from F is bounded, crudely, by $\binom{n}{v_G}$, while the number of copies of G which intersect F on a given subgraph H can be bounded as follows: there are at most $N(n, e_F, H) \leq N(n, (m-1)e_G, H)$ choices of H in F , at most $\binom{n-v_H}{v_G-v_H} = \binom{n}{v_G} / \binom{n}{v_H}$ choices of the remaining

vertices in G and, finally, at most $v_G!$ orderings. Hence

$$\begin{aligned} \mathbf{E}Y_G^m &= \sum_{i_1, \dots, i_m} \mathbf{E}(Y_{i_1} \cdots Y_{i_m}) = \sum_{i_1, \dots, i_m} p^{e(G_{i_1} \cup \dots \cup G_{i_m})} \\ &= \sum_{i_1, \dots, i_{m-1}} p^{e(F)} \sum_{i_m} p^{e_G - e(F \cap G_{i_m})} \\ &\leq \sum_{i_1, \dots, i_{m-1}} p^{e(F)} \left((n)_{v_G} p^{e_G} + \sum_{H \subseteq G} \sum_{G_i \cap F \cong H} p^{e_G - e_H} \right) \\ &\leq \sum_{i_1, \dots, i_{m-1}} p^{e(F)} \left(\mu_G + v_G! \sum_{H \subseteq G} N(n, (m-1)e_G, H) \frac{\mu_G}{\mu_H} \right) \\ &= \mathbf{E}Y_G^{m-1} \cdot \mu_G \left(1 + v_G! \sum_{H \subseteq G} \frac{N(n, (m-1)e_G, H)}{\mu_H} \right), \end{aligned}$$

and (2.1) follows by induction.

By Markov's inequality, we then have

$$(2.2) \quad \mathbf{P}(Y_G \geq t\mu_G) \leq t^{-m} \left(1 + v_G! \sum_{H \subseteq G} \frac{N(n, (m-1)e_G, H)}{\mu_H} \right)^{m-1}.$$

For $p \geq n^{-2}$, choose $m = \lceil c' M_G^* \rceil$, with $c' = c'(G, t) < 1/e_G$ chosen below. We then have by Lemma 2.1 and the inequality $N(n, M_G^*, H) \leq \Psi_H$,

$$\begin{aligned} \frac{N(n, (m-1)e_G, H)}{\mu_H} &\leq \frac{N(n, c' M_G^* e_G, H)}{\mu_H} \leq C_H c' e_G \frac{N(n, M_G^*, H)}{\mu_H} \\ &\leq C_H c' e_G \frac{\Psi_H}{\mu_H} = C_H c' e_G \frac{n^{v_H}}{(n)_{v_H}} \\ &\leq C_H c' e_G \frac{n^{v_G}}{(n)_{v_G}} \leq C_H c' e_G \frac{v_G^{v_G}}{v_G!}. \end{aligned}$$

Hence, if $c' = (e_G v_G^{v_G} \sum_{H \subseteq G} C_H)^{-1} \min(t^{1/2} - 1, 1)$, then

$$1 + v_G! \sum_{H \subseteq G} \frac{N(n, c' e_G M_G^*, H)}{\mu_H} \leq t^{1/2},$$

and (2.2) yields the upper bound in Theorem 1.2, with $c(t, G) = c'^{\frac{1}{2}} \log t$. The exceptional case $p < n^{-2}$ is trivial with $m = 1$. \blacksquare

3. Proof of the lower bound

The lower bound in Theorem 1.2 will be a consequence of the following result. As in the proof of the upper bound, Y_H stands for the number of ordered copies of H in $G(n, p)$.

THEOREM 3.1: *Suppose that $0 < p < 1$ and $n \geq v_G$. If there exist $H \subseteq G$, $t > 0$ and m such that*

$$N(n, m, H) \geq 2t\Psi_H,$$

then

$$\mathbf{P}(Y_G \geq t\mathbf{E}Y_G) \geq \frac{1}{4}p^{m+\epsilon_G}.$$

To prove Theorem 3.1 we will use a lower tail estimate for a sum of indicator random variables. It is a consequence of the following Paley–Zygmund type inequality (see, e.g., [14, p. 8]) which generalizes a strengthening of Chebyshev’s inequality (see, e.g., [11, page 54, Remark 3.1].) We include a simple proof for completeness.

LEMMA 3.2: *Let X be a random variable with $\mathbf{E}X > 0$ and $\mathbf{E}X^2 < \infty$. Then, for every $0 \leq \delta < 1$,*

$$\mathbf{P}(X > \delta\mathbf{E}X) \geq (1 - \delta)^2 \frac{(\mathbf{E}X)^2}{\mathbf{E}X^2}.$$

Proof: Let $\mathbf{1}_\mathcal{E}$ denote the indicator of an event \mathcal{E} . We have $X\mathbf{1}_{\{X \leq \delta\mathbf{E}X\}} \leq \delta\mathbf{E}X$, and thus

$$\mathbf{E}(X\mathbf{1}_{\{X > \delta\mathbf{E}X\}}) = \mathbf{E}X - \mathbf{E}(X\mathbf{1}_{\{X \leq \delta\mathbf{E}X\}}) \geq \mathbf{E}X - \delta\mathbf{E}X.$$

The Cauchy–Schwarz inequality now gives

$$((1 - \delta)\mathbf{E}X)^2 \leq (\mathbf{E}(X\mathbf{1}_{\{X > \delta\mathbf{E}X\}}))^2 \leq \mathbf{E}X^2 \mathbf{P}(X > \delta\mathbf{E}X). \quad \blacksquare$$

LEMMA 3.3: *If $X = \sum_{i=1}^N X_i$, where X_i are arbitrary indicator random variables with $\mathbf{E}X_i = \mathbf{P}(X_i = 1) = p_i \geq p_0$ for some $p_0 \geq 0$, then*

$$\mathbf{P}(X > \frac{1}{2}Np_0) \geq \mathbf{P}(X > \frac{1}{2}\mathbf{E}X) \geq \frac{1}{4}p_0.$$

Proof: Clearly, $\mathbf{E}X = \sum_{i=1}^N p_i \geq Np_0$. Moreover, by the Cauchy–Schwarz inequality, $X^2 \leq N \sum_{i=1}^N X_i^2$, and thus

$$\mathbf{E}X^2 \leq N \sum_{i=1}^N \mathbf{E}X_i^2 = N \sum_{i=1}^N \mathbf{E}X_i = N\mathbf{E}X.$$

Hence, by Lemma 3.2,

$$\mathbf{P}(X > \frac{1}{2}Np_0) \geq \mathbf{P}(X > \frac{1}{2}\mathbf{E}X) \geq \frac{(\mathbf{E}X)^2}{4\mathbf{E}X^2} \geq \frac{\mathbf{E}X}{4N} \geq \frac{1}{4}p_0. \quad \blacksquare$$

Proof of Theorem 3.1: By assumption, there exists a graph $F \subseteq K_n$ with $e_F \leq m$ and such that

$$N(F, H) \geq 2t\Psi_H \geq 2t\mathbf{E}Y_H.$$

We fix one such subgraph F of K_n . Note that F contains many more copies of H than $G(n, p)$ does on average. Thus our strategy will be to guarantee that $G(n, p) \supseteq F$, and, using the lower tail estimate of Lemma 3.3, that a fraction of the expected number of extensions of these copies of H to copies of G will indeed appear in $G(n, p)$.

Fix one copy of H in G and ignore the other copies of H , if there are any. Then each ordered copy of G in another graph contains a unique corresponding ordered copy of H . Call a copy of G in K_n *F-rooted*, if this distinguished copy of H is a subgraph of F . Note that there are precisely $\text{aut}(H)N(F, H)$ ordered copies of H in F , and thus

$$N := \text{aut}(H)N(F, H)(n - v_H)_{v_G - v_H}$$

F -rooted copies of G in K_n ; we denote them by G_1, \dots, G_N . Note also that

$$(3.1) \quad Np^{e_G - e_H} \geq 2t\mathbf{E}Y_H(n - v_H)_{v_G - v_H}p^{e_G - e_H} = 2t\mathbf{E}Y_G.$$

Let $G_F(n, p)$ be the random graph $G(n, p)$ conditioned on $G(n, p) \supseteq F$, i.e. the random graph obtained by adding to F each of the remaining $\binom{n}{2} - e_F$ edges with probability p , independently of the others. Let Z_i be the indicator that $G_F(n, p)$ contains G_i . Then

$$\mathbf{P}(Z_i = 1) = p^{e(G_i \setminus F)} \geq p^{e_G - e_H},$$

and thus, by Lemma 3.3,

$$(3.2) \quad \begin{aligned} \mathbf{P}(Y_G \geq \frac{1}{2}Np^{e_G - e_H} \mid G(n, p) \supseteq F) &\geq \mathbf{P}\left(\sum_{i=1}^N Z_i \geq \frac{1}{2}Np^{e_G - e_H}\right) \\ &\geq \frac{1}{4}p^{e_G - e_H} \geq \frac{1}{4}p^{e_G}. \end{aligned}$$

Consequently, by (3.1) and (3.2),

$$\begin{aligned} \mathbf{P}(Y_G \geq t\mathbf{E}Y_G) &\geq \mathbf{P}(Y_G \geq \frac{1}{2}Np^{e_G - e_H}) \geq \frac{1}{4}p^{e_G}\mathbf{P}(G(n, p) \supseteq F) \\ &= \frac{1}{4}p^{e_G}p^{e_F} \geq \frac{1}{4}p^{e_G + m}. \quad \blacksquare \end{aligned}$$

To complete the proof of the lower bound in Theorem 1.2, note first that, by the definition (1.2) of M_G^* , there exists $H \subseteq G$ such that $N(n, M_G^* + 1, H) > \Psi_H$, except when $M_G^* = \binom{n}{2}$. It follows from Lemma 2.1 that with $C' = 4 \max_{H \subseteq G} C_H$, either

$$N(n, C'tM_G^*, H) \geq 2tN(n, 2M_G^*, H) \geq 2tN(n, M_G^* + 1, H) > 2t\Psi_H$$

or $C'tM_G^* > \binom{n}{2}$. In the second case, the lower bound is a consequence of the inequality

$$(3.3) \quad \mathbf{P}(X_G \geq t\mathbf{E}X_G) \geq \mathbf{P}(X_G = N(K_n, G)) \geq \mathbf{P}(G(n, p) = K_n) = p^{\binom{n}{2}}$$

and thus

$$\mathbf{P}(X_G \geq t\mathbf{E}X_G) \geq p^{C'tM_G^*}.$$

In the first case, Theorem 3.1 yields

$$\mathbf{P}(X_G \geq t\mathbf{E}X_G) \geq \frac{1}{4}p^{C'tM_G^* + e_G} \geq \frac{1}{4}p^{C''tM_G^*}$$

with $C'' = C' + e_G$, and it remains only to take care of the factor $1/4$. If $p \leq 1/2$, we can use $1/4 \geq p^2 \geq p^{2M_G^*}$ and the result follows.

For $p > 1/2$ we note that Lemma 2.1 implies that if $c' = 1/\max_{H \subseteq G} C_H$, then for every $H \subseteq G$

$$N(n, c' \binom{n}{2} p^{e_G}, H) \leq p^{e_G} N(n, \binom{n}{2}, H) = p^{e_G} N(K_n, H) \leq \Psi_H$$

and thus $M_G^* \geq c' \binom{n}{2} p^{e_G} \geq 2^{-e_G} c' \binom{n}{2}$. Consequently, the result follows again from (3.3). ■

4. Subgraph counts in $G(n, M)$

In this section we consider the random graph $G(n, M)$ with n (labelled) vertices and M edges, chosen uniformly at random from all sets of M edges in K_n . (Here n and M are integers with $0 \leq M \leq \binom{n}{2}$.) Thus, X_G now denotes the random variable which is the number of copies of G in $G(n, M)$. Note that the maximum possible value of X_G is $N(n, M, G)$.

It is well known that in many respects, $G(n, M)$ and $G(n, p)$ with $p = M/\binom{n}{2}$ behave similarly, see [11, Section 1.4]. For example, this is true for the property $X_G > 0$ of containing a copy of G . Under some mild restrictions, the lower tail estimates carry over from $G(n, p)$ to $G(n, M)$, see [11, Section 3.1]. It is not obvious that this correspondence holds also for the upper tails, but indeed it does. We have the following version of Theorem 1.2 for $G(n, M)$.

THEOREM 4.1: *For every graph G and for every $t > 1$ there exist constants $c(t, G) > 0$ and $C(t, G) > 0$ such that for all $n \geq v_G$ and $e_G \leq M \leq \binom{n}{2}$, with $p := M/\binom{n}{2}$,*

$$\mathbf{P}(X_G \geq t\mathbf{E}X_G) \leq \exp\{-c(t, G)M_G^*(n, p)\},$$

and, provided $t\mathbf{E}X_G \leq N(n, M, G)$,

$$\mathbf{P}(X_G \geq t\mathbf{E}X_G) \geq p^{C(t, G)M_G^*(n, p)}.$$

Proof: Note first that, taking $H = K_2$ in (1.2),

$$(4.1) \quad M_G^* = N(n, M_G^*, K_2) \leq \Psi_{K_2} = n^2 p = O(M).$$

For the upper bound we follow the proof in Section 2, now taking expectations for $G(n, M)$. Let again $\mu_H := \binom{n}{v_H} p^{e_H}$ but observe that this now is only an approximation to $\mathbf{E}Y_H$. We have

$$\mathbf{E}Y_G = \binom{n}{v_G} \frac{\binom{M}{e_G}}{\binom{\binom{n}{2}}{e_G}}$$

and thus, by simple estimates,

$$(4.2) \quad \mu_G \left(1 - \binom{e_G}{2} \frac{1}{M}\right) \leq \mathbf{E}Y_G \leq \mu_G.$$

Similarly we see that $\mathbf{E}Y_G^m$ is not larger for $G(n, M)$ than for $G(n, p)$. Hence (2.1) still holds. By Markov's inequality we now have

$$\mathbf{P}(Y_G \geq t\mathbf{E}Y_G) \leq t^{-m} \frac{\mathbf{E}Y_G^m}{\mu_G^m} \left(\frac{\mu_G}{\mathbf{E}Y_G}\right)^m,$$

yielding the same estimate as in (2.2), except for a factor $(\mu_G/\mathbf{E}Y_G)^m$.

If $\mu_G/\mathbf{E}Y_G \leq t^{1/4}$, say, this completes the proof of the upper bound as in Section 2.

If $t^{1/4} < \mu_G/\mathbf{E}Y_G$, (4.2) implies $t = 1 + O(1/M)$ and thus by (4.1)

$$(t - 1)^2 M_G^* = O((t - 1)^2 M) = O(t - 1) = O(\log t).$$

Taking $c(t, G) = c_1(t - 1)^2$ with c_1 small, we thus have $c(t, G)M_G^* \leq \log t$ and our bound follows by Markov's inequality $\mathbf{P}(X_G \geq t\mathbf{E}X_G) \leq 1/t$.

For the lower bound, we note first that if we choose a subgraph $F \subseteq K_n$ with $e_F = M$ and $N(F, G) = N(n, M, G)$, then, using Stirling's formula,

(4.3)

$$\mathbf{P}(X_G = N(n, M, G)) \geq \mathbf{P}(G(n, M) \supseteq F) = \binom{\binom{n}{2}}{M}^{-1} \geq \frac{(M/e)^M}{\binom{n}{2}^M} = \left(\frac{p}{e}\right)^M.$$

In the case $p \geq 1/2$ we have $M_G^* = \Theta(n^2) = \Theta(M)$ by Theorem 1.5, and the lower bound follows by (4.3), if we note that the assumption $t\mathbf{E}X_G \leq N(n, M, G) \leq N(K_n, G)$ implies that $p \leq c'(t, G)$ for some $c'(t, G) < 1$.

Assume now $p < 1/2$. By Lemma 2.1 and the argument in Section 3, there exists a subgraph $H \subseteq G$ and a constant $C' = 2^{e_G+3}C_H$ such that either

$$(4.4) \quad N(n, C'tM_G^*, H) > 2^{e(G)+2}t\Psi_H \geq 2^{e(G)+2}t\mathbf{E}Y_H$$

or $C'tM_G^* > \binom{n}{2}$. In the second case, and more generally if $C'tM_G^* \geq M/2$, the lower bound follows again from (4.3).

In the remaining case, (4.4) holds and $C'tM_G^* < M/2$. Now fix a graph $F \subseteq K_n$ with $e_F \leq C'tM_G^* < M/2$ such that

$$(4.5) \quad N(F, H) \geq 2^{e(G)+2}t\mathbf{E}Y_H.$$

We condition on $F \subseteq G(n, M)$, which means that we consider the random graph $G_F(n, M)$ obtained by adding $M - e_F \geq M/2$ random edges to F . We count F -rooted copies of G as in the proof of Theorem 3.1 in Section 3, now with

$$\mathbf{P}(Z_i = 1) \geq (1 + O(M^{-1}))(p/2)^{e(G_i \setminus F)} \geq 2^{-e_G-1}p^{e_G-e_H},$$

if M is large enough. Denoting the number of F -rooted copies of G in K_n by N as before, we now have by (4.5), instead of (3.1), $Np^{e_G-e_H} \geq 2^{e_G+2}t\mathbf{E}Y_G$. Consequently, just as in Section 3,

$$\frac{1}{2}\mathbf{E} \sum_{i=1}^N Z_i \geq \frac{1}{2}N2^{-e_G-1}p^{e_G-e_H} \geq t\mathbf{E}Y_G,$$

and the result again follows by Lemma 3.3, since

$$\mathbf{P}(G(n, M) \supset F) \geq (p/2)^{e_F} \geq (p/2)^{M/2}. \quad \blacksquare$$

5. Proof of Theorem 1.3

We follow [7] with minor modifications. Consider the linear programming problem of finding $x_v, v \in V(H)$, that maximize $\sum_v x_v$ subject to constraints

$$(5.1) \quad 0 \leq x_v \leq \log n, \quad v \in V(H),$$

and

$$(5.2) \quad x_v + x_w \leq \log m, \quad vw \in E(H).$$

Let $\gamma = \gamma(n, m, H)$ be the maximum value of $\sum_{v \in V} x_v$.

Given an optimal solution $x_v, v \in V(H)$, we construct a graph F by blowing up each vertex $v \in V(H)$ to a set of $n_v = \lceil ce^{x_v} \rceil$ vertices, where c is a small constant depending on H only ($c = v_H^{-2}$ will do). In other words, let $V(F) = \bigcup_{v \in V(H)} V_v$, where V_v are disjoint sets, $|V_v| = n_v$, and the pair (V_v, V_w) spans a complete bipartite graph if $vw \in E(H)$ and the empty graph otherwise. It is clear by the above constraints that $v_F \leq n, e_F \leq m$, while

$$N(F, H) \geq \prod_v n_v \geq c^{v_H} e^\gamma.$$

Conversely, consider the dual linear programming problem which is to find $y_v, v \in V(H)$, and $z_e, e \in E(H)$, that minimize

$$\sum_v y_v \log n + \sum_e z_e \log m$$

under the constraints

$$(5.3) \quad y_v, z_e \geq 0, \quad v \in V(H), e \in E(H),$$

and

$$(5.4) \quad y_v + \sum_{e \ni v} z_e \geq 1, \quad v \in V(H).$$

Let (y_v) and (z_e) denote an optimal solution. Then, by the duality theorem (see, e.g., [16]),

$$(5.5) \quad \sum_v y_v \log n + \sum_e z_e \log m = \gamma.$$

Let F be a graph with $v_F \leq n$ and $e_F \leq m$. For any partition $V = V(F) = \bigcup_{v \in V(H)} V_v$, let \mathcal{W} be the set of all ordered copies of H in F with each $v \in V(H)$ mapped to a vertex in the corresponding set V_v . An elementary application of the probabilistic method (see, e.g., [7, Proposition 1.3]) yields that there is always a partition with

$$(5.6) \quad |\mathcal{W}| \geq v_H^{-v_H} N(F, H).$$

Let us fix one such partition. Define a hypergraph \mathcal{W}' on V by

$$\mathcal{W}' = \{V(H') : H' \in \mathcal{W}\}.$$

Since each $H' \in \mathcal{W}$ is determined by $V(H')$, we have $|\mathcal{W}'| = |\mathcal{W}|$.

For a subset $U \subseteq V$ define the trace of \mathcal{W}' on U by

$$\text{Tr}(\mathcal{W}', U) = \{W \cap U : W \in \mathcal{W}'\} = \{V(H') \cap U : H' \in \mathcal{W}'\}.$$

Let t be a large integer and let $s_v = \lceil y_v t \rceil$, $u_e = \lceil z_e t \rceil$, where we recall that (y_v) and (z_e) is a solution satisfying (5.5). We form a sequence U_1, \dots, U_s of subsets of V by taking each V_v precisely s_v many times and each $V_v \cup V_w$ precisely u_{vw} many times. Note that, by (5.5), each vertex $v \in V_v$ is contained in $s_v + \sum_{e \ni v} u_e \geq y_v t + \sum_{e \ni v} z_e t \geq t$ of these sets.

By Shearer's Lemma ([4]; cf. [7, Lemma 1.2]),

$$\begin{aligned} |\mathcal{W}|^t &= |\mathcal{W}'|^t \leq \prod_{j=1}^s |\text{Tr}(\mathcal{W}', U_j)| \\ (5.7) \qquad &= \prod_{v \in V(H)} |\text{Tr}(\mathcal{W}', V_v)|^{s_v} \prod_{vw \in E(H)} |\text{Tr}(\mathcal{W}', V_v \cup V_w)|^{u_{vw}}. \end{aligned}$$

Now, as $\text{Tr}(\mathcal{W}', V_v)$ is a set of singletons, $|\text{Tr}(\mathcal{W}', V_v)| \leq |V_v| \leq v_F \leq n$, while $\text{Tr}(\mathcal{W}', V_v \cup V_w)$ is a set of edges of F , so $|\text{Tr}(\mathcal{W}', V_v \cup V_w)| \leq e_F \leq m$. Hence, by (5.7),

$$|\mathcal{W}| \leq \prod_v n^{s_v/t} \prod_{vw} m^{u_{vw}/t} \leq \prod_v n^{y_v+1/t} \prod_{vw} m^{z_{vw}+1/t} \rightarrow e^\gamma$$

as $t \rightarrow \infty$. Finally, by (5.6),

$$(5.8) \qquad N(F, H) \leq v_H^{y_H} e^\gamma,$$

which implies that $N(n, m, H) \leq v_H^{y_H} e^\gamma$. This, together with the previously established lower bound, yields that

$$N(n, m, H) = \Theta(e^\gamma).$$

It remains to calculate γ , for which we have to further study the system of constraints (5.1) and (5.2). In the case $m \leq n$, after rescaling by $\log m$, it is straightforward to see that $\gamma = \alpha_H^* \log m$. In the trivial case $m \geq n^2$, the optimal assignment is $x_v = \log n$ for all v , and so $\gamma = v_H \log n$.

In the remaining case $n \leq m < n^2$, any optimal assignment must satisfy $x_v \geq \log m - \log n$ for all v , because otherwise we could increase an x_v to $\log m - \log n$ without violating any constraints. Consequently, we may write

$$x_v = \log m - \log n + (2 \log n - \log m) \xi_v,$$

where $0 \leq \xi_v \leq 1$. Then the conditions (5.1) and (5.2) become, in terms of ξ_v , the standard conditions defining the fractional independence number α_H^* . Hence, in this case,

$$\begin{aligned} \gamma &= \sum_v x_v = (\log m - \log n)v_H + (2 \log n - \log m) \sum_v \xi_v \\ &= (\log m - \log n)v_H + (2 \log n - \log m)\alpha_H^*. \quad \blacksquare \end{aligned}$$

6. Remaining proofs

In this section we prove Theorem 1.5 and its corollaries. We consider only subgraphs $H \subseteq G$ without isolated vertices.

LEMMA 6.1: *For every subgraph $H \subseteq G$,*

$$(6.1) \quad e_H \leq \Delta_G(v_H - \alpha_H^*),$$

and equality holds for at least one H with $e_H > 0$.

Proof: The inequality follows immediately from Lemma A.1 and the obvious fact that $\Delta_H \leq \Delta_G$. To obtain equality, let K be a maximal star in G , i.e. a vertex v with maximal degree and its Δ_G neighbours together with the Δ_G edges from v . Then $e_K = \Delta_G$, $v_K = \Delta_G + 1$ and $\alpha_K^* = \Delta_G$. \blacksquare

Proof of Theorem 1.5: First, if $p < n^{-1/m_G}$, then $\Psi_H < 1$ for some $H \subseteq G$. Since $N(n, e_H, H) \geq N(H, H) = 1 > \Psi_H$, (1.2) yields $M_G^* < e_H \leq e_G$ (except when $e_H = 1$ and $M_G^* = 1$).

Next, assume that $n^{-1/m_G} \leq p \leq n^{-1/\Delta_G}$ and let $m = \min_{H \subseteq G} \Psi_H^{1/\alpha_H^*}$. Taking $H \subseteq G$ which yields equality in (6.1), we have $e_H/\Delta_G = v_H - \alpha_H^*$ and thus

$$\Psi_H = n^{v_H} p^{e_H} \leq n^{v_H - e_H/\Delta_G} = n^{\alpha_H^*}.$$

Hence, $m \leq \Psi_H^{1/\alpha_H^*} \leq n$. If $m \geq e_G$, Theorem 1.3 thus yields that, for every $H \subseteq G$,

$$(6.2) \quad N(n, m, H) = \Theta(m^{\alpha_H^*}) = O(\Psi_H)$$

and for some $H_0 \subseteq G$ (the one achieving the minimum in the definition of m)

$$(6.3) \quad N(n, m, H_0) = \Theta(m^{\alpha_{H_0}^*}) = \Theta(\Psi_{H_0}).$$

It is now easy to see, by Lemma 2.1, that, for suitable constants c and C , $N(n, cm, H) \leq \Psi_H$ for all $H \subseteq G$, while either $N(n, Cm, H_0) > \Psi_{H_0}$ or $Cm >$

$\binom{n}{2}$). Hence, $cm \leq M_G^* < Cm$. In the case $m < e_G$, we similarly find $M_G^* < Ce_G$ and thus $M_G^* = \Theta(1) = \Theta(m)$.

Finally, assume that $p \geq n^{-1/\Delta_G}$, and let now $m = n^2 p^{\Delta_G}$; thus $m \geq n$. If $m \geq e_G$, Theorem 1.3 now yields, for every $H \subseteq G$,

$$N(n, m, H) = \Theta(m^{v_H - \alpha_H^*} n^{2\alpha_H^* - v_H}) = \Theta(n^{v_H} p^{\Delta_G(v_H - \alpha_H^*)}).$$

By Lemma 6.1, we have $\Delta_G(v_H - \alpha_H^*) \geq e_H$ for every $H \subseteq G$ with equality for some H_0 (at least for $H_0 = K$ — the maximal star in G). Hence, we again have (6.2) for every H and (6.3) for some H_0 (ignoring the middle terms), and $M_G^* = \Theta(m)$ follows as before. Finally, in the case $m < e_G$, we easily find $M_G^* = \Theta(1) = \Theta(m)$. ■

LEMMA 6.2: If $H \subseteq G$, and $p \leq n^{-1/\Delta_G}$, then $\Psi_H^{1/\alpha_H^*} \geq n^2 p^{\Delta_G}$.

Proof: We have $e_H \leq \Delta_G v_H / 2 \leq \Delta_G \alpha_H^*$ and thus

$$(6.4) \quad \begin{aligned} \Psi_H / (n^2 p^{\Delta_G})^{\alpha_H^*} &= n^{v_H - 2\alpha_H^*} p^{e_H - \Delta_G \alpha_H^*} \\ &\geq n^{v_H - 2\alpha_H^*} (n^{-1/\Delta_G})^{e_H - \Delta_G \alpha_H^*} = n^{v_H - \alpha_H^* - e_H / \Delta_G} \geq 1, \end{aligned}$$

since $e_H / \Delta_G \leq v_H - \alpha_H^*$ by Lemma 6.1. ■

Proof of Corollary 1.7: Suppose that $\Delta_G = k$ and that H is a k -regular subgraph of G . Then $e_H = kv_H/2$ and, by Lemma A.2, $\alpha_H^* = v_H/2$. Thus

$$\Psi_H^{1/\alpha_H^*} = (n^{v_H} p^{kv_H/2})^{2/v_H} = n^2 p^k.$$

Consequently, Lemma 6.2 yields, for $p \leq n^{-1/k}$,

$$\min_{H' \subseteq G} \Psi_{H'}^{1/\alpha_{H'}^*} = \Psi_H^{1/\alpha_H^*} = n^2 p^k,$$

and the result follows from Theorem 1.5.

Conversely, if $M_G^* = \Theta(n^2 p^k)$ for all $p \geq n^{-1/m_G}$, then Theorem 1.5 implies first $\Delta_G = k$ and secondly that for $n^{-1/m_G} \leq p \leq n^{-1/k}$ there must be equality in (6.4) for some $H \subseteq G$, which is possible only if $e_H = k\alpha_H^*$ and $v_H = 2\alpha_H^*$. But then $e_H = kv_H/2$; since $\Delta_H \leq \Delta_G = k$, this means that H is a k -regular subgraph of G . ■

Corollary 1.8 follows almost immediately from Theorem 1.5, using $\alpha_G^* = k$, and we omit the details.

Proof of Corollary 1.9: The case $p \geq n^{-1/2}$ is immediate from Theorem 1.5.

In the case $p \leq n^{-1/2}$, let us for ease of notation set $\tilde{\Psi}_k := \Psi_{P_k}^{1/\alpha_{P_k}^*}$ and $\tilde{M}_k := \min_{H \subseteq P_k} \Psi_H^{1/\alpha_H^*}$. We claim that if $k \geq 3$ is odd, then

$$(6.5) \quad \tilde{M}_k = n^{2\frac{k}{k+1}} p^{2\frac{k-1}{k+1}} \quad \text{if } n^{-1-1/(k-1)} \leq p \leq n^{-1/2},$$

and if $k \geq 4$ is even, then

$$(6.6) \quad \tilde{M}_k = \begin{cases} n^2 p^{2\frac{k-1}{k}} & \text{if } n^{-1-1/(k-1)} \leq p \leq n^{-1}, \\ n^{2\frac{k-1}{k}} p^{2\frac{k-2}{k}} & \text{if } n^{-1} \leq p \leq n^{-1/2}. \end{cases}$$

The result then follows by Theorem 1.5.

To show (6.5) and (6.6), we observe by Remark 1.6 that we only have to consider connected subgraphs $H \subseteq P_k$ in the definition of \tilde{M}_k . The only such H are themselves paths, and thus, for $k \geq 3$,

$$(6.7) \quad \tilde{M}_k = \min_{2 \leq j \leq k} \tilde{\Psi}_j = \min(\tilde{M}_{k-1}, \tilde{\Psi}_k).$$

If k is odd, $\alpha_{P_k}^* = (k + 1)/2$ and

$$\tilde{\Psi}_k = (n^k p^{k-1})^{2/(k+1)} = n^{2\frac{k}{k+1}} p^{2\frac{k-1}{k+1}},$$

while if k is even, $\alpha_{P_k}^* = k/2$ and

$$\tilde{\Psi}_k = (n^k p^{k-1})^{2/k} = n^2 p^{2\frac{k-1}{k}}.$$

It is now elementary to verify (6.5) and (6.6) by induction using (6.7), starting with $\tilde{M}_3 = \min(\tilde{M}_2, \tilde{\Psi}_3) = \min(\tilde{\Psi}_2, \tilde{\Psi}_3) = \tilde{\Psi}_3$. ■

7. The phases

We study in this section the exponent $M_G^*(n, p)$ as a function of p , for a given graph G . Compare the corresponding results for the lower tail in [11, §3.2]. We use the notation $M_G^* \asymp n^x p^y$ for $M_G^* = \Theta(n^x p^y)$.

Let $p = n^{-z}$ (where $z \geq 0$ does not have to be regarded as fixed). Define $\ell_H(z) := (v_H - e_H z)/\alpha_H^*$. Thus $\Psi_H^{1/\alpha_H^*} = n^{\ell_H(z)}$, and it follows from Theorem 1.5 that

$$(7.1) \quad M_G^*(n, n^{-z}) \asymp n^{L(z)},$$

where

$$(7.2) \quad L(z) := \begin{cases} 2 - \Delta_G z & \text{if } 0 \leq z \leq 1/\Delta_G, \\ \min_{H \subseteq G} \ell_H(z) & \text{if } 1/\Delta_G \leq z \leq 1/m_G, \\ 0 & \text{if } 1/m_G \leq z. \end{cases}$$

By Remark 1.6, it suffices to consider connected subgraphs H .

Clearly, $L(z)$ is a piecewise linear, continuous function, which is strictly decreasing on $[0, 1/m_G]$. See the example in Figure 1.

Let q denote (the minimal) number of line segments in $L(z)$ for $z \leq 1/m_G$. There exist real numbers $\{x_i\}_1^q$, $\{y_i\}_1^q$ and $\{z_i\}_1^{q+1}$, with $x_i > 0$, $y_i > 0$ and $1/m_G = z_1 > z_2 > \dots > z_{q+1} = 0$, such that $L(z) = x_i - y_i z$ for $z_i \geq z \geq z_{i+1}$ and thus, by (7.1),

$$(7.3) \quad M_G^*(n, p) \asymp n^{x_i} p^{y_i} \quad \text{if } n^{-z_i} \leq p \leq n^{-z_{i+1}}.$$




Figure 1. $L(z)$ (thick curve) for P_6 , with the corners marked by circles. The dotted line is $2 - \Delta_G z$; the thin lines are $\ell_H(z)$ for the connected subgraphs of P_6 ; the intersections with the z -axis are in order P_6, P_5, P_4, P_3, P_2 .

The above intervals for p , or z , will be referred to as the *phases* of the upper tail (of X_G , or G , for short). The interval $I_i = [z_{i+1}, z_i]$ will be called the i -th phase, $i = 1, \dots, q$, and in particular $I_q = [0, z_q]$ will be called the last phase. So, q stands for the number of phases. For instance, by Corollaries 1.7–1.9, regular graphs have one phase, stars (except K_2) and odd paths have two, while even paths have three phases of the upper tail (as in Figure 1).

If $z_q > 1/\Delta_G$, then it follows from the proof of Lemma 6.2 that $z_q = z_1 = m_G$ and $q = 1$. Hence, the graphs with only one phase of the upper tail are fully characterized by Corollary 1.7, and for all graphs G with more than one phase, the last phase is $[0, 1/\Delta_G]$, i.e. $n^{-1/\Delta_G} \leq p \leq 1$.

From now on let us assume that $q \geq 2$, and so $z_q = 1/\Delta_G$. For each phase I_i , $i = 1, \dots, q-1$, there exists a (not necessarily unique) subgraph $H_i \subseteq G$ such that $L(z) = \ell_{H_i}(z)$ for $z \in I_i$ and thus $x_i = v_{H_i}/\alpha_{H_i}^*$ and $y_i = e_{H_i}/\alpha_{H_i}^*$. We call these subgraphs H_i the *leading subgraphs* (for the upper tail). Note that the leading subgraphs are the subgraphs H that can be used in our construction for the lower bound in Theorem 1.2 for the respective ranges of p (see Theorem 3.1).

Returning to Corollaries 1.8 and 1.9, we see from the proofs that if G is a star or a path, it is itself the unique leading subgraph for the first phase. If G is an even path of order $k \geq 4$, the leading subgraph for the second phase is P_{k-1} (unique as an abstract graph, although there are two copies of it in G). For P_6 in Figure 1, the leading subgraphs are thus P_6 and P_5 .

7.1. THE SUBGRAPH PLOT. There is another, more graphic way to depict the phases. The subgraph plot in [11, §3.2], which yields information about the phases and leading subgraphs for the lower tail, can be adapted to the upper tail. In this version, let $Q_H := (v_H/\alpha_H^*, e_H/\alpha_H^*) \in \mathbb{R}^2$ and consider the set $\Sigma := \{Q_H : H \subseteq G, e_H > 0\}$ in the xy -plane. See Figures 2 and 3 below.

Note that m_G is the maximal slope of the lines from $(0, 0)$ to the points in Σ , while, by Lemma 6.1, Δ_G is the maximal slope of the lines from $(1, 0)$ to the points in Σ . Let $L^{(m)}$ and $L^{(\Delta)}$ be the lines $y = m_G x$ and $y = \Delta_G(x - 1)$, resp., and let C be their common point. Moreover, let $D := (2, \Delta_G)$. By the relation between ℓ_H and Q_H , it is easy to see the following.

There is only one phase for the upper tail of G if and only if $C = D$, and the graphs for which this happens are characterized by Corollary 1.7. Otherwise, let A be the point in $\Sigma \cap L^{(m)}$ nearest to C (on the right), and let B be the point in $\Sigma \cap L^{(\Delta)}$ nearest to C (on the left). If $C \in \Sigma$, then $A = B = C$ and there are exactly two phases, corresponding to C and D . The leading subgraphs for the first phase correspond to $A = B = C$ in the sense that $Q_H = C$.

If $C \notin \Sigma$, then there are at least three phases. Our “battlefield” is confined to the (geometric) triangle $T := ABC$. Let Σ^* be the convex hull $\text{Conv}(\Sigma \cap T)$. The extreme points of Σ^* , ordered with decreasing y -coordinate, form a sequence beginning with $Q_1 = A$ and ending with $Q_{q-1} = B$. It is then easily seen that these extreme points, supplemented by $Q_q = D$, correspond to the phases, in

the sense that $Q_i = (x_i, y_i)$ and that the leading subgraphs for the i -th phase ($1 \leq i \leq q-1$) are the subgraphs H for which $Q_H = Q_i$. Moreover, the slopes of the line segments joining the points Q_1, \dots, Q_{q-1} form an increasing sequence strictly contained between m_G and Δ_G , and their reciprocals are the numbers z_2, \dots, z_{q-1} , which together with $z_1 = 1/m_G$ and $z_q = 1/\Delta_G$ (the reciprocals of the slopes of $L^{(m)}$ and $L^{(\Delta)}$) determine the intervals I_1, \dots, I_q . Note that it follows that the exponents in (7.3) satisfy $x_1 > x_2 > \dots > x_{q-1} < x_q$ and $y_1 > y_2 > \dots > y_{q-1} < y_q$, with the monotonicity broken at the last phase.

In particular (when $C \neq D$), A and B represent the first and the last but one phases. For these two phases we can be more explicit about the leading subgraphs. It is seen from the subgraph plot and the definition of $Q_1 = A$ above that the leading subgraphs for the first phase are the subgraphs H with α_H^*/v_H maximal among the subgraphs with e_H/v_H maximal.

If G is strictly balanced (see [11] for definition), the leading subgraph in the first phase is thus G itself, and it is unique (see Examples 7.1–7.4). If G is balanced but not strictly balanced, G may be or not be a leading subgraph in the first phase, but is not a leading subgraph in any later phase (see Examples 7.5 and 7.6). If G is not balanced, then it may only become its own leading subgraph at a later phase (see Example 7.7). All of this contrasts with the lower tail, where G always is a (unique) leading subgraph in the first phase.

Similarly, for the last but one phase, the leading subgraphs H are those with α_H^*/v_H minimal among the subgraphs with $e_H/(v_H - \alpha^*) = \Delta_G$ (compare Lemma 6.1). The star $K = K_{1,\Delta}$ may or may not be among them (see Examples 7.2 – 7.7).

7.2. EXAMPLES. We finish this section with some examples illustrating this geometric approach. In each example, we provide the parameters $m = m_G$, $\Delta = \Delta_G$, the points A, B, C, D and $Q_i = (x_i, y_i)$, $i = 1, \dots, q$, and z_1, \dots, z_q ; the values of $M^* = M_G^*$ follow from these by (7.3). We also discuss the leading subgraphs, and in two cases include figures with the subgraph plots.

Example 7.1: We choose $G = K_3$ as a graph with just one phase. Here $m = 1$, $\Delta = 2$, $Q_1 = A = B = C = D = (2, 2)$, $z_1 = 1$, and $M^* \asymp n^2 p^2$ for all $p \geq 1/n$. We have not defined any leading subgraphs when there is only one phase; a natural definition would be the subgraphs corresponding to C , in this case K_3 itself.

Example 7.2: The star $G = K_{1,3}$ has just two phases (Corollary 1.8). We have $m = 3/4$, $\Delta = 3$, $Q_1 = A = B = C = (4/3, 1)$, $Q_2 = D = (2, 3)$, $z_1 = 4/3$,

$z_2 = 1/3$. Thus $M^* \asymp n^{4/3}p$ for $n^{-4/3} \leq p \leq n^{-1/3}$ and $M^* \asymp n^2p^3$ for $n^{-1/3} \leq p < 1$. The only leading subgraph is $H = G$.

Example 7.3: Also the path on five vertices $G = P_5$ has two phases with the leading subgraph $H = G$, which, however, is not the star $K = K_{1,\Delta}$. We have $m = 4/5$, $\Delta = 2$, $Q_1 = A = B = C = (5/3, 4/3)$, $Q_2 = D = (2, 2)$, $z_1 = 5/4$, $z_2 = 1/2$.

Example 7.4: The path on six vertices $G = P_6$ enjoys three phases with $H = G$ and P_5 being the two leading subgraphs. Here $m = 5/6$, $\Delta = 2$, $Q_1 = A = (2, 5/3)$, $Q_2 = B = (5/3, 4/3)$, $C = (12/7, 10/7)$, $Q_3 = D = (2, 2)$, and $z_1 = 6/5$, $z_2 = 1$, $z_3 = 1/2$. The subgraph plot is shown in Figure 2; the points Q_H in Σ (for connected subgraphs only) are shown as large crosses, and the lines $L^{(m)}$ and $L^{(\Delta)}$ together with Σ^* (in this case the line joining A and B) are added. Note that this figure is dual to Figure 1 (in the sense of projective geometry); the points (crosses) in Figure 2 correspond to the lines in Figure 1 and, conversely, a line such as AB through two points in in Figure 2 corresponds to the intersection point of the corresponding lines in Figure 1.




Figure 2. Subgraph plot for P_6 . Σ^* is the line AB .

Example 7.5: If G is a K_3 with an attached pendant vertex, both G and the subgraph K_3 are leading subgraphs in the first phase. This graph has three phases; $K = K_{1,3}$ is the unique leading subgraph for the second phase $n^{-2/3} \leq p \leq n^{-1/3}$. Here $m = 1$, $\Delta = 3$, $Q_1 = A = (2, 2)$, $Q_2 = B = (4/3, 1)$, $C = (3/2, 3/2)$, $Q_3 = D = (2, 3)$, $z_1 = 1$, $z_2 = 2/3$, $z_3 = 1/3$.

Example 7.6: If G_1 is a C_4 with an attached pendant vertex, G_1 is the unique leading subgraph in the first phase. This graph too has three phases, with $K = K_{1,3}$ the unique leading subgraph for the second phase $n^{-1/2} \leq p \leq n^{-1/3}$. Here $m = 1$, $\Delta = 3$, $Q_1 = A = (5/3, 5/3)$, $Q_2 = B = (4/3, 1)$, $C = (3/2, 3/2)$, $Q_3 = D = (2, 3)$, $z_1 = 1$, $z_2 = 1/2$, $z_3 = 1/3$.

If G_2 is a C_4 with a path of length 2 attached, the unique leading subgraph in the first phase is the subgraph G_1 . In fact, G_2 has the same phases and leading subgraphs, and indeed the same $L(z)$, as G_1 . Hence also m , Δ , A , B , C , D , Q_i and z_i are the same as for G_1 .




Figure 3. Subgraph plot for Example 7.7. Σ^* is shaded.

Example 7.7: Let G be the “kite graph”, formed of a copy of K_4 without one edge e and a pendant edge attached at a vertex not adjacent to e . We have $m = 5/4, \Delta = 4, Q_1 = A = (2, 5/2), Q_3 = B = (5/4, 1), C = (16/11, 20/11)$ and $Q_4 = D = (2, 4)$; in addition, $Q_2 = (5/3, 2)$ and $z_1 = 4/5, z_2 = 2/3, z_3 = 5/12, z_4 = 1/4$. This graph has four phases, and the three leading subgraphs are $K_4 - e, G$ and $K = K_{1,4}$ (all unique). The subgraph plot (connected subgraphs only) is shown in Figure 3.

7.3. GRAPHS WITH MANY PHASES. Our last example describes an infinite family of trees T^1, T^2, \dots , such that for every integer $k \geq 2$ the tree T^k has $k + 1$ phases for the upper tail. (For the lower tail, graphs with arbitrarily many phases were constructed in [21]. These are not trees, since a tree has only two phases for the lower tail.)

Example 7.8: Let T^k be the tree obtained by taking k stars $K_{1,i}, i = 1, \dots, k$, and tying them up by merging one pendant vertex from each star into one vertex. The knot vertex then has k neighbors v_1, \dots, v_k , and each v_i is, in turn, attached to $i - 1$ leaves. The total number of vertices is $1 + k + \binom{k}{2}$. We may call T^k a star cluster.

It is convenient to label the vertices by $1, \dots, 1 + k + \binom{k}{2}$, in the depth-first manner, beginning with the knot and searching the stars in order of decreasing size, i.e. taking the knot, then v_k followed by its attached leaves, then v_{k-1} followed by its leaves, and so on until v_1 . Figure 4 presents T^4 and the labelling of its vertices.




Figure 4. The tree T^4 with 5 phases.

PROPOSITION 7.9: *For every $k \geq 2$, the graph T^k described above has $k + 1$ phases for the upper tail.*

Proof: By Remark 1.6, it is enough to consider only connected subgraphs of $T = T^k$, that is, subtrees. Moreover, as is the case for all bipartite graphs, for trees $\alpha^* = \alpha$ — the size of the largest independent set. So, if H is a subtree with

h vertices, then $\Psi_H = (n^h p^{h-1})^{1/\alpha_H}$, and the minimum of Ψ_H over all H with given h is obtained by maximizing α_H . Let $\alpha(h) := \max \alpha_H$ over all subgraphs H of T with $v_H = h$. Furthermore, let $Q_h^* = (h/\alpha(h), (h-1)/\alpha(h))$. In the upper tail subgraph plot, we thus only have to consider the points Q_1^*, \dots, Q_v^* , where $v = v_T$.

Let T_h be the subtree of T induced by the vertices $1, \dots, h$ (in the labelling described in Example 7.8). It is straightforward to show that for each $h = 1, \dots, v$, we have $\alpha(h) = \alpha_{T_h} = h - d(h)$, where $d(h)$ is the degree in T_h of the knot. (A maximal independent set in T_h is obtained by taking all vertices of T_h except the v_i 's.)

Let us now distinguish those numbers h among $1, \dots, v$ for which T_h consists of a number of complete stars in T , i.e. define

$$h(j) = 1 + k + (k - 1) + \dots + (k - j + 1) = 1 + kj - j(j - 1)/2,$$

$j = 0, \dots, k$. We thus have $\alpha(h(j)) = h(j) - j$.

Note that if h is *not* of the form $h(j)$, one has $\alpha(h + 1) = \alpha(h) + 1$; it follows by a simple calculation that Q_h^* lies to the right of Q_{h+1}^* and that the slope from Q_{h+1}^* to Q_h^* is

$$\frac{h - \alpha(h) - 1}{h - \alpha(h)} = \frac{d(h) - 1}{d(h)} < \frac{v - 1}{v} = m.$$

This means that for such h , the points Q_h^* do not appear as extreme points of the convex hull Σ^* ; consequently, the extreme points Q_i in the upper tail subgraph plot are among the points $Q_{h(j)}^*$, $j = 1, \dots, k$. We will now show that these points lie on a (strictly) concave curve, proving that all these points are extreme points and that T thus has $k + 1$ phases, given by them and by $D = (2, k)$; more precisely, $Q_i = Q_{h(k+1-i)}^*$, $1 \leq i \leq k$. It will also follow that $T_{h(k+1-j)}$ is a leading subgraph for the i -th phase. (Except for the first phase, the leading subgraph is not unique.)

To see that the points $Q_{h(j)}^*$, $j = 0, \dots, k$, lie on a concave curve, regard $h(t) := 1 + (k + 1/2)t - t^2/2$ as a function of a real variable t , define $x(t) = h(t)/(h(t) - t)$ and $y(t) = (h(t) - 1)/(h(t) - t)$ and note that $Q_{h(j)}^* = (x(j), y(j))$. For $t \in [0, k]$ we have $h(t) - t > 0$, and by standard calculus, $dx/dt > 0$ and

$$dy/dx = (dy/dt)/(dx/dt) = 1 + (k - 1/2 - t)/(1 + t^2/2).$$

Hence, x is an increasing function of t , dy/dx is a decreasing function of t , and thus dy/dx is a decreasing function of x , showing that y is a strictly concave function of x . ■

8. Further remarks

Remark 8.1: The main remaining problem for the upper tail, as we see it, is the gap between the upper and lower bounds in Theorem 1.2. For constant t , the gap amounts to a factor $O(\log(1/p))$ in the exponent. In a few special cases, we can say more.

When $G = K_2$, and more generally in the rather trivial case when G is a matching, it is easy to see that for $p \geq n^{-1/m_G} = n^{-2}$, there is a lower bound $\exp\{-C(t, G)M_G^*(n, p)\}$ of the same type as the upper bound. In particular, the lower bound in Theorem 1.2 is not sharp for these graphs.

In the case when p is below the threshold n^{-1/m_G} , we have $M^* = \Theta(1)$, our upper bound is $\Theta(1)$, and the lower bound is $p^{\Theta(1)}$. On the other hand, at least for balanced G , the correct probability is just $\mathbf{P}(X_G > 0)$, which, by the result of [10], see (1.1), is $\Theta(\mathbf{E}X_G)$. Indeed, when $\mathbf{E}X_G < 1/t$, to have more than $t\mathbf{E}X_G$ copies of G is the same as to have at least one. Thus the upper bound is not sharp here, but the case when 1 is much more than $t\mathbf{E}X_G$ is not very interesting.

Returning to the case when $\mathbf{E}X_G$ is large, it is shown in [13] that for $G = K_4$ and $n^{-2/3} \log^{1/6} n \ll p \leq n^{-1/2-\epsilon}$, there is an upper bound

$$\mathbf{P}(X_G \geq 2\mathbf{E}X_G) \leq \exp\{-cM_G^*(n, p) \log^{1/2} n\},$$

which is half-way (on a doubly logarithmic scale) between the upper and lower bounds in Theorem 1.2. Hence the upper bound in Theorem 1.2 is not sharp in this case. The same holds for $G = C_4$ in a similar range. The proof uses a different method than the present paper and does not easily extend to other graphs. We do not know exactly what happens for K_4 when $p \geq n^{-1/2}$.

These examples show that neither the upper nor lower bound in Theorem 1.2 is sharp in general.

Remark 8.2: Our proof of the upper bound in Theorem 1.2 yields the estimate $c(t, G) = \Omega(t - 1)^2$ for $1 < t \leq 2$, and $c(t, G) = \Omega(\log t)$ for $t \geq 2$. Using a version of Lemma 2.1 for $m_1 > m_2$, the latter can easily be improved to

$$(8.1) \quad c(t, G) = \Omega(t^{1/\max_{H \subseteq G} \alpha_H^*}), \quad t \geq 2.$$

Note that for the trivial case $G = K_2$, when X_G has a binomial distribution, the upper bound in Theorem 1.2 is optimal (assuming p and t are not too extreme), with $c(t, G) = \Theta(t - 1)^2$ for $1 < t \leq 2$ and $c(t, G) = \Theta(t \log t)$ for $t \geq 2$. It follows that if G is a matching with j disjoint edges, then $c(t, G)$ can

be at most $c(t, G) = \Theta(t^{1/j} \log t)$ for $t \geq 2$. This suggests that the best $c(t, G)$ for large t in general might be of the form (8.1), possibly with an additional factor $\log t$, but we have no proof of that.

Our proof of the lower bound yields $C(t, G) = O(t)$, but we have not investigated this constant further.

Remark 8.3: It follows from our results that the upper tail probabilities typically are larger than the lower tail probabilities in (1.1). More precisely, suppose that $\Delta_G \geq 2$ and that $n \rightarrow \infty$ and $p \rightarrow 0$ with $p \gg n^{-1/m_G}$ (i.e. $pn^{-1/m_G} \rightarrow \infty$). Then

$$(8.2) \quad M_G^* \ll \min_{H \subseteq G, e_H > 0} \Psi_H.$$

This can be shown directly from Definition (1.2), but we use instead Theorem 1.5:

Let $H \subseteq G$ with $e_H > 0$. First, consider the case $n^{-1/m_G} \ll p \leq n^{-1/\Delta_G}$. If $\alpha_H^* > 1$ we have, since $p \gg n^{-1/m_G}$ implies $\Psi_H \rightarrow \infty$, $\Psi_H \gg \Psi_H^{1/\alpha_H^*}$. On the other hand, the only H with $e_H > 0$ and $\alpha_H^* < 1$ is $H = K_2$. In this case, let $S = K_{1, \Delta_G}$ be a maximal star in G and note that

$$\Psi_H = n^2 p \gg n^{1+1/\Delta_G} p = \Psi_K^{1/\alpha_K^*}.$$

Hence, $\min_H \Psi_H \gg \min_H \Psi_H^{1/\alpha_H^*} = \Theta(M_G^*)$.

In the case $n^{-1/\Delta_G} \leq p \ll 1$, we note that $e_H \leq \Delta_G v_H / 2$. Thus, if $v_H > 2$,

$$\Psi_H \geq n^{v_H} p^{\Delta_G v_H / 2} = (n^2 p^{\Delta_G})^{v_H / 2} \gg n^2 p^{\Delta_G} = \Theta(M_G^*).$$

The remaining case is again $H = K_2$, but then

$$\Psi_H = n^2 p \gg n^2 p^{\Delta_G} = \Theta(M_G^*).$$

Hence (8.2) holds for $n^{-1/m_G} \ll p \ll 1$, which shows that the bound for the lower tail in (1.1) is smaller than the upper bound for the upper tail in Theorem 1.2. It is easily seen by arguments as above that if we further require $n^{-1/m_G} \log n \ll p \ll 1$, then also the lower bound in Theorem 1.2 is larger than (1.1), so the two tails are of different orders.

A heuristic reason for the difference is that it is possible to create many copies of G using comparatively few excess edges, as in our construction in Section 3, while there seems to be no similar way to greatly reduce the number of copies by deleting rather few edges.

Remark 8.4: As mentioned in the introduction, Spencer [22], see also [2, §8.5] and [23, §5.2], studied a generalization of subgraph counts X_G , viz. the number of *extensions* isomorphic to G of a given set of vertices in $G(n, p)$. Although we have not examined this more general problem, we believe that our methods apply here too.

Remark 8.5: Alon's result in [1], which is the case $n \geq m$ of our Theorem 1.3, was extended to hypergraphs by Friedgut and Kahn [7]. We guess that Theorem 1.3 too has an extension to hypergraphs, in particular, since we use the method of proof in [7]. However, we have not pursued this.

Problem 8.6: Graphs with arbitrarily many phases were constructed in Example 7.8 above. *How many phases are possible for graphs with a given number of vertices?* (For the lower tail, see [21].)

Appendix A. Fractional independence number

For a given graph H , the *fractional independence number* of H , denoted by α_H^* , is the largest value of $\sum_v \alpha_v$ over all assignments of weights $\alpha_v \in [0, 1]$ to the vertices of H satisfying the condition $\alpha_v + \alpha_u \leq 1$ for all edges uv of H . It follows immediately from this definition that if $e_H > 0$ then

$$(A.1) \quad 1 \leq v_H/2 \leq \alpha_H^* \leq v_H - 1.$$

The upper bound can be strengthened.

LEMMA A.1: *For every graph H with $e_H > 0$,*

$$(A.2) \quad \alpha_H^* \leq v_H - \frac{e_H}{\Delta_H}.$$

Proof: For any feasible assignment of weights α_v to $V(H)$, if $d_H(v)$ is the degree of v in H ,

$$\begin{aligned} e_H &\leq \sum_{vu \in E(H)} (1 - \alpha_v + 1 - \alpha_u) = \sum_{v \in V(H)} d_H(v)(1 - \alpha_v) \\ &\leq \Delta_H \sum_{v \in V(H)} (1 - \alpha_v) = \Delta_H \left(v_H - \sum_v \alpha_v \right). \end{aligned}$$

Taking an optimal assignment, we obtain (A.2). ■

For regular graphs, the lower bound in (A.1) and the upper in (A.2) coincide, so we have a simple formula:

LEMMA A.2: *If the graph H is regular and $e_H > 0$, then*

$$\alpha_H^* = v_H/2. \quad \blacksquare$$

Note that the bound in (A.2) also is achieved by, among others, stars and paths of even length. Incidentally, in both these cases, α_H^* is just equal to the independence number α_H , that is, there is a 0-1 optimal assignment. It is well-known that every vertex of the polytope defined by $x_v \geq 0$ ($v \in V(H)$) and $x_v + x_u \leq 1$ ($uv \in E(H)$) is half-integral, and consequently α_H^* can always be realized by an optimal assignment with $\alpha_v \in \{0, \frac{1}{2}, 1\}$. (This result is sometimes attributed to Lovász, who, however, denies it, saying that it was already folklore before.) It is characteristic of all bipartite graphs that α_H^* has an integral solution (but there are, clearly, non-bipartite graphs, like the one in Example 7.5, for which this is true as well).

PROPOSITION A.3: *If H is bipartite, then $\alpha_H^* = \alpha_H$.*

Proof: If H has isolated vertices, they contribute equally to both parameters. So, assume that the minimum degree $\delta_H \geq 1$, and consider the dual problem of minimizing $\sum_e \rho_e$ over all assignments of weights $\rho_e \in [0, 1]$ satisfying the condition $\sum_{e \ni v} \rho_e \geq 1$ for all vertices v of H . This minimum value is denoted by ρ_H^* and called the *fractional edge covering number*. By the duality theorem of linear programming (see, e.g., [16]) $\alpha_H^* = \rho_H^*$ for all graphs H . Trivially, we also have $\rho_H \geq \rho_H^* = \alpha_H^* \geq \alpha_H$, where ρ_H is the edge covering number of H , i.e. the smallest number of edges of H whose union is $V(H)$. But, for bipartite graphs H , as a consequence of Gallai’s Theorem, we have in fact $\alpha_H = \rho_H$ (see, e.g., Corollary 1.1.7 in [16]), which implies that indeed $\alpha_H^* = \alpha_H$. \blacksquare

Finally, let us characterize the graphs for which $\alpha_H^* = v_H/2$. By Lemma A.2, all regular graphs are among them (and in fact, every graph with a regular *spanning* subgraph). A *perfect 2-matching* of a graph H (see [16], page 216) is a spanning subgraph of H whose every component is either a cycle or an edge. The following result is implicit in [1], but it can also be deduced from [16]. Here we provide a simple proof based on the fact that a graph has a perfect 2-matching if and only if it satisfies the Hall condition that every independent set A of vertices has at least $|A|$ neighbors — see Corollary 6.1.5 in [16]. (Note that for regular graphs this condition holds.)

PROPOSITION A.4: *We have $\alpha_H^* = v_H/2$ if and only if H has a perfect 2-matching.*

Proof: If an independent set A violates the Hall condition then assign 1 to all vertices in A , 0 to all neighbors of A and $\frac{1}{2}$ otherwise. This is a feasible assignment of total weight greater than $v_H/2$. Conversely, if $\alpha_H^* > v_H/2$, then, in an optimal, half-integral assignment, the set A of vertices with weight 1 must have bigger size than the set B of vertices with weight 0. But A is an independent set with no neighbors outside B . ■

Alternatively, Proposition A.4 can be proved without invoking Hall’s condition, by using instead the duality theorem and the fact that ρ_H^* can also be realized by a half-integral assignment.

Appendix B. An application of upper tail

The following illustration is a simplified form of a problem appearing in [8] (see also [12]). The main difference is that here we deal with induced subgraphs of a random graph, while the real problem in [8] involves non-induced subgraphs, and so a regularity lemma for sparse graphs has to be used in order to count the edges.

Given a graph G , its *base* $B(G)$ is defined as the graph with vertex set $V(G)$ and the edges taken as all pairs $\{u, v\}$ such that for some $w \in V(G)$ we have $uw, vw \in E(G)$. It is straightforward to prove that if, say, $p \geq (2.1 \log n/n)^{1/2}$ then, a.a.s. $B(G(n, p)) = K_n$, while already for $p = \Theta(1/\sqrt{n})$, we have a.a.s. $B = \Theta(n^2)$, where we for simplicity write $B := e(B(G(n, p)))$.

Does the last estimate hold for *all* sufficiently large induced subgraphs of $G(n, p)$, say of size $n/2$? As there are roughly 2^n such subgraphs, it is sufficient to show that if $a > 0$, the random graph $G(n/2, p)$, or, equivalently, $G(n, p)$, where $p = a/\sqrt{n}$, satisfies $B = \Theta(n^2)$ with probability at least $1 - o(2^{-n})$. Below we provide a short proof of this fact based on our upper tail estimate.

Let M be the number of edges and let $X = X_{C_4}$ be the number of 4-cycles C_4 in $G(n, p)$. Furthermore, let D_v be the degree of vertex v and T_{uv} the number of “tepees” over uv , that is, the number of vertices w such that uw, vw are edges of $G(n, p)$. Then, by double counting and convexity of $\binom{x}{2}$, we have, provided $M \geq n$,

$$\sum_{uv} T_{uv} = \sum_v \binom{D_v}{2} \geq n \binom{\sum D_v/n}{2} = n \binom{2M/n}{2} \geq \frac{M^2}{n}$$

and then either $B > M^2/2n$ or $\sum_{uv} T_{uv}/B \geq M^2/nB \geq 2$. In the latter case

we similarly have

$$2X = \sum_{uv} \binom{T_{uv}}{2} \geq B \binom{\sum T_{uv}/B}{2} \geq \frac{M^4}{4Bn^2},$$

where the sums \sum_{uv} are taken over all pairs uv with $T_{uv} > 0$, and thus have precisely B terms. Hence (still deterministically), if $M \geq n$,

$$B \geq \min \left(\frac{M^4}{8n^2 X}, \frac{M^2}{2n} \right).$$

By Chernoff's bound (see, e.g., [11]), $M \geq n^2 p/3$ with probability at least $1 - e^{-\Theta(n^{3/2})}$. Noticing that $\mathbf{E}X = \Theta(n^4 p^4)$, all we need to complete the proof is a bound of the form $\mathbf{P}(X \geq t\mathbf{E}X) \leq e^{-n}$, which is provided by our Theorem 1.2 (upper bound) with t constant and large enough (see Remark 8.2) together with Corollary 1.7.

References

- [1] N. Alon, *On the number of subgraphs of prescribed type of graphs with a given number of edges*, Israel Journal of Mathematics **38** (1981), 116–130.
- [2] N. Alon and J. Spencer, *The Probabilistic Method*, 2nd ed., John Wiley and Sons, New York, 2000.
- [3] B. Bollobás, *Threshold functions for small subgraphs*, Mathematical Proceedings of the Cambridge Philosophical Society **90** (1981), 197–206.
- [4] F. R. K. Chung, R. L. Graham, P. Frankl and J. B. Shearer, *Some intersection theorems for ordered sets and graphs*, Journal of Combinatorial Theory, Series A **43** (1986), 23–37.
- [5] P. Collet and J.-P. Eckmann, *The number of large graphs with a positive density of triangles*, Journal of Statistical Physics **109** (2002), 923–943.
- [6] P. Erdős and A. Rényi, *On the evolution of random graphs*, Publications of the Mathematical Institute of the Hungarian Academy of Sciences **5** (1960), 17–61.
- [7] E. Friedgut and J. Kahn, *On the number of copies of one hypergraph in another*, Israel Journal of Mathematics **105** (1998), 251–256.
- [8] E. Friedgut, V. Rödl, A. Ruciński and P. Tetali, *A sharp threshold for random graphs with a monochromatic triangles in every edge coloring*, Submitted.
- [9] S. Janson, *Poisson approximation for large deviations*, Random Structures and Algorithms **1** (1990), 221–230.

- [10] S. Janson, T. Luczak and A. Ruciński, *An exponential bound for the probability of nonexistence of a specified subgraph in a random graph*, in *Random Graphs '87 (Proceedings, Poznań 1987)* (M. Karoński, J. Jaworski and A. Ruciński, eds.), John Wiley and Sons, Chichester, 1990, pp. 73–87.
- [11] S. Janson, T. Luczak and A. Ruciński, *Random Graphs*, John Wiley and Sons, New York, 2000.
- [12] S. Janson and A. Ruciński, *The infamous upper tail*, *Random Structures and Algorithms* **20** (2002), 317–342.
- [13] S. Janson and A. Ruciński, *The deletion method for upper tail estimates*, Technical report U.U.D.M. 2000:28, Uppsala. Available from <http://www.math.uu.se/~svante/>.
- [14] J.-P. Kahane, *Some Random Series of Functions*, 2nd ed., Cambridge University Press, Cambridge, 1985.
- [15] J. H. Kim and V. H. Vu, *Divide and conquer martingales and the number of triangles in a random graph*, *Random Structures and Algorithms* (2004), to appear.
- [16] L. Lovász and M. D. Plummer, *Matching Theory*, North-Holland Mathematics Studies 121, Annals of Discrete Mathematics 29, North-Holland Publishing Co., Amsterdam; Akadémiai Kiadó, Budapest, 1986.
- [17] D. Panchenko, *Deviation inequality for monotonic Boolean functions with application to a number of k -cycles in a random graph*, *Random Structures and Algorithms* **24** (2004), 65–74.
- [18] V. Rödl and A. Ruciński, *Random graphs with monochromatic triangles in every edge coloring*, *Random Structures and Algorithms* **5** (1994), 253–270.
- [19] V. Rödl and A. Ruciński, *Threshold functions for Ramsey properties*, *Journal of the American Mathematical Society* **8** (1995), 917–942.
- [20] A. Ruciński, *When are small subgraphs of a random graph normally distributed?* *Probability Theory and Related Fields* **78** (1988), 1–10.
- [21] A. Ruciński, *Convex hulls of graphs*, *Ars Combinatoria* **32** (1991), 293–300.
- [22] J. Spencer, *Counting extensions*, *Journal of Combinatorial Theory, Series A* **55** (1990), 247–255.
- [23] J. Spencer, *The Strange Logic of Random Graphs*, Springer-Verlag, Berlin, 2001.
- [24] V. H. Vu, *A large deviation result on the number of small subgraphs of a random graph*, *Combinatorics, Probability and Computing* **10** (2001), 79–94.