

UPPER TAILS VIA HIGH MOMENTS AND ENTROPIC STABILITY

MATAN HAREL, FRANK MOUSSET, AND WOJCIECH SAMOTIJ

ABSTRACT. Suppose that X is a bounded-degree polynomial with nonnegative coefficients on the p -biased discrete hypercube. Our main result gives sharp estimates on the logarithmic upper tail probability of X whenever an associated extremal problem satisfies a certain entropic stability property. We apply this result to solve two long-standing open problems in probabilistic combinatorics: the upper tail problem for the number of arithmetic progressions of a fixed length in the p -random subset of the integers and the upper tail problem for the number of cliques of a fixed size in the random graph $G_{n,p}$. We also make significant progress on the upper tail problem for the number of copies of a fixed regular graph H in $G_{n,p}$. To accommodate readers who are interested in learning the basic method, we include a short, self-contained solution to the upper tail problem for the number of triangles in $G_{n,p}$ for all $p = p(n)$ satisfying $n^{-1} \log n \ll p \ll 1$.

CONTENTS

1. Introduction	1
2. Triangles in random graphs	9
3. The main technical result: ‘entropic stability implies localisation’	13
4. Arithmetic progressions in random sets of integers	19
5. Counting small subgraphs—a graph-theoretic interlude	25
6. Cliques in random graphs	34
7. Extensions to regular graphs	47
8. The Poisson regime	55
9. Beyond polynomials with nonnegative coefficients	63
10. Concluding remarks	64
References	65

1. INTRODUCTION

Suppose that $Y = (Y_1, \dots, Y_N)$ is a sequence of independent Bernoulli random variables with mean p and that $X = X(Y)$ is an N -variate polynomial with nonnegative real coefficients. Perhaps the simplest question that can be asked about the typical behaviour of X is whether it satisfies a law of large numbers, i.e., whether $X \rightarrow \mathbb{E}[X]$ in probability as $N \rightarrow \infty$. Once this is established, it is natural to ask for quantitative estimates of the probability of the event that X differs from its mean by a significant amount. In the special case where $Y \mapsto X(Y)$ is a linear function, this problem is addressed by the classical theory of large deviations, see [14, 28]. This theory shows that, under mild conditions on the coefficients of the linear function X and the assumptions $p \rightarrow 0$ and $Np \rightarrow \infty$,

$$\mathbb{P}(|X - \mathbb{E}[X]| \geq \delta \mathbb{E}[X]) = e^{-(I(\delta) + o(1))Np}$$

Date: July 2, 2021.

This research is supported in part by the Israel Science Foundation grants 1147/14 (FM and WS), 1028/16 (FM), and 1145/18 (FM and WS), the ERC Starting Grant 633509 (FM) and 678520 (MH), and the Zuckerman Postdoctoral Fellowship Program (MH).

for an explicitly computable function $I: (0, \infty) \rightarrow (0, \infty]$. However, there are many natural situations where one would like to consider nonlinear polynomials $X(Y)$, as in the following two examples. We use the notation $\llbracket N \rrbracket = \{1, \dots, N\}$.

Example 1.1 (Arithmetic progressions in random sets of integers). A k -term arithmetic progression is a sequence of k integers of the form $(a, a+b, a+2b, \dots, a+(k-1)b)$, where $b > 0$. We write $\llbracket N \rrbracket_p$ for the random subset of $\llbracket N \rrbracket$ obtained by including every element of $\llbracket N \rrbracket$ independently with probability p . Let $X_{N,p}^{k\text{-AP}}$ denote the number of k -term arithmetic progressions in $\llbracket N \rrbracket_p$. Then $X_{N,p}^{k\text{-AP}}$ can be considered as a polynomial with nonnegative coefficients and degree k in the independent $\text{Ber}(p)$ random variables Y_1, \dots, Y_N , where Y_i is the indicator variable of the event that $i \in \llbracket N \rrbracket_p$; explicitly,

$$X_{N,p}^{k\text{-AP}} = \sum_{b>0} \sum_{\substack{a \geq 1 \\ a+(k-1)b \leq N}} \prod_{i=0}^{k-1} Y_{a+ib}.$$

We remark that, unlike [10] and several other works, we count only genuine arithmetic progressions (i.e., we do not consider degenerate progressions of the form (a, \dots, a)) and we count every progression only once (as opposed to counting (a_1, a_2, \dots, a_k) and $(a_k, a_{k-1}, \dots, a_1)$ as two different progressions).

Example 1.2 (Subgraph counts in random graphs). Fix a nonempty graph H and let $X_{n,p}^H$ be the number of copies of H in the random graph $G_{n,p}$, that is, the number of subgraphs of $G_{n,p}$ isomorphic to H . Then $X_{n,p}^H$ can be written as a polynomial with nonnegative coefficients and degree e_H in the $N = \binom{n}{2}$ indicator random variables of the possible edges of $G_{n,p}$. More precisely, fix an arbitrary bijection $\sigma_n: \binom{\llbracket n \rrbracket}{2} \rightarrow \llbracket N \rrbracket$ (the precise choice is irrelevant) and, for every $i \in \llbracket N \rrbracket$, let Y_i be the indicator variable of the event that $\sigma_n^{-1}(i)$ is an edge in $G_{n,p}$. Then Y_1, \dots, Y_N are independent $\text{Ber}(p)$ random variables and

$$X_{n,p}^H = \sum_{\substack{H' \subseteq K_n \\ H' \cong H}} \prod_{e \in E(H')} Y_{\sigma_n(e)},$$

where K_n denotes the complete graph on the vertex set $\llbracket n \rrbracket$ and \cong denotes the isomorphism relation on graphs.

In this paper, we will always assume that δ is fixed and $p \rightarrow 0$ as $N \rightarrow \infty$.

The large deviation problem for the variables described above is significantly more involved than the linear case; in particular, the lower and upper tail probabilities—that is, $\mathbb{P}(X \leq (1 - \delta) \mathbb{E}[X])$ and $\mathbb{P}(X \geq (1 + \delta) \mathbb{E}[X])$, respectively—exhibit dramatically different behaviours. On the one hand, using a combination of Harris’s inequality [41] and Janson’s inequality [43], one can show that $X = X_{N,p}^{k\text{-AP}}$ satisfies

$$e^{-C_1(\delta) \min\{\mathbb{E}[X], Np\}} \leq \mathbb{P}(X \leq (1 - \delta) \mathbb{E}[X]) \leq e^{-C_2(\delta) \min\{\mathbb{E}[X], Np\}} \quad (1)$$

for some positive $C_1(\delta)$ and $C_2(\delta)$; similar bounds are available for $X = X_{n,p}^H$.¹ On the other hand, there are no comparably simple tools that allow one to easily obtain similar estimates on the logarithm of the upper tail probability. The standard concentration inequalities due to Azuma–Hoeffding [42], Talagrand [64], or Kim–Vu [52, 68] yield bounds that are far from tight in Examples 1.1 and 1.2. For a survey discussing these and other classical approaches to the ‘infamous upper tail’ problem, see [45].

Unlike the lower tail, the upper tail is susceptible to the influence of small structures whose appearance increases the value X atypically, a phenomenon that we refer to as *localisation*. For example, in the case of $X = X_{N,p}^{k\text{-AP}}$ where $k \geq 3$, a typical subset of size $m = o(N)$ contains

¹For more precise results, we refer the interested reader to [48, 54, 58, 71].

$\Theta(N^2(m/N)^k) = o(m^2)$ k -term arithmetic progressions, whereas some very rare subsets (notably an interval of length m) can contain as many as $\Theta(m^2)$ such progressions. The event that $\llbracket N \rrbracket_p$ contains an interval of length $\Theta(\sqrt{\mathbb{E}[X]})$ thus provides a lower bound on the upper tail probability. More precisely, $\mathbb{P}(X \geq (1 + \delta) \mathbb{E}[X]) \geq \exp(-O(\sqrt{\mathbb{E}[X]} \log(1/p)))$, which is significantly larger than the lower tail probability (1) for most p . In order to properly analyse the upper tail event, one must account for these local effects, which frequently requires understanding the peculiar combinatorial nature of the random variable X .

The last decade has seen the development of an increasingly powerful theory of ‘nonlinear large deviations’, which began with the work of Chatterjee–Dembo [18] and was further developed by Eldan [30], Cook–Dembo [24], Augeri [2, 3], and Cook–Dembo–Pham [25]. Whenever a general function f of i.i.d. random variables satisfies certain complexity and regularity conditions, these results can be used to express the upper tail probability $\mathbb{P}(f \geq (1 + \delta) \mathbb{E}[f])$ in terms of an associated variational problem. In the case where f is a polynomial with nonnegative coefficients on the hypercube, this variational problem is able to capture the presence of localisation, if it occurs. In the two examples mentioned above, nonlinear large deviation theory gives tight control of the upper tail probabilities whenever $p \geq N^{-\alpha}$ for some constant $\alpha > 0$. However, the best-known constant α is not optimal in both examples.

Our main contribution is a general method for proving sharp bounds on the upper tail probability of the polynomial $X = X(Y)$ in the presence of localisation. In many cases where localisation occurs, our approach can also give a coarse description of the tail event. At the heart of our method lies an adaptation of the classical moment argument of Janson, Oleszkiewicz, and Ruciński [44], which we use to formalise the intuition that the upper tail event is dominated by the appearance of near-minimisers of the combinatorial optimisation problem

$$\Phi_X(\delta) = \min \left\{ |I| \log(1/p) : I \subseteq \llbracket N \rrbracket \text{ and } \mathbb{E}[X \mid Y_i = 1 \text{ for all } i \in I] \geq (1 + \delta) \mathbb{E}[X] \right\}. \quad (2)$$

Roughly speaking, we say that $I \subseteq \llbracket N \rrbracket$ is a *core* if it is a feasible set for the above optimisation problem, its size is $O(\Phi_X(\delta))$, and it satisfies a certain natural rigidity condition that arises from requiring every element in the set to contribute a sizeable amount to the expectation. The constraints used to define a core are loosely analogous to the complexity conditions used in nonlinear large deviation theory; we will give a more precise definition of a core and discuss its relations to earlier work in more detail in Section 3.

We show that the upper tail probability is approximately equal to the probability of the appearance of a core. In particular, when the number of cores of size m is $(1/p)^{o(m)}$, a property we term *entropic stability*, then a union bound implies that $-\log \mathbb{P}(X \geq (1 + \delta) \mathbb{E}[X])$ is well approximated by $\Phi_X(\delta)$. We will verify that the random variables $X_{N,p}^{k\text{-AP}}$ and $X_{n,p}^H$ (for a large class of graphs H) satisfy the entropic stability condition under optimal, or nearly optimal, assumptions on p .²

One important caveat that we have ignored so far is that the upper tail exhibits localisation only when the expectation of X tends to infinity sufficiently quickly. In fact, if $\mathbb{E}[X]$ is of constant order, then, under relatively mild conditions, X converges in distribution to a Poisson random variable and no localisation occurs. We show that, for the two examples discussed above, the upper tail continues to have Poisson behaviour even when $\mathbb{E}[X]$ goes to infinity sufficiently slowly. In the cases of k -term arithmetic progressions in $\llbracket N \rrbracket_p$ and cliques in $G_{n,p}$, our results for the Poisson and localised regimes cover almost the whole range of densities $p \rightarrow 0$ with $\mathbb{E}[X] \rightarrow \infty$,

²We use the phrase *entropic stability* in a similar sense to the notion of stability in extremal combinatorics. More precisely, we are considering situations where the probability that $\prod_{i \in I} Y_i = 1$ for some minimiser I of (2) is asymptotically equal to the probability of appearance of one such minimiser—in other words, the *energy* of such configurations (given by the number of elements involved) dominates over the *entropy* (that is, the number of such configurations). Then *entropic stability* means that the entropy term remains negligible even as we move away from true minimisers of (2) to sets that are merely close to being minimisers (the cores).

leaving the upper tail probability undetermined only at densities for which it is believed that the two behaviours coexist.

1.1. Arithmetic progressions in random sets of integers. Let $X = X_{N,p}^{k\text{-AP}}$ denote the number of k -term arithmetic progressions in $\llbracket N \rrbracket_p$. It is not hard to see that $\mathbb{E}[X] = \Theta(N^2 p^k)$. Whenever this expectation vanishes, the upper tail event is commensurate with the probability of $X \geq 1$, which can be controlled using Markov's inequality. More generally, if $\mathbb{E}[X]$ is bounded, then it follows from standard techniques that X is asymptotically Poisson [7]; in this case, the large deviations of X are those of a Poisson random variable with mean $\mathbb{E}[X]$. For the remainder of this section, we shall thus assume that $\mathbb{E}[X] \rightarrow \infty$, i.e., that $p^{k/2} \gg N^{-1}$.

Improving an earlier estimate due to Janson and Ruciński [47], Warnke [69] proved that under fairly general assumptions (in particular, for constant $\delta > 0$ and all p bounded away from 1),

$$-\log \mathbb{P}(X \geq (1 + \delta) \mathbb{E}[X]) = \Theta \left(\min \left\{ ((1 + \delta) \log(1 + \delta) - \delta) \mathbb{E}[X], \sqrt{\delta \mathbb{E}[X]} \log(1/p) \right\} \right), \quad (3)$$

where the constants implicit in the Θ -notation are independent of δ . Note that the two terms of the minimum correspond to the dominance of the Poisson and the localised regimes, respectively.

Since then, it has been an open problem to determine the missing constant factor in (3). Using the above-mentioned framework of Eldan [30], Bhattacharya, Ganguly, Shao, and Zhao [10] were able to do so in the range $N^{-\frac{1}{12(k-1)}} (\log N)^{O(1)} \ll p^{k/2} \ll 1$. This was subsequently improved by Briët–Gopi [15] to the slightly wider range $N^{-\frac{1}{12\lceil(k-1)/2\rceil}} (\log N)^{O(1)} \ll p^{k/2} \ll 1$, also using Eldan's result. The two theorems below improve significantly on these results and determine the precise rate of the upper tail for all $N^{-1} \ll p^{k/2} \ll 1$, excepting the case $p^{k/2} = \Theta(N^{-1} \log N)$. The first result concerns the range where the minimum in (3) is $\sqrt{\delta \mathbb{E}[X]} \log(1/p)$.

Theorem 1.3. *Let $k \geq 3$ be an integer and let $X = X_{N,p}^{k\text{-AP}}$ denote the number of k -term arithmetic progressions in $\llbracket N \rrbracket_p$. Then, for every fixed positive constant δ and all $p = p(N)$ satisfying $N^{-1} \log N \ll p^{k/2} \ll 1$,*

$$\lim_{N \rightarrow \infty} \frac{-\log \mathbb{P}(X \geq (1 + \delta) \mathbb{E}[X])}{N p^{k/2} \log(1/p)} = \sqrt{\delta}.$$

Observe that Theorem 1.3 shows that the upper tail probability is well approximated by the probability of appearance of an interval (or, more generally, an arithmetic progression) of length $\sqrt{\delta N^2 p^k}$ in $\llbracket N \rrbracket_p$. Since each such interval contains approximately $\delta \mathbb{E}[X]$ arithmetic progressions of length k , it is not hard to see that conditioning $\llbracket N \rrbracket_p$ on the appearance of such a set will cause the upper tail event to occur with sizable probability. Conversely, our methods may be used to prove that the upper tail event is dominated by the appearance of some set of size $(1 + o(1)) \sqrt{\delta N^2 p^k}$ that contains nearly $\delta \mathbb{E}[X]$ arithmetic progressions of length k . It seems natural to guess that each such set is, in some sense, close to an arithmetic progression. However, this is not the case, as was shown by Green–Sisask [37]. We currently do not know a structural characterisation of the sets described above, which prevents us from proving a qualitative description of the upper tail event. For further discussion, we refer to Section 10.

The second result treats the complementary range $N^{-1} \ll p^{k/2} \ll N^{-1} \log N$, where the upper tail has Poisson behaviour.

Theorem 1.4. *Let $k \geq 3$ be an integer and let $X = X_{N,p}^{k\text{-AP}}$ denote the number of k -term arithmetic progressions in $\llbracket N \rrbracket_p$. Then, for every fixed positive constant δ and all $p = p(N)$ satisfying $N^{-1} \ll p^{k/2} \ll N^{-1} \log N$,*

$$\lim_{N \rightarrow \infty} \frac{-\log \mathbb{P}(X \geq (1 + \delta) \mathbb{E}[X])}{\mathbb{E}[X]} = (1 + \delta) \log(1 + \delta) - \delta.$$

1.2. Subgraph counts in random graphs. Let $X = X_{N,p}^H$ be the number of copies of a fixed graph H in $G_{n,p}$. Note that $\mathbb{E}[X] = \Theta(n^{v_H} p^{e_H})$. Since controlling the distribution of X for completely general graphs involves many technical difficulties (see for example [12, 66]), we will restrict our attention to connected, Δ -regular graphs H . If the expected value of X is bounded, then X converges to a Poisson random variable, as was shown independently by Bollobás [11] and by Karoński–Ruciński [49]. In view of this, for the remainder of this section, we shall assume that $\mathbb{E}[X] \rightarrow \infty$, i.e., that $p^{\Delta/2} \gg n^{-1}$. As mentioned before, we will also assume that $p \rightarrow 0$; the case where $p \in (0, 1)$ is fixed, which is fundamentally different, was addressed in [21, 56, 59].

The problem of controlling the upper tail of X has a long history. A sequence of papers [46, 53, 68], culminating in the work of Janson, Oleszkiewicz, and Ruciński [44], resulted in upper and lower bounds on the logarithmic upper tail probability that differed by a multiplicative factor of $\log(1/p)$. In the case where H is a clique, Chatterjee [16] and DeMarco–Kahn [26] independently added the missing logarithmic factor to the upper bound, thus establishing the order of magnitude of the logarithmic upper tail probability. The theory of nonlinear large deviations (discussed above) provides a variational description of the dependence of the upper tail probability $\mathbb{P}(X \geq (1 + \delta) \mathbb{E}[X])$ on δ for a certain range of $p \rightarrow 0$, as established in [3, 18, 24, 25, 30]; the strongest of these results require $p \gg n^{-1/(\Delta+1)}$ for general graphs of maximal degree Δ [24, 25], and $p^{\Delta/2} \gg n^{-1/2}$ for the case where H is a cycle [3, 24] (disregarding polylogarithmic factors). The associated variational problem was solved by Lubetzky–Zhao [57] when H is a clique and by Bhattacharya, Ganguly, Lubetzky, and Zhao [9] for general H . For a more detailed overview of these techniques, we refer the reader to the book of Chatterjee [17].

The solution to the variational problem is expressed in terms of the independence polynomial of a graph. For any H , define $P_H(x) = \sum_k i_k(H) x^k$, where $i_k(H)$ is the number of independent sets of H of size k , and let $\theta = \theta(\delta)$ be the unique positive solution to $P_H(\theta) = 1 + \delta$.³ There are two constructions that yield lower bounds for the upper tail probability (see Figure 1). In both cases, one plants a ‘small’ subgraph whose presence ensures that $G_{n,p}$ contains $(1 + \delta) \mathbb{E}[X]$ copies of H with good probability. The first of these subgraphs is a clique on $\delta^{1/v_H} n p^{\Delta/2}$ vertices (as in the left side of the figure), which contains the extra $\delta \mathbb{E}[X]$ copies of H required by the upper tail event (up to lower-order corrections). The second subgraph (which is often called a ‘hub’) is a complete bipartite graph with parts of size $\theta n p^\Delta$ and $n - \theta n p^\Delta$, respectively (as in the right side of the figure); since we are implicitly assuming that $\theta n p^\Delta$ is an integer, rounding errors play a significant role here unless $n p^\Delta \gg 1$. A short calculation shows that the expected number of copies of H which intersect this graph is approximately $\delta \mathbb{E}[X]$ and thus the actual number of such copies is almost $\delta \mathbb{E}[X]$ with good probability. In both cases, the complement of the planted subgraph typically contains approximately $\mathbb{E}[X]$ copies of H . Formalising this argument, one obtains the two lower bounds

$$\mathbb{P}(X \geq (1 + \delta) \mathbb{E}[X]) \geq p^{(\delta^{2/v_H + o(1)}) n^2 p^\Delta} \quad \text{and} \quad \mathbb{P}(X \geq (1 + \delta) \mathbb{E}[X]) \geq p^{(\theta + o(1)) n^2 p^\Delta},$$

which correspond to planting the clique and the complete bipartite graph, respectively. (Recall that the latter bound is valid only when $n p^\Delta \gg 1$.)

Our main result is that, when H is not bipartite, one of the above bounds is tight in nearly the whole range of densities. When H is bipartite, we prove tight bounds on the logarithmic upper tail probability only when $p^{\Delta/2} \geq n^{-1/2 - o(1)}$.

Theorem 1.5. *Let $\Delta \geq 2$ be an integer, let H be a connected, nonbipartite, Δ -regular graph, and let $X = X_{N,p}^H$ denote the number of copies of H in $G_{n,p}$. Then for every fixed positive constant δ*

³We note that $i_0(H) = 1$ for every graph H , so that, for example, $P_{K_r}(x) = 1 + rx$.

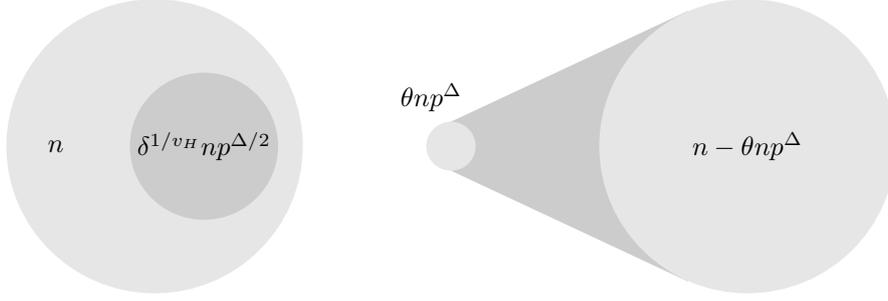


FIGURE 1. The two constructions giving lower bounds for the upper tail probability of $X = X_{n,p}^H$.

and all $p = p(n)$ satisfying $n^{-1}(\log n)^{\Delta v_H^2} \ll p^{\Delta/2} \ll 1$,

$$\lim_{n \rightarrow \infty} \frac{-\log \mathbb{P}(X \geq (1 + \delta) \mathbb{E}[X])}{n^2 p^\Delta \log(1/p)} = \begin{cases} \delta^{2/v_H} / 2 & \text{if } np^\Delta \rightarrow 0, \\ \min \{ \delta^{2/v_H} / 2, \theta \} & \text{if } np^\Delta \rightarrow \infty, \end{cases}$$

where θ is the unique positive solution to $P_H(\theta) = 1 + \delta$. Additionally, if $p^{\Delta/2} \geq n^{-1/2-o(1)}$, then the assumption that H is nonbipartite is not necessary.

We note that the theorem leaves open the case where $np^\Delta \rightarrow c \in (0, \infty)$. In this regime, the explicit dependence of the upper tail probability on δ involves various integrality conditions and is therefore quite complicated. In the next subsection, we explicitly treat this regime when H is a clique. The assumption of nonbipartiteness is not phenomenological, but only technical. The aforementioned entropic stability condition, which plays a crucial role in our proof, ceases to hold when H is bipartite as soon as $p^{\Delta/2} \leq n^{-1/2-\Theta(1)}$, see Section 10.

Our next result concerns the Poisson regime of the upper tail.

Theorem 1.6. *Let $\Delta \geq 2$ be an integer, let H be a connected, Δ -regular graph, and let $X = X_{n,p}^H$ denote the number of copies of H in $G_{n,p}$. Then, for every fixed positive constant δ and all $p = p(n)$ satisfying $n^{-1} \ll p^{\Delta/2} \ll n^{-1}(\log n)^{\frac{1}{v_H-2}}$,*

$$\lim_{n \rightarrow \infty} \frac{-\log \mathbb{P}(X \geq (1 + \delta) \mathbb{E}[X])}{\mathbb{E}[X]} = (1 + \delta) \log(1 + \delta) - \delta.$$

In the case where H is a clique, DeMarco and Kahn [26] proved that the logarithmic upper tail probability is of order $\mathbb{E}[X]$ throughout the regime covered by Theorem 1.6. For other Δ -regular graphs H , the analogous fact was known only in the range $n^{-1} \ll p^{\Delta/2} \ll n^{-1}(\log n)^{c_H}$, for some positive constant $c_H < 1/(v_H - 2)$, see [46, 62, 67, 70].

Finally, we point out that the powers of the logarithms in the assumptions of Theorems 1.5 and 1.6 do not match. After a preprint of this paper appeared, Basak and Basu [8] combined a generalised notion of entropic stability with a more refined version of the approach used in this paper to prove tight bounds on the logarithmic upper tail probability for the subgraph count of any Δ -regular graph H in the entire localised regime (see Section 3 for a more detailed discussion). Specifically, they remove the assumption that H is nonbipartite, and improve the assumed lower bound on the density p in Theorem 1.5 to $n^{-1}(\log n)^{1/(v_H-2)} \ll p^{\Delta/2}$, thus matching the assumptions of Theorem 1.6.

1.3. Clique counts in random graphs. We now consider the case of $X = X_{n,p}^H$ where H is a clique on $r \geq 3$ vertices. Thanks to the simpler structure of these graphs, we are able to prove significantly stronger results in this setting. First, we are able to determine the explicit dependence of the logarithmic upper tail probability on δ even when $np^{r-1} \rightarrow c \in (0, \infty)$. Moreover, we resolve the upper tail problem for the optimal range of densities $n^{-1}(\log n)^{\frac{1}{r-2}} \ll p^{\frac{r-1}{2}} \ll 1$,

complementing the range covered by Theorem 1.6. Finally, we give a structural description of $G_{n,p}$ conditioned on the upper tail event.

In order to formally state the theorem, it is convenient to define

$$\psi_r(\delta, c, x) = \frac{(\delta(1-x))^{2/r}}{2} + \frac{\lfloor x\delta c/r \rfloor + \{x\delta c/r\}^{\frac{1}{r-1}}}{c}, \quad (4)$$

where δ and c are nonnegative reals, $x \in [0, 1]$, and $\{a\}$ denotes the fractional part of a , and

$$\varphi_r(\delta, c) = \min_{x \in [0, 1]} \psi_r(\delta, c, x). \quad (5)$$

For an intuitive explanation of the combinatorial meaning of these definitions, we refer to the discussion at the beginning of Section 6. An easy convexity argument shows that the minimum in the definition of φ_r is attained when $x \in \{0, r\lfloor \delta c/r \rfloor / (\delta c), 1\}$, see Lemma 6.1. This leads to the explicit formula

$$\varphi_r(\delta, c) = \min \left\{ \frac{\delta^{2/r}}{2}, \frac{\lfloor \delta c/r \rfloor + \{\delta c/r\}^{1/(r-1)}}{c}, \frac{\lfloor \delta c/r \rfloor}{c} + \frac{(r\{\delta c/r\}/c)^{2/r}}{2} \right\}.$$

Theorem 1.7. *Let $r \geq 3$ be an integer and let $X = X_{n,p}^{K_r}$ denote the number of r -vertex cliques in the random graph $G_{n,p}$. Then, for every fixed positive constant δ and all $p = p(n)$ satisfying $n^{-1}(\log n)^{\frac{1}{r-2}} \ll p^{\frac{r-1}{2}} \ll 1$,*

$$\lim_{n \rightarrow \infty} \frac{-\log \mathbb{P}(X \geq (1+\delta)\mathbb{E}[X])}{n^2 p^{r-1} \log(1/p)} = \begin{cases} \delta^{2/r}/2 & \text{if } np^{r-1} \rightarrow 0, \\ \varphi_r(\delta, c) & \text{if } np^{r-1} \rightarrow c \in (0, \infty), \\ \min\{\delta^{2/r}/2, \delta/r\} & \text{if } np^{r-1} \rightarrow \infty. \end{cases}$$

Our next result describes the typical structure of the random graph $G_{n,p}$ conditioned upon the upper tail event. We write $G_{n,p}[U]$ for the subgraph of $G_{n,p}$ induced by U and $e(A, B)$ for the number of edges in $G_{n,p}$ with one endpoint in A and another in B . Define the following three events:

- (i) Let $\text{UT}(\delta)$ be the upper tail event $\{X \geq (1+\delta)\mathbb{E}[X]\}$.
- (ii) Let $\text{Clique}_\varepsilon(x)$ be the event that $G_{n,p}$ contains a set $U \subseteq [n]$ of size at least $(1-\varepsilon)x^{1/r}np^{(r-1)/2}$ such that $G_{n,p}[U]$ has minimum degree at least $(1-\varepsilon)|U|$.
- (iii) Let $\text{Hub}_\varepsilon(x)$ be the event that $G_{n,p}$ contains a set $U \subseteq [n]$ such that at least $\lfloor (1-\varepsilon)|U| \rfloor$ vertices in U have degree at least $(1-\varepsilon)n$ and

$$e(U, [n] \setminus U) \geq (1-\varepsilon)n(\lfloor xnp^{r-1}/r \rfloor + \{xnp^{r-1}/r\}^{\frac{1}{r-1}}).$$

Observe that $\text{Clique}_\varepsilon(0)$ and $\text{Hub}_\varepsilon(0)$ hold vacuously.

Theorem 1.8. *Let $r \geq 3$ be an integer and let δ , ε , and c be fixed positive constants. The following holds for all $p = p(n)$ satisfying $n^{-1}(\log n)^{\frac{1}{r-2}} \ll p^{\frac{r-1}{2}} \ll 1$.*

- (i) *If $np^{r-1} \rightarrow 0$, then*

$$\mathbb{P}(\text{Clique}_\varepsilon(\delta) \mid \text{UT}(\delta)) \rightarrow 1.$$

- (ii) *If $np^{r-1} \rightarrow c$, then, letting $x^* = r\lfloor \delta c/r \rfloor / (\delta c)$,*

$$\mathbb{P}\left(\bigcup_{x \in \{0, x^*, 1\}} (\text{Clique}_\varepsilon(\delta(1-x)) \cap \text{Hub}_\varepsilon(\delta x)) \mid \text{UT}(\delta)\right) \rightarrow 1,$$

Moreover, if $x \mapsto \psi_r(\delta, c, x)$ has a unique minimiser $x \in \{0, x^, 1\}$, then*

$$\mathbb{P}(\text{Clique}_\varepsilon(\delta(1-x)) \cap \text{Hub}_\varepsilon(\delta x) \mid \text{UT}(\delta)) \rightarrow 1.$$

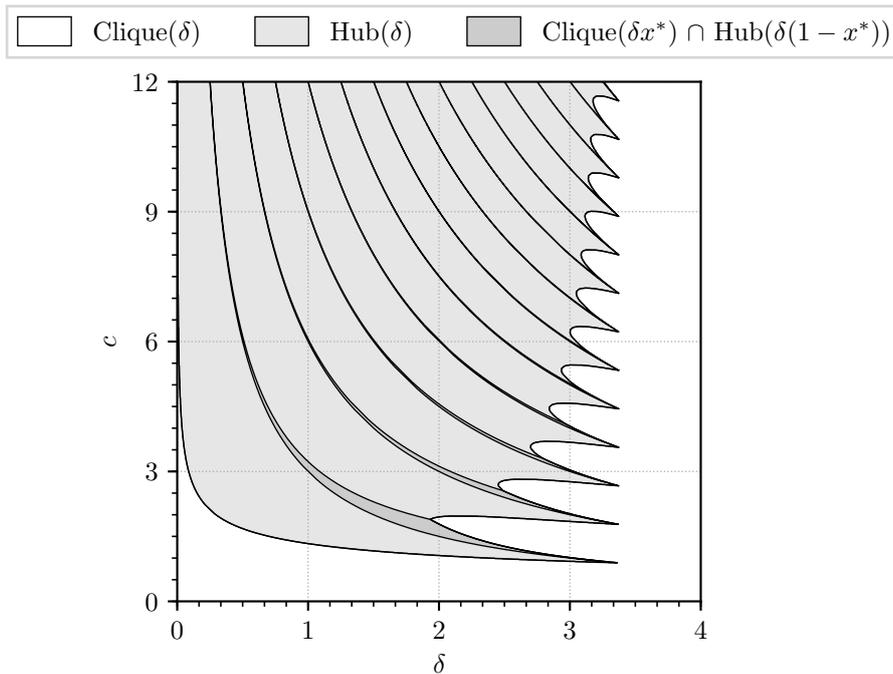


FIGURE 2. Asymptotic structure of $G_{n,p}$ conditioned upon the upper tail event $\text{UT}(\delta) = \{X_{n,p}^{K_3} \geq (1 + \delta) \mathbb{E}[X_{n,p}^{K_3}]\}$ when $np^2 \rightarrow c$ as $n \rightarrow \infty$. In this conditional model, it is highly probable that we observe either $\text{Clique}_\varepsilon(\delta)$, $\text{Hub}_\varepsilon(\delta)$, or $\text{Clique}_\varepsilon(\delta x^*) \cap \text{Hub}_\varepsilon(\delta(1 - x^*))$ with $x^* = 3\lfloor \delta c/3 \rfloor / (\delta c) \in (0, 1)$, depending on the values of δ and c (all regions are open).

(iii) If $np^{r-1} \rightarrow \infty$, then

$$\mathbb{P}(\text{Clique}_\varepsilon(\delta) \cup \text{Hub}_\varepsilon(\delta) \mid \text{UT}(\delta)) \rightarrow 1.$$

Moreover,

$$\mathbb{P}(\text{Clique}_\varepsilon(\delta) \mid \text{UT}(\delta)) \rightarrow \begin{cases} 1 & \text{if } \delta^{2/r}/2 < \delta/r, \\ 0 & \text{if } \delta^{2/r}/2 > \delta/r. \end{cases}$$

Note that Theorem 1.8 remains agnostic about the exact structure of the conditional model in the case where there are multiple minimisers to $x \mapsto \psi_r(\delta, np^{r-1}, x)$. However, it is not too difficult to show that for every r , the set of $(\delta, c) \in (0, \infty)^2$ for which $x \mapsto \psi_r(\delta, c, x)$ has multiple minimisers has Lebesgue measure zero. Figure 2 gives a graphical representation of the assertion of Theorem 1.8 in the case where $r = 3$ and $np^2 \rightarrow c$. As the figure illustrates, the conditional model undergoes infinitely many phase transitions if $\delta^{2/3}/2 > \delta/3$ (that is, $\delta < 3.375$) and no phase transition at all if $\delta^{2/3}/2 < \delta/3$.

1.4. Organisation of the paper. In Section 2, we present a short and self-contained solution to the upper tail problem for triangle counts in $G_{n,p}$. This section is somewhat redundant, since its content is just a special case of the more general Proposition 6.4. We include it in order to demonstrate our method in a simple setting that conveniently avoids many technical complications that arise in the general case.

Section 3 introduces a concentration inequality that gives a general condition under which the logarithmic upper tail probability can be approximated by $\Phi_X(\delta)$, the solution to the optimisation problem (2).

In Section 4, we use the inequality developed in Section 3 to determine the asymptotics of the logarithmic upper tail probability of $X_{N,p}^{k\text{-AP}}$ in the complete range of densities where localisation occurs. After collecting some graph-theoretic tools in Section 5, we study the localised regime of the upper tails of $X_{n,p}^{K_r}$ and $X_{n,p}^H$ for connected, Δ -regular graphs H in Sections 6 and 7, respectively. We note that the three Sections 4, 6, and 7 are logically independent and may be read in any order; however, both Sections 6 and 7 rely on the tools of Section 5.

In Section 8, we prove various results related to the Poisson regime; in particular, we give the proofs of Theorems 1.4 and 1.6. The arguments we use there do not rely on the methods developed in Section 3, but rather on explicit calculations of high factorial moments. Section 9 contains a brief discussion on extending the result from Section 3 to the more general case of nonnegative random variables on the hypercube. Finally, in Section 10 we make some concluding remarks and discuss open problems.

1.5. Notation. Before ending the introduction, we collect some notation which will be used throughout the paper. We write K_r for the complete graph on r vertices and $K_{s,t}$ for the complete bipartite graph with parts of size s and t . For any graph G , let $V(G)$ and $E(G)$ denote the vertex and edge sets of G , respectively, and set $v_G = |V(G)|$ and $e_G = |E(G)|$. For two graphs J and G , we let $N(J, G)$ be the number of copies of J in G , and $\text{Emb}(J, G)$ be the set of embeddings of J into G —i.e., injective maps from $V(J)$ to $V(G)$ that map edges of J to edges of G . For an edge $uv \in E(G)$, we also let $N(J, G; uv)$ and $\text{Emb}(J, G; uv)$ be the number of copies of J that contain the edge uv , and the set of embeddings that map an edge of J to uv , respectively. Finally, for a subset I of $\llbracket N \rrbracket$, we let $Y_I = \prod_{i \in I} Y_i$, and $\mathbb{E}_I[X] = \mathbb{E}[X \mid Y_I = 1]$. If subsets $I \subseteq \llbracket N \rrbracket$ can be identified with subgraphs $G \subseteq K_n$, as in Example 1.2, we will write $\mathbb{E}_G[X]$ instead of $\mathbb{E}_I[X]$.

2. TRIANGLES IN RANDOM GRAPHS

Assume that $n^{-1} \log n \ll p \ll 1$ and let X denote the number of triangles in $G_{n,p}$. Using the shorthand notation $\mathbb{E}_G[X] = \mathbb{E}[X \mid G \subseteq G_{n,p}]$, we define, for each positive δ ,

$$\Phi_X(\delta) = \min \{e_G \log(1/p) : G \subseteq K_n \text{ and } \mathbb{E}_G[X] \geq (1 + \delta) \mathbb{E}[X]\}. \quad (6)$$

Note that this agrees with the definition (2). The goal of this section is to prove that, for every fixed positive ε and all large enough n ,

$$(1 - \varepsilon) \Phi_X(\delta - \varepsilon) \leq -\log \mathbb{P}(X \geq (1 + \delta) \mathbb{E}[X]) \leq (1 + \varepsilon) \Phi_X(\delta + \varepsilon). \quad (7)$$

At this point, we do not address the issue of evaluating $\Phi_X(\delta)$. For the sake of completeness, let us mention that a special case of a more general result of Lubetzky–Zhao [57] is that, when $n^{-1} \ll p \ll 1$,

$$\lim_{n \rightarrow \infty} \frac{\Phi_X(\delta)}{n^2 p^2 \log(1/p)} = \begin{cases} \delta^{2/3}/2 & \text{if } np^2 \rightarrow 0 \\ \min \{\delta/3, \delta^{2/3}/2\} & \text{if } np^2 \rightarrow \infty. \end{cases}$$

In Section 6, we shall fill in the gap at $p = \Theta(n^{-1/2})$ to obtain an asymptotic formula for $\Phi_X(\delta)$ in the full range of interest.

We now give a proof of (7), where we may assume without loss of generality that $\varepsilon \leq \delta/10$. All statements that we make in this section should be understood to hold only for sufficiently large n . We start with some easy observations. First, for every graph G with $O(1)$ edges, we have $\mathbb{E}_G[X] - \mathbb{E}[X] = O(1 + np^2) \ll (\delta - 2\varepsilon) \mathbb{E}[X]$, and so $\Phi_X(\delta - 2\varepsilon) \gg \log(1/p) \gg 1$. Second, the condition in (6) is satisfied when G is a clique on $(1 + \delta)^{1/3} np$ vertices, and therefore $\Phi_X(\delta - \varepsilon) \leq (1 + \delta)^{2/3} n^2 p^2 \log(1/p)/2$.

The easier of the two inequalities in (7) is the upper bound. To prove it, let G be a graph attaining the minimum in the definition of $\Phi_X(\delta + \varepsilon)$ and let $\mathbb{P}_G(\cdot) = \mathbb{P}(\cdot \mid G \subseteq G_{n,p})$. Since X

never exceeds $\binom{n}{3}$,

$$(1 + \delta + \varepsilon) \mathbb{E}[X] \leq \mathbb{E}_G[X] \leq \binom{n}{3} \mathbb{P}_G(X \geq (1 + \delta) \mathbb{E}[X]) + (1 + \delta) \mathbb{E}[X],$$

and so

$$\mathbb{P}_G(X \geq (1 + \delta) \mathbb{E}[X]) \geq \varepsilon \mathbb{E}[X] / \binom{n}{3} = \varepsilon p^3.$$

Hence,

$$\begin{aligned} -\log \mathbb{P}(X \geq (1 + \delta) \mathbb{E}[X]) &\leq -\log \left(\mathbb{P}(G \subseteq G_{n,p}) \cdot \mathbb{P}_G(X \geq (1 + \delta) \mathbb{E}[X]) \right) \\ &\leq e_G \log(1/p) + \log(1/\varepsilon p^3). \end{aligned}$$

Since $e_G \log(1/p) = \Phi_X(\delta + \varepsilon) \geq \Phi_X(\delta - 2\varepsilon) \gg \log(1/p)$, this establishes the lower bound in (7).

We now turn to proving the lower bound. Let $C = C(\varepsilon, \delta)$ denote a sufficiently large positive constant. We call a graph $G \subseteq K_n$ a *seed* if

- (S1) $\mathbb{E}_G[X] \geq (1 + \delta - \varepsilon) \mathbb{E}[X]$ and
- (S2) $e_G \leq Cn^2 p^2 \log(1/p)$.

We make the following claim:

Claim 2.1. $\mathbb{P}(X \geq (1 + \delta) \mathbb{E}[X]) \leq (1 + \varepsilon) \mathbb{P}(G_{n,p} \text{ contains a seed})$.

Remark 2.2. Given this claim, one may be tempted to simply apply the union bound over all seeds. Using such a strategy, one would find that

$$\begin{aligned} \mathbb{P}(X \geq (1 + \delta) \mathbb{E}[X]) &\stackrel{\text{Claim 2.1}}{\leq} (1 + \varepsilon) \mathbb{P}(G_{n,p} \text{ contains a seed}) \\ &\leq (1 + \varepsilon) \sum_{m=m_{\min}}^{\infty} p^m \cdot |\{G \subseteq K_n : G \text{ is a seed with } m \text{ edges}\}|, \end{aligned} \quad (8)$$

where $m_{\min} = \Phi_X(\delta - \varepsilon) / \log(1/p)$ is the minimal number of edges in a seed. Unfortunately, such a strategy is bound to fail, as there are far too many seeds. To see this, we observe that if G satisfies (S1), then so does any supergraph of G . In particular, if we take a seed with m_{\min} edges and add an arbitrary set of Km_{\min} edges (where K is a large constant), the resulting graph remains a seed. Therefore, we can (rather loosely) bound the number of seeds with $\tilde{m} = (K + 1)m_{\min}$ edges from below:

$$|\{G \subseteq K_n : G \text{ is a seed with } \tilde{m} \text{ edges}\}| \geq \binom{\binom{n}{2} - m_{\min}}{Km_{\min}} \geq \left(\frac{1}{3(1 + \delta)^{2/3} K p^2} \right)^{Km_{\min}},$$

using the bound $\binom{x}{y} \geq (x/2y)^y$ for any $y < x/2$ and $m_{\min} \leq (1 + \delta)^{2/3} n^2 p^2 / 2$. If we choose K to be large enough, we may conclude that

$$|\{G \subseteq K_n : G \text{ is a seed with } \tilde{m} \text{ edges}\}| \geq (1/p)^{3\tilde{m}/2},$$

for all large enough n . This shows that the \tilde{m} th term of the final sum in (8) is arbitrarily large, making the entire approach fruitless.

From Remark 2.2, it is clear that seeds may include ‘extraneous’ edges that have no structural role in the upper tail event. We wish to consider more constrained structures which exclude constructions like the one outlined above. To that end, we call a graph $G^* \subseteq K_n$ a *core* if

- (C1) $\mathbb{E}_{G^*}[X] \geq (1 + \delta - 2\varepsilon) \mathbb{E}[X]$,
- (C2) $e_{G^*} \leq Cn^2 p^2 \log(1/p)$, and
- (C3) $\min_{e \in E(G^*)} (\mathbb{E}_{G^*}[X] - \mathbb{E}_{G^* \setminus e}[X]) \geq \varepsilon \mathbb{E}[X] / (Cn^2 p^2 \log(1/p))$.

Condition (C3) requires that every edge contributed meaningfully to the expectation, and thus excludes the pathological seeds described above.

Claim 2.3. *Every seed contains a core.*

Finally, we claim that the additional constraint (C3) allows us to get a very strong bound on the number of cores with a given number of edges, as it ensures that either the product of the degrees of the endpoints of any edge in G^* must be nearly as large as e_{G^*} , or the sum of these degrees is nearly linear in n (note that the former condition holds when G^* is a clique, and the latter when G^* is a hub). This, in turn, implies that the number of cores with a given set of vertices and m edges is $\exp(O(m \log \log(1/p)))$. Once we show that the number of choices for the vertex set of a core with m edges is $(1/p)^{o(m)}$, we prove the following claim.

Claim 2.4. *For every m , there are at most $(1/p)^{\varepsilon m}$ cores with exactly m edges.*

As will be discussed in Section 3, we refer to the property described in Claim 2.4 as *entropic stability*.

Let us show how these three claims imply the lower bound in (7):

$$\begin{aligned} \mathbb{P}(X \geq (1 + \delta) \mathbb{E}[X]) &\stackrel{\text{Claim 2.1}}{\leq} (1 + \varepsilon) \mathbb{P}(G_{n,p} \text{ contains a seed}) \\ &\stackrel{\text{Claim 2.3}}{\leq} (1 + \varepsilon) \mathbb{P}(G_{n,p} \text{ contains a core}) \\ &\leq (1 + \varepsilon) \sum_{m=0}^{\infty} p^m \cdot |\{G^* \subseteq K_n : G^* \text{ is a core with } m \text{ edges}\}| \\ &\stackrel{\text{Claim 2.4}}{\leq} (1 + \varepsilon) \sum_{m=m_{\min}}^{\infty} p^{(1-\varepsilon)m}, \end{aligned}$$

where m_{\min} is the minimal number of edges in a core. Since (C1) implies that $\Phi_X(\delta - 2\varepsilon) \leq m_{\min} \cdot \log(1/p)$, the assumption $p \ll 1$ yields

$$\mathbb{P}(X \geq (1 + \delta) \mathbb{E}[X]) \leq (1 + 2\varepsilon) \exp(- (1 - \varepsilon) \Phi_X(\delta - 2\varepsilon)).$$

Finally, as $\Phi_X(\delta - 2\varepsilon) \gg 1$, we obtain

$$-\log \mathbb{P}(X \geq (1 + \delta) \mathbb{E}[X]) \geq (1 - 2\varepsilon) \Phi_X(\delta - 2\varepsilon),$$

thus proving (7) with 2ε instead of ε .

Remark 2.5. Conditions (S2) and (C2) require both seeds and cores to have no more than $Cn^2 p^2 \log(1/p)$ edges, forcing us to consider graphs which are larger than the minimiser of $\Phi_X(\delta)$ by a multiplicative factor of $\log(1/p)$. This is optimal in the following sense. On the one hand, replacing $\log(1/p)$ with a larger function would weaken the conclusion of Claim 2.1 and make Claim 2.4 hold only for a smaller range of densities p . On the other hand, if we could replace $\log(1/p)$ by $\ell(p) = o(\log(1/p))$ without altering the validity of Claim 2.1, then Claim 2.4 (with an appropriately adjusted definition of a core) would hold for any $p \gg \ell(p)/n$; this would imply an upper bound on the upper tail probability that is stronger than the lower bound proven in Theorem 1.6 for $\ell(p)/n \ll p \ll \log n/n$.

Proof of Claim 2.1. We refine a classical argument due to Janson, Oleszkiewicz, and Ruciński [44]. Let Z be the indicator random variable of the event that $G_{n,p}$ does not contain a seed and let $\ell = \lceil (C/3)n^2 p^2 \log(1/p) \rceil$. Since $XZ \geq 0$ and $Z^\ell = Z$, Markov's inequality gives

$$\mathbb{P}(X \geq (1 + \delta) \mathbb{E}[X] \text{ and } G_{n,p} \text{ contains no seed}) = \mathbb{P}(XZ \geq (1 + \delta) \mathbb{E}[X]) \leq \frac{\mathbb{E}[X^\ell Z]}{(1 + \delta)^\ell \mathbb{E}[X]^\ell}. \quad (9)$$

We write $X = \sum_T Y_T$, where the sum ranges over all triangles T in K_n and Y_T is the indicator random variable of the event that T is contained in $G_{n,p}$. For every subgraph $G \subseteq K_n$, let Z_G be the indicator random variable of the event that $G \cap G_{n,p}$ does not contain a seed. Observe that

$Z_{G'} \leq Z_G$ whenever $G \subseteq G'$. In particular, since $Z = Z_{K_n}$, we have, for every $k \in \llbracket \ell \rrbracket$,

$$\begin{aligned} \mathbb{E}[X^k Z] &= \sum_{T_1, \dots, T_k} \mathbb{E}[Y_{T_1} \cdots Y_{T_k} \cdot Z] \\ &\leq \sum_{T_1, \dots, T_k} \mathbb{E}[Y_{T_1} \cdots Y_{T_k} \cdot Z_{T_1 \cup \dots \cup T_k}] \\ &\leq \sum_{T_1, \dots, T_{k-1}} \mathbb{E}[Y_{T_1} \cdots Y_{T_{k-1}} \cdot Z_{T_1 \cup \dots \cup T_{k-1}}] \cdot \sum_{T_k} \mathbb{E}[Y_{T_k} \mid Y_{T_1} \cdots Y_{T_{k-1}} \cdot Z_{T_1 \cup \dots \cup T_{k-1}} = 1], \end{aligned}$$

where we can let the first sum in the last line range only over sequences T_1, \dots, T_{k-1} for which the event $\{Y_{T_1} \cdots Y_{T_{k-1}} \cdot Z_{T_1 \cup \dots \cup T_{k-1}} = 1\}$ has positive probability. This is equivalent to saying that the graph $T_1 \cup \dots \cup T_{k-1}$ does not contain a seed and thus $Y_{T_1} \cdots Y_{T_{k-1}} \cdot Z_{T_1 \cup \dots \cup T_{k-1}} = Y_{T_1} \cdots Y_{T_{k-1}}$. Moreover, since $e_{T_1 \cup \dots \cup T_{k-1}} \leq 3(k-1) \leq 3(\ell-1) \leq Cn^2 p^2 \log(1/p)$,

$$\sum_{T_k} \mathbb{E}[Y_{T_k} \mid Y_{T_1} \cdots Y_{T_{k-1}} = 1] = \mathbb{E}_{T_1 \cup \dots \cup T_{k-1}}[X] < (1 + \delta - \varepsilon) \mathbb{E}[X],$$

as otherwise $T_1 \cup \dots \cup T_{k-1}$ would be a seed, see (S1) and (S2). Therefore,

$$\sum_{T_1, \dots, T_k} \mathbb{E}[Y_{T_1} \cdots Y_{T_k} \cdot Z_{T_1 \cup \dots \cup T_k}] < (1 + \delta - \varepsilon) \mathbb{E}[X] \cdot \sum_{T_1, \dots, T_{k-1}} \mathbb{E}[Y_{T_1} \cdots Y_{T_{k-1}} \cdot Z_{T_1 \cup \dots \cup T_{k-1}}]$$

and it follows by induction that $\mathbb{E}[X^\ell Z] < (1 + \delta - \varepsilon)^\ell \mathbb{E}[X]^\ell$. Substituting this inequality into (9) gives

$$\mathbb{P}(X \geq (1 + \delta) \mathbb{E}[X] \text{ and } G_{n,p} \text{ contains no seed}) \leq \left(\frac{1 + \delta - \varepsilon}{1 + \delta} \right)^\ell.$$

Since the probability that $G_{n,p}$ contains a seed is at least $e^{-\Phi_X(\delta - \varepsilon)}$, the probability that $G_{n,p}$ contains a given seed of smallest size, the bounds $1 \ll \Phi_X(\delta - \varepsilon) \leq (1 + \delta)^{2/3} n^2 p^2 \log(1/p)/2$ imply that, for all sufficiently large n ,

$$\left(\frac{1 + \delta - \varepsilon}{1 + \delta} \right)^\ell \leq \left(\frac{1 + \delta - \varepsilon}{1 + \delta} \right)^{(C/3)n^2 p^2 \log(1/p)} \leq \varepsilon e^{-\Phi_X(\delta - \varepsilon)} \leq \varepsilon \mathbb{P}(G_{n,p} \text{ contains a seed})$$

whenever the constant C is sufficiently large⁴. This implies the assertion of the claim. \square

Proof of Claim 2.3. Let G be a seed. Define a sequence $G = G_0 \supseteq G_1 \supseteq \dots \supseteq G_s = G^*$ of subgraphs of G by repeatedly setting $G_{i+1} = G_i \setminus e$ for some edge $e \in G_i$ such that $\mathbb{E}_{G_i} - \mathbb{E}_{G_i \setminus e}[X] < \varepsilon \mathbb{E}[X]/(Cn^2 p^2 \log(1/p))$, as long as such an edge e exists. By construction, G^* clearly satisfies (C3). Since $e_{G^*} \leq e_G \leq Cn^2 p^2 \log(1/p)$, we see that (C2) holds as well. Finally, as $s \leq e_G \leq Cn^2 p^2 \log(1/p)$, we have

$$\mathbb{E}_G[X] - \mathbb{E}_{G^*}[X] = \sum_{i=0}^{s-1} (\mathbb{E}_{G_i}[X] - \mathbb{E}_{G_{i+1}}[X]) < \varepsilon \mathbb{E}[X].$$

Since G is a seed, $\mathbb{E}_G[X] \geq (1 + \delta - \varepsilon) \mathbb{E}[X]$ and we obtain (C1). \square

Proof of Claim 2.4. We bound the number of cores with m edges from above. This number is zero whenever $m > Cn^2 p^2 \log(1/p)$, by (C2). We may thus assume that $m \leq Cn^2 p^2 \log(1/p)$. Given a core G^* , we denote by A_{G^*} the set of vertices of G^* with degree at least $\varepsilon n p / (30C \log(1/p))$ and by $B_{G^*} \subseteq A_{G^*}$ the set of vertices of G^* with degree at least $\varepsilon n / (30C \log(1/p))$, where the degree is taken in G^* . Since G^* has m edges,

$$|A_{G^*}| \leq a := \frac{60Cm \log(1/p)}{\varepsilon n p} \quad \text{and} \quad |B_{G^*}| \leq b := \frac{60Cm \log(1/p)}{\varepsilon n}.$$

The key observation, which we will now verify, is that every edge of G^* is either contained in A_{G^*} or has an endpoint in B_{G^*} , see Figure 3 for an illustration. To see this, consider some edge

⁴We note that the requirement for C to be large is only used here.

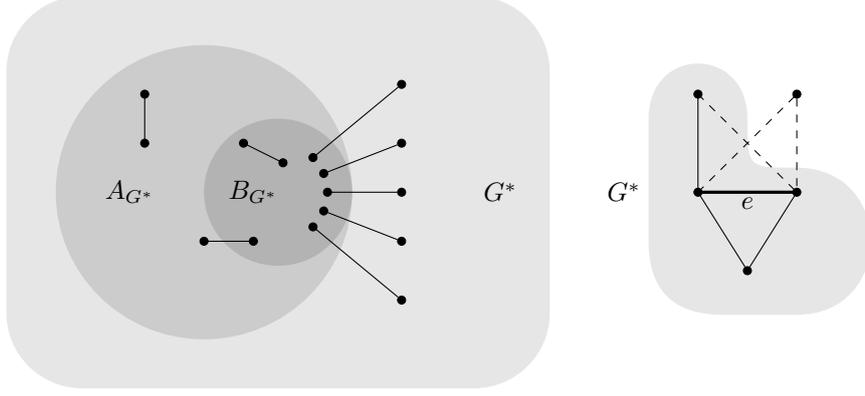


FIGURE 3. Left: The sets A_{G^*} and B_{G^*} of high-degree vertices capture the edges of the core. Right: Three different types of triangles containing an edge e in the core G^* .

$e \in E(G^*)$. For every nonempty graph $F \subseteq K_3$, let $N(F, G^*; e)$ denote the number of copies of F in G^* that contain e . By considering how the $n - 2$ triangles in K_n that contain e intersect G^* (see Figure 3), one can see that

$$\mathbb{E}_{G^*}[X] - \mathbb{E}_{G^* \setminus e}[X] \leq (N(K_3, G^*; e) + N(K_{1,2}, G^*; e) \cdot p + np^2) \cdot (1 - p).$$

Using $\mathbb{E}_{G^*}[X] - \mathbb{E}_{G^* \setminus e}[X] \geq \varepsilon \mathbb{E}[X]/(Cn^2 p^2 \log(1/p))$ and $\mathbb{E}[X] \geq (1 - o(1))n^3 p^3/6$, we thus get

$$\frac{\varepsilon np}{7C \log(1/p)} \leq N(K_3, G^*; e) + N(K_{1,2}, G^*; e) \cdot p + np^2.$$

Since $p \ll 1$ implies that $np^2 \ll np/\log(1/p)$, we find that either

$$N(K_3, G^*; e) \geq \frac{\varepsilon np}{15C \log(1/p)} \quad \text{or} \quad N(K_{1,2}, G^*; e) \geq \frac{\varepsilon n}{15C \log(1/p)}. \quad (10)$$

Since $N(K_3, G^*; uv) \leq \min\{\deg_{G^*} u, \deg_{G^*} v\}$ and $N(K_{1,2}, G^*; uv) \leq \deg_{G^*} u + \deg_{G^*} v$, the first inequality in (10) implies that e is contained in A_{G^*} whereas the second inequality implies that e has an endpoint in B_{G^*} , as claimed.

Our key observation implies that for fixed sets $B \subseteq A \subseteq \llbracket n \rrbracket$ with $|A| = a$ and $|B| = b$, there are at most $\binom{a^2/2 + bn}{m}$ cores G^* with m edges that satisfy $A_{G^*} \subseteq A$ and $B_{G^*} \subseteq B$. We can thus (generously) upper bound the number of cores with m edges by

$$\binom{n}{a} \binom{n}{b} \binom{a^2/2 + bn}{m}.$$

Recalling the inequality $\binom{x}{y} \leq (ex/y)^y$, valid for all nonnegative integers x and y , we may conclude that the number of cores with m edges is at most

$$n^{\frac{120Cm \log(1/p)}{\varepsilon np}} \cdot \left(\frac{e(60C)^2 m (\log(1/p))^2}{2\varepsilon^2 n^2 p^2} + \frac{e60C \log(1/p)}{\varepsilon} \right)^m.$$

Since $p \gg n^{-1} \log n$, the first factor is at most $e^{o(m \log(1/p))}$. Using $m \leq Cn^2 p^2 \log(1/p)$, the second factor is at most $e^{O(m \log \log(1/p))} = e^{o(m \log(1/p))}$. This shows that the number of cores with m edges is indeed at most $(1/p)^{\varepsilon m}$, as claimed. \square

3. THE MAIN TECHNICAL RESULT: ‘ENTROPIC STABILITY IMPLIES LOCALISATION’

The goal of this section is to state a general result that allows one, in many cases of interest, to reduce the problem of determining the precise asymptotics of the logarithmic upper tail probability of a polynomial (with nonnegative coefficients) of independent Bernoulli random variables to a

counting problem. Since the main technical lemmas also apply to non-product measures on the hypercube, we phrase the basic definitions in this broader context.

We denote by Y a random variable taking values in the discrete N -dimensional cube $\{0, 1\}^N$ and by $X = X(Y)$ a real-valued, increasing function of Y with positive expectation. Given a subset $I \subseteq \llbracket N \rrbracket$, we write $Y_I = \prod_{i \in I} Y_i$ for the indicator random variable of the event $\{Y_i = 1 \text{ for all } i \in I\}$. Using the shorthand notation $\mathbb{E}_I[X] = \mathbb{E}[X \mid Y_I = 1]$,⁵ we define a function $\Phi_X: \mathbb{R} \rightarrow \mathbb{R}_{\geq 0} \cup \{\infty\}$ by⁶

$$\Phi_X(\delta) = \min \{ -\log \mathbb{P}(Y_I = 1) : I \subseteq \llbracket N \rrbracket \text{ and } \mathbb{E}_I[X] \geq (1 + \delta) \mathbb{E}[X] \}. \quad (11)$$

It is easy to see that Φ_X is a nondecreasing function satisfying $\Phi_X(\delta) > 0$ for all $\delta > 0$. We say that a function $X: \{0, 1\}^N \rightarrow \mathbb{R}_{\geq 0}$ is a polynomial with nonnegative coefficients and degree at most d if it admits a representation $X = \sum_{I \subseteq \llbracket N \rrbracket} \alpha_I Y_I$, where each coefficient α_I is nonnegative and $\alpha_I = 0$ whenever $|I| > d$.

Let \mathcal{I} be a collection of subsets $I \subseteq \llbracket N \rrbracket$. Given $\varepsilon > 0$ and $p > 0$, we say that \mathcal{I} is an *entropically stable* family (with respect to ε and p) if, for every integer m ,

$$|\{I \in \mathcal{I} : |I| = m\}| \leq (1/p)^{\varepsilon m/2}.$$

For the sake of brevity, we will suppress the dependence of this property on ε and p .

The following statement is the main technical result of our work.

Theorem 3.1. *For every positive integer d and all positive real numbers ε and δ with $\varepsilon < 1/2$, there is a positive $K = K(d, \varepsilon, \delta)$ such that the following holds. Let Y be a sequence of N independent $\text{Ber}(p)$ random variables for some $p \in (0, 1 - \varepsilon]$ and let $X = X(Y)$ be a nonzero polynomial with nonnegative coefficients and degree at most d such that $\Phi_X(\delta - \varepsilon) \geq K \log(1/p)$. Denote by \mathcal{I}^* the collection of all subsets $I \subseteq \llbracket N \rrbracket$ satisfying*

- (C1) $\mathbb{E}_I[X] \geq (1 + \delta - \varepsilon) \mathbb{E}[X]$,
- (C2) $|I| \leq K \cdot \Phi_X(\delta + \varepsilon)$, and
- (C3) $\min_{i \in I} (\mathbb{E}_I[X] - \mathbb{E}_{I \setminus \{i\}}[X]) \geq \mathbb{E}[X] / (K \cdot \Phi_X(\delta + \varepsilon))$,

and assume \mathcal{I}^* is an entropically stable family. Then

$$(1 - \varepsilon) \Phi_X(\delta - \varepsilon) \leq -\log \mathbb{P}(X \geq (1 + \delta) \mathbb{E}[X]) \leq (1 + \varepsilon) \Phi_X(\delta + \varepsilon) \quad (12)$$

and, writing \mathcal{J}^* for the collection of those $I \in \mathcal{I}^*$ with $-\log \mathbb{P}(Y_I = 1) \leq (1 + \varepsilon) \Phi_X(\delta + \varepsilon)$,

$$\mathbb{P}(X \geq (1 + \delta) \mathbb{E}[X] \text{ and } Y_I = 0 \text{ for all } I \in \mathcal{J}^*) \leq \varepsilon \mathbb{P}(X \geq (1 + \delta) \mathbb{E}[X]). \quad (13)$$

Remark 3.2. Observe that (12) gives tight bounds on the logarithmic upper tail probability of X , provided that $\Phi_X(\delta) = (I(\delta) + o(1))f(N, p)$ for a continuous, positive function I and some function f . Equation (13) states that the upper tail event is (almost) contained in the event that $Y_I = 1$ for some $I \in \mathcal{J}^*$; note that each such I is a near-minimiser of $\Phi_X(\delta - \varepsilon)$. In some cases, it is possible to classify these near-minimisers and thereby obtain a rough structural characterisation of the upper tail event.

Remark 3.3. The assumption $\Phi_X(\delta - \varepsilon) \geq K \log(1/p)$ means that conditioning on $Y_I = 1$ for any constant-size subset $I \subseteq \llbracket N \rrbracket$ cannot increase the expected value of X by more than $(\delta - \varepsilon) \mathbb{E}[X]$; it is very easy to verify this for the applications we have in mind. The more onerous task is verifying that \mathcal{I}^* is an entropically stable family. In fact, a large part of this paper is dedicated to counting *cores* (as we call the elements of \mathcal{I}^*). Frequently, there are very few minimisers of $\Phi_X(\delta)$, for every $\delta > 0$. Entropic stability quantifies the notion that there are few near-minimisers as well.

⁵Strictly speaking, $\mathbb{E}_I[X]$ is well defined only if $\mathbb{P}(Y_I = 1) > 0$. However, the value of $\mathbb{E}_I[X]$ for sets I with $\mathbb{P}(Y_I = 1) = 0$ does not affect any of our statements.

⁶We use the standard convention that $\min \emptyset = \infty$.

Remark 3.4. In the following sections, we will compute the logarithmic upper tail probabilities in various settings by estimating the function Φ_X and verifying that the random variable X in question satisfies the assumptions of Theorem 3.1. As will be seen in the proof of Theorem 3.1, entropic stability implies that

$$-\log \mathbb{P}(Y_{I^*} = 1 \text{ for some } I^* \in \mathcal{I}^*) \geq (1 - \varepsilon/2) \cdot \Phi_X(\delta - \varepsilon). \quad (14)$$

However, there are many natural contexts where the entropic stability assumption fails despite the fact that (14) remains true. For example, when $H = C_4$, then every copy of the complete bipartite graph $K_{2, cn^2 p^2}$ in K_n is a core, provided that c is large enough. There are $\binom{n}{cn^2 p^2}$ such copies in K_n and $\binom{n}{cn^2 p^2}$ is larger than $(1/p)^{\Omega(n^2 p^2)}$ when $p \ll n^{-1/2-\xi}$ for some small $\xi > 0$. In order to study such scenarios, one may search out a weaker condition which still implies (14), and employ Theorem 3.1 *mutatis mutandis*. One such modification is to allow the number of cores with m edges to be as large as $(1/p)^{m - m_{\min} + o(m_{\min})}$; such a generalisation was introduced in the work of Basak–Basu [8].

Remark 3.5. Let $\mathcal{M}(\{0, 1\}^N)$ be the set of measures on the N -dimensional hypercube. For any random variable $X = X(Y)$ defined on the p -biased hypercube, it is possible to give an abstract description of the probability of the upper tail event via the Gibbs variational principle, which states that

$$-\log \mathbb{P}(X \geq (1 + \delta) \mathbb{E}[X]) \geq \inf_{\nu \in \mathcal{M}_\delta^X} D_{\text{KL}}(\nu \parallel \mu_p),$$

where μ_p is the product of N Bernoulli measures of mean p , $D_{\text{KL}}(\cdot \parallel \cdot)$ is the relative entropy

$$D_{\text{KL}}(\nu \parallel \mu) = \sum_{y \in \{0, 1\}^N} \nu(y) \log \left(\frac{\nu(y)}{\mu(y)} \right),$$

and

$$\mathcal{M}_\delta^X := \left\{ \nu \in \mathcal{M}(\{0, 1\}^N) : \sum_{y \in \{0, 1\}^N} X(y) \nu(y) \geq (1 + \delta) \mathbb{E}[X] \right\}.$$

The *naïve mean field approximation* holds for the upper tail of X if one can replace the infimum over all measures in \mathcal{M}_δ^X by an infimum over all *product* measures in \mathcal{M}_δ^X , while incurring only lower order errors.

The seminal work of Chatterjee–Dembo [17], further developed by Eldan [30] and Augeri [3], provides very general sufficient conditions that imply the naïve mean field approximation for a general function f of Bernoulli random variables, stated in terms of the ‘smoothness’ of f and the ‘complexity’ of its gradient. Although the various works consider different notions of low complexity, all of them seem to imply the heuristic statement that the set of all directions of the gradient of f is an extremely sparse subset of the $(N - 1)$ -dimensional sphere. An alternate approach to the naïve mean field approximation is used by Cook–Dembo [24], which specializes to the case of subgraph counts in Erdős–Rényi random graphs. Instead of appealing to gradient complexity bounds, they construct an efficient covering of (most of) the hypercube by convex sets on which the subgraph counts are nearly constant, in the spirit of the regularity-based approach of Chatterjee–Varadhan [21].

Although its formulation is rather different, Theorem 3.1 can also be considered in the context of the naïve mean field approximation, coverings of the hypercube, and low-complexity gradients. Given a subset $I \subseteq \llbracket N \rrbracket$, we define a product measure ν_I by

$$\nu_I(Y_i = 1) = \begin{cases} 1 & \text{if } i \in I, \\ p & \text{otherwise.} \end{cases}$$

A straightforward computation shows that $D_{\text{KL}}(\nu_I \parallel \mu_p) = |I| \log(1/p)$, and so

$$\Phi_X(\delta) = \inf \left\{ D_{\text{KL}}(\nu_I \parallel \mu_p) : \nu_I \in \mathcal{M}_\delta^X \right\}.$$

In these terms, Theorem 3.1 shows that, if \mathcal{I}^* an entropically stable family, then a particularly simple form of the naïve mean field approximation holds: one must only consider product measures that assign edges probability p or 1. Our adaptation of the high moment argument of Janson, Oleszkiewicz, and Ruciński, used in Lemma 3.7, constructs a covering of (most of) the upper tail event by a family \mathcal{I} of small subsets I with $\nu_I \in \mathcal{M}_\delta^X$. The extraction of cores from these subsets corresponds to identifying the directions in which the possible partial derivatives are large; in this perspective, entropic stability is analogous to the sparseness property that is encoded by the low-complexity gradient condition.

The final stipulation of Theorem 3.1 gives a structural description of the measure conditioned on the upper tail event. More specifically, a sample from the conditional measure will contain an element of \mathcal{J}^* with high probability. Working from the naïve mean-field approximation, one may consider the more subtle question of the exact relationship between the conditional measure and the family of product measures in \mathcal{M}_δ^X that attain the minimal relative entropy from μ_p . The work of Eldan–Gross [31] shows that, under certain conditions, the conditional measure is close to a mixture of such product measures, in the sense of optimal transport; Austin [4] proves similar results for a broader class of measures (not necessarily on the hypercube).

The upper bound on $-\log \mathbb{P}(X \geq (1 + \delta) \mathbb{E}[X])$ stated in (12) will follow easily from the following simple lemma.

Lemma 3.6. *Let Y be a random variable taking values in $\{0, 1\}^N$ and let $X = X(Y)$ be a real-valued function of Y . Suppose that $\mathbb{E}[X] > 0$ and that $X \leq M$ always. Then for all positive ε and δ ,*

$$-\log \mathbb{P}(X \geq (1 + \delta) \mathbb{E}[X]) \leq \Phi_X(\delta + \varepsilon) + \log \left(\frac{M}{\varepsilon \mathbb{E}[X]} \right).$$

Proof. Let $t = (1 + \delta) \mathbb{E}[X]$. If $\Phi_X(\delta + \varepsilon) = \infty$, then the assertion of the lemma is vacuous. Otherwise, there exists a set $I \subseteq \llbracket N \rrbracket$ with $-\log \mathbb{P}(Y_I = 1) = \Phi_X(\delta + \varepsilon)$ and $\mathbb{E}_I[X] \geq t + \varepsilon \mathbb{E}[X]$. As $\mathbb{E}_I[X] \leq M \cdot \mathbb{P}(X \geq t \mid Y_I = 1) + t$, it follows that

$$\mathbb{P}(X \geq t) \geq \mathbb{P}(Y_I = 1) \cdot \mathbb{P}(X \geq t \mid Y_I = 1) \geq \mathbb{P}(Y_I = 1) \cdot \frac{\varepsilon \mathbb{E}[X]}{M}.$$

Taking the negative logarithm of both sides gives the assertion of the lemma. \square

The next lemma lies at the heart of the matter. In very broad terms, it states that the upper tail event $\{X \geq (1 + \delta) \mathbb{E}[X]\}$, viewed as a subset of the cube $\{0, 1\}^N$, may be covered almost completely by a union of subcubes of small codimension, where, crucially, the average value of X on each of these subcubes is at least $(1 + \delta - \varepsilon) \mathbb{E}[X]$. The proof uses a variant of the moment argument of Janson, Oleszkiewicz, and Ruciński [44].

Lemma 3.7. *Let Y be a random variable taking values in $\{0, 1\}^N$ and let $X = X(Y)$ be a nonzero polynomial with nonnegative coefficients and degree at most d . Then for every positive integer ℓ and all positive real numbers ε and δ ,*

$$\mathbb{P}(X \geq (1 + \delta) \mathbb{E}[X] \text{ and } Y_I = 0 \text{ for all } I \in \mathcal{I}) \leq \left(\frac{1 + \delta - \varepsilon}{1 + \delta} \right)^\ell,$$

where $\mathcal{I} = \{I \subseteq \llbracket N \rrbracket : |I| \leq d\ell \text{ and } \mathbb{E}_I[X] \geq (1 + \delta - \varepsilon) \mathbb{E}[X]\}$.

Proof. Given $S \subseteq \llbracket N \rrbracket$, let Z_S be the indicator random variable of the event that $Y_I = 0$ for all $I \in \mathcal{I}$ with $I \subseteq S$. Note that $I' \subseteq I$ implies $Z_{I'} \leq Z_I$ and let $Z = Z_{\llbracket N \rrbracket}$. Since $XZ \geq 0$ and $Z^\ell = Z$, Markov's inequality gives

$$\mathbb{P}(X \geq (1 + \delta) \mathbb{E}[X] \text{ and } Z = 1) = \mathbb{P}(XZ \geq (1 + \delta) \mathbb{E}[X]) \leq \frac{\mathbb{E}[X^\ell Z]}{((1 + \delta) \mathbb{E}[X])^\ell}. \quad (15)$$

Write $X = \sum_I \alpha_I Y_I$, where the sum ranges over all subsets $I \subseteq \llbracket N \rrbracket$, each coefficient α_I is nonnegative, and $\alpha_I = 0$ unless $|I| \leq d$. Then for every $k \in \llbracket \ell \rrbracket$,

$$\begin{aligned} \mathbb{E}[X^k Z] &= \sum_{I_1, \dots, I_k} \alpha_{I_1} \cdots \alpha_{I_k} \mathbb{E}[Y_{I_1} \cdots Y_{I_k} \cdot Z] \\ &\leq \sum_{I_1, \dots, I_k} \alpha_{I_1} \cdots \alpha_{I_k} \mathbb{E}[Y_{I_1} \cdots Y_{I_k} \cdot Z_{I_1 \cup \dots \cup I_k}] \\ &\leq \sum_{I_1, \dots, I_{k-1}} \alpha_{I_1} \cdots \alpha_{I_{k-1}} \mathbb{E}[Y_{I_1} \cdots Y_{I_{k-1}} \cdot Z_{I_1 \cup \dots \cup I_{k-1}}] \cdot \mathbb{E}[X \mid Y_{I_1} \cdots Y_{I_{k-1}} \cdot Z_{I_1 \cup \dots \cup I_{k-1}} = 1], \end{aligned}$$

where we may let the third sum range only over sequences I_1, \dots, I_{k-1} for which the event $\{Y_{I_1} \cdots Y_{I_{k-1}} \cdot Z_{I_1 \cup \dots \cup I_{k-1}} = 1\}$ has a positive probability of occurring. Note that for any such sequence, $Y_{I_1} \cdots Y_{I_{k-1}} \cdot Z_{I_1 \cup \dots \cup I_{k-1}} = Y_{I_1} \cdots Y_{I_{k-1}}$ and $I_1 \cup \dots \cup I_{k-1} \notin \mathcal{I}$. Since $|I_1 \cup \dots \cup I_{k-1}| \leq d(k-1) \leq d\ell$, we have

$$\mathbb{E}[X \mid Y_{I_1} \cdots Y_{I_{k-1}} = 1] = \mathbb{E}_{I_1 \cup \dots \cup I_k}[X] < (1 + \delta - \varepsilon) \mathbb{E}[X],$$

as otherwise $I_1 \cup \dots \cup I_{k-1}$ would belong to \mathcal{I} . It follows that

$$\begin{aligned} \sum_{I_1, \dots, I_k} \alpha_{I_1} \cdots \alpha_{I_k} \mathbb{E}[Y_{I_1} \cdots Y_{I_k} \cdot Z_{I_1 \cup \dots \cup I_k}] \\ < (1 + \delta - \varepsilon) \mathbb{E}[X] \cdot \sum_{I_1, \dots, I_{k-1}} \alpha_{I_1} \cdots \alpha_{I_{k-1}} \mathbb{E}[Y_{I_1} \cdots Y_{I_{k-1}} \cdot Z_{I_1 \cup \dots \cup I_{k-1}}]. \end{aligned}$$

By induction, we see that $\mathbb{E}[X^\ell Z] < ((1 + \delta - \varepsilon) \mathbb{E}[X])^\ell$. Substituting this inequality into (15) completes the proof. \square

The following easy lemma will be used to relate the family \mathcal{I} from the statement of Lemma 3.7 to the family \mathcal{I}^* of cores.

Lemma 3.8. *Let Y be a random variable taking values in $\{0, 1\}^N$ and let $X = X(Y)$ be a real-valued function of Y . Then for every $I \subseteq \llbracket N \rrbracket$ and every nonnegative real number s , there exists some $I^* \subseteq I$ such that*

- (i) $\mathbb{E}_{I^*}[X] \geq \mathbb{E}_I[X] - s$ and
- (ii) $\min_{i \in I^*} (\mathbb{E}_{I^*}[X] - \mathbb{E}_{I^* \setminus \{i\}}[X]) \geq s/|I|$.

Proof. Define a sequence $I = I_0 \supseteq I_1 \supseteq \dots \supseteq I_r = I^*$ by repeatedly setting $I_{k+1} = I_k \setminus \{i\}$ for some $i \in I_k$ such that $\mathbb{E}_{I_k}[X] - \mathbb{E}_{I_k \setminus \{i\}}[X] < s/|I|$, as long as such an i exists. By construction, the set I^* satisfies (ii). Finally, since $r \leq |I|$, we have

$$\mathbb{E}_I[X] - \mathbb{E}_{I^*}[X] = \sum_{k=0}^{r-1} (\mathbb{E}_{I_k}[X] - \mathbb{E}_{I_{k+1}}[X]) \leq s,$$

which is (i). \square

Proof of Theorem 3.1. Let $t = (1 + \delta) \mathbb{E}[X]$. We first prove the upper bound in (12). Let $\mathbf{1}$ denote the N -dimensional all-ones vector. Since X is an increasing function of Y , we have $X \leq X(\mathbf{1})$ always. In particular, Lemma 3.6 implies that

$$-\log \mathbb{P}(X \geq t) \leq \Phi_X(\delta + \varepsilon) + \log \left(\frac{X(\mathbf{1})}{\varepsilon \mathbb{E}[X]} \right).$$

As X has degree at most d and nonnegative coefficients, we have $\mathbb{E}[X] \geq X(\mathbf{1}) \cdot p^d$ and thus

$$-\log \mathbb{P}(X \geq t) \leq \Phi_X(\delta + \varepsilon) + \log(1/(\varepsilon p^d)) \leq (1 + \varepsilon/8) \cdot \Phi_X(\delta + \varepsilon), \quad (16)$$

where the second inequality holds provided that K is sufficiently large, as we have assumed that $\Phi_X(\delta + \varepsilon) \geq K \log(1/p)$ and $p \leq 1 - \varepsilon$.

For the rest of the proof let $\ell = \lceil \varepsilon(3d)^{-1}K \cdot \Phi_X(\delta + \varepsilon) \rceil$ and define

$$\mathcal{I} = \{I \subseteq \llbracket N \rrbracket : |I| \leq d\ell \text{ and } \mathbb{E}_I[X] \geq (1 + \delta - \varepsilon/2)\mathbb{E}[X]\}.$$

It follows from Lemma 3.7 (invoked with ε replaced by $\varepsilon/2$) that

$$\mathbb{P}(X \geq t \text{ and } Y_I = 0 \text{ for all } I \in \mathcal{I}) \leq \left(1 - \frac{\varepsilon/2}{1 + \delta}\right)^\ell.$$

Since we have already shown that $\mathbb{P}(X \geq t) \geq \exp(- (1 + \varepsilon)\Phi_X(\delta + \varepsilon))$, see (16), we find that letting K be sufficiently large ensures

$$\mathbb{P}(X \geq t \text{ and } Y_I = 0 \text{ for all } I \in \mathcal{I}) \leq (1 - \varepsilon/(2 + 2\delta))^\ell \leq (\varepsilon/2) \cdot \mathbb{P}(X \geq t).$$

Note next that every $I \in \mathcal{I}$ satisfies $|I| \leq d\ell \leq (\varepsilon K/2) \cdot \Phi_X(\delta + \varepsilon)$ and hence, by Lemma 3.8 applied with $s = \varepsilon \mathbb{E}[X]/2$, there is a subset $I^* \subseteq I$ satisfying the conditions (C1), (C2), and (C3). It follows that

$$\mathbb{P}(X \geq t \text{ and } Y_{I^*} = 0 \text{ for all } I^* \in \mathcal{I}^*) \leq (\varepsilon/2) \cdot \mathbb{P}(X \geq t). \quad (17)$$

Let $\mathcal{I}_m^* := \{I^* \in \mathcal{I}^* : |I^*| = m\}$ and recall that we assume $|\mathcal{I}_m^*| \leq (1/p)^{\varepsilon m/2}$ for all $m \in \mathbb{N}$.

We now prove the upper bound in (12). It follows from (17) that

$$\mathbb{P}(X \geq t) \leq (1 - \varepsilon/2)^{-1} \cdot \mathbb{P}(Y_{I^*} = 1 \text{ for some } I^* \in \mathcal{I}^*).$$

Moreover, the definitions of \mathcal{I}^* and $\Phi_X(\delta - \varepsilon)$ imply that every core $I^* \in \mathcal{I}^*$ satisfies $|I^*| \log(1/p) = -\log \mathbb{P}(Y_{I^*} = 1) \geq \Phi_X(\delta - \varepsilon)$, see (C1). Hence, taking the union bound over all cores and using $|\mathcal{I}_m^*| \leq (1/p)^{\varepsilon m/2}$, we find that

$$\begin{aligned} \mathbb{P}(X \geq t) &\leq (1 - \varepsilon/2)^{-1} \sum_{I^* \in \mathcal{I}^*} p^{|I^*|} \leq (1 - \varepsilon/2)^{-1} \sum_m |\mathcal{I}_m^*| \cdot p^m \\ &\leq (1 - \varepsilon/2)^{-1} \sum_{m=\frac{\Phi_X(\delta - \varepsilon)}{\log(1/p)}}^{\infty} p^{(1 - \varepsilon/2)m} = \frac{e^{-(1 - \varepsilon/2)\Phi_X(\delta - \varepsilon)}}{(1 - \varepsilon/2)(1 - p^{1 - \varepsilon/2})}. \end{aligned}$$

Taking logarithms and using $p \leq 1 - \varepsilon$ and $\Phi_X(\delta - \varepsilon) \geq K \log(1/p)$, we see that a large enough choice of K ensures that $-\log \mathbb{P}(X \geq t) \geq (1 - \varepsilon)\Phi_X(\delta - \varepsilon)$, as required.

Finally, let us prove (13). Using (17), we obtain

$$\mathbb{P}(X \geq t \text{ and } Y_{I^*} = 0 \text{ for all } I^* \in \mathcal{J}^*) \leq (\varepsilon/2) \cdot \mathbb{P}(X \geq t) + \mathbb{P}(Y_{I^*} = 1 \text{ for some } I^* \in \mathcal{I}^* \setminus \mathcal{J}^*).$$

Noting that every $I^* \in \mathcal{I}^* \setminus \mathcal{J}^*$ satisfies $|I^*| \log(1/p) = -\mathbb{P}(Y_{I^*} = 1) > (1 + \varepsilon)\Phi_X(\delta + \varepsilon)$, we may employ a union bound again to show that

$$\mathbb{P}(Y_{I^*} = 1 \text{ for some } I^* \in \mathcal{I}^* \setminus \mathcal{J}^*) \leq \sum_{I^* \in \mathcal{I}^* \setminus \mathcal{J}^*} p^{|I^*|} \leq \frac{e^{-(1 - \varepsilon/2)(1 + \varepsilon)\Phi_X(\delta + \varepsilon)}}{1 - p^{1 - \varepsilon/2}}.$$

In order to complete the proof, it now suffices to show that

$$\frac{e^{-(1 + \varepsilon)(1 - \varepsilon/2)\Phi_X(\delta + \varepsilon)}}{1 - p^{1 - \varepsilon/2}} \leq (\varepsilon/2) \cdot \mathbb{P}(X \geq t). \quad (18)$$

To see that this inequality holds, note first that $(1 + \varepsilon)(1 - \varepsilon/2) > 1 + \varepsilon/4$ as $\varepsilon < 1/2$ and therefore, by (16),

$$e^{-(1 + \varepsilon)(1 - \varepsilon/2)\Phi_X(\delta + \varepsilon)} \leq \mathbb{P}(X \geq t) \cdot e^{-(\varepsilon/8)\Phi_X(\delta + \varepsilon)}.$$

As $p \leq 1 - \varepsilon$ and $\Phi_X(\delta + \varepsilon) \geq K \log(1/p)$, we can choose K so large that $e^{-(\varepsilon/8)\Phi_X(\delta + \varepsilon)}/(1 - p^{1 - \varepsilon/2}) \leq \varepsilon/2$, proving (18). \square

4. ARITHMETIC PROGRESSIONS IN RANDOM SETS OF INTEGERS

Fix an integer $k \geq 3$ and let $X = X_{N,p}^{k\text{-AP}}$ be the number of k -term arithmetic progressions (k -APs) in the random set $\llbracket N \rrbracket_p$. The goal of this section is to study the upper tail of X in the regime where Theorem 3.1 is applicable. In particular, we will prove Theorem 1.3, which we restate here for convenience.

Theorem 1.3. *Let $k \geq 3$ be an integer and let $X = X_{N,p}^{k\text{-AP}}$ denote the number of k -term arithmetic progressions in $\llbracket N \rrbracket_p$. Then, for every fixed positive constant δ and all $p = p(N)$ satisfying $N^{-1} \log N \ll p^{k/2} \ll 1$,*

$$\lim_{N \rightarrow \infty} \frac{-\log \mathbb{P}(X \geq (1 + \delta) \mathbb{E}[X])}{N p^{k/2} \log(1/p)} = \sqrt{\delta}.$$

To prove the theorem, we will use Theorem 3.1 to relate $-\log \mathbb{P}(X \geq (1 + \delta) \mathbb{E}[X])$ to the solution of the optimisation problem

$$\Phi_X(\delta) = \min \{ |I| \log(1/p) : I \subseteq \llbracket N \rrbracket \text{ and } \mathbb{E}_I[X] \geq (1 + \delta) \mathbb{E}[X] \}.$$

More precisely, we shall prove the following statement, which is the main result of this section.

Proposition 4.1. *For every integer $k \geq 3$ and all positive real numbers ε and δ , there exists a positive constant C such that the following holds. Suppose that $N \in \mathbb{N}$ and $p \in (0, 1)$ satisfy $C N^{-1} \log N \leq p^{k/2} \leq 1/C$. Then $X = X_{N,p}^{k\text{-AP}}$ satisfies*

$$(1 - \varepsilon) \Phi_X(\delta - \varepsilon) \leq -\log \mathbb{P}(X \geq (1 + \delta) \mathbb{E}[X]) \leq (1 + \varepsilon) \Phi_X(\delta + \varepsilon).$$

The variational problem $\Phi_X(\delta)$ is a discretisation of the variational problem considered by Bhattacharya, Ganguly, Shao, and Zhao [10]. In their setup, one minimizes over the set of all product measures on $\{0, 1\}^{\llbracket N \rrbracket}$, whereas we only consider ‘planting’ constructions; in other words, we restrict our attention to products of $\text{Ber}(p)$ and $\text{Ber}(1)$ measures. The result below can be easily deduced from [10, Theorem 2.2], but we will reprove it in Section 4.2, for completeness.

Proposition 4.2. *For every integer $k \geq 3$ and all positive real numbers ε and δ , there exists a positive constant C such that the following holds. Suppose that $N \in \mathbb{N}$ and $p \in (0, 1)$ satisfy $C N^{-1} \leq p^{k/2} \leq 1/C$. Then $X = X_{N,p}^{k\text{-AP}}$ satisfies*

$$1 - \varepsilon \leq \frac{\Phi_X(\delta)}{\sqrt{\delta} \cdot N p^{k/2} \log(1/p)} \leq 1 + \varepsilon.$$

Clearly, Propositions 4.1 and 4.2 imply Theorem 1.3.

4.1. Proof outline. The proof of Proposition 4.2 will be relatively straightforward: On the one hand, since every interval (or, more generally, every arithmetic progression) of length $\sqrt{\delta} N p^{k/2}$ contains approximately $\delta \mathbb{E}[X]$ arithmetic progressions of length k , we have $\mathbb{E}_I[X] \geq (1 + \delta - o(1)) \mathbb{E}[X]$ for each such interval I . Consequently, $\Phi_X(\delta) \leq (1 + o(1)) \sqrt{\delta} N p^{k/2} \log(1/p)$. On the other hand, a simple calculation shows that, for every set $I \subseteq \llbracket N \rrbracket$ with $O(N p^{k/2})$ elements, $\mathbb{E}_I[X] - \mathbb{E}[X]$ is asymptotically equal to the number of k -APs in I . Therefore, $\Phi_X(\delta) / \log(1/p)$ is bounded from below (asymptotically) by the minimal size of a set of integers that contains at least $\delta \mathbb{E}[X]$ k -APs. We will show that this minimum is achieved by an interval, see Theorem 4.3 below; thus, we conclude that $\Phi_X(\delta) \geq (1 - o(1)) \sqrt{\delta} N p^{k/2} \log(1/p)$.

In order to derive Proposition 4.1 from Theorem 3.1, we will need to show that the family \mathcal{I}^* of cores from the statement of the theorem is entropically stable. In order to bound the number of cores of a given size, we will first observe that, for every $I \in \mathcal{I}^*$ and each $i \in I$, the difference $\mathbb{E}_I[X] - \mathbb{E}_{I \setminus \{i\}}[X]$ is asymptotically equal to the number of k -APs in I that contain the element i . In particular, condition (C3) and $\Phi_X(\delta) = O(N p^{k/2} \log(1/p))$ can be combined to conclude that

each element of every cores lies in $\Omega(Np^{k/2}/\log(1/p))$ arithmetic progressions of length k that are fully contained in the core; see Claim 4.5 below.

The heart of the proof of the proposition is a counting argument showing that very few sets have this combinatorial property. Let us first sketch a simplified version of this argument that would be sufficient to prove the proposition under the slightly stronger assumption that $Np^{k/2} \gg (\log N)^{k+1}$. Suppose that I is a core of cardinality m and let I' be a random subset of I with $m/\log(1/p)$ elements. For every $i \in I$, we expect that there will be $\Omega(Np^{k/2}/\log(1/p))$ arithmetic progressions of length k in I that contain i , and that a $(\log(1/p))^{1-k}$ -proportion of these will be contained in $I' \cup \{i\}$. A standard application of Janson's inequality yields that the above description holds with probability very close to one, simultaneously for all $i \in I$. In particular, I contains some subset I' with $m/\log(1/p)$ elements and the property that, for every $i \in I \setminus I'$, there are at least $\Omega(Np^{k/2}/(\log(1/p))^k)$ arithmetic progressions of length k that comprise i and some $k - 1$ elements of I' .

We may now enumerate all possible cores I in two steps: First, there are at most $\binom{N}{m/\log(1/p)} \leq \exp(O(m))$ choices of I' . Second, since I' intersects at most $O(|I'|^2)$ arithmetic progressions of length k at $k - 1$ elements and each $i \in I$ is contained in at least $\Omega(Np^{k/2}/(\log(1/p))^k)$ such progressions, the elements of $I \setminus I'$ must all come from a set of size

$$O\left(\frac{|I'|^2(\log(1/p))^k}{Np^{k/2}}\right) \leq O(m(\log(1/p))^{k-1})$$

that depends solely on I' . Thus, the number of choices for $I \setminus I'$, the remaining elements of the core, is at most $\binom{O(m(\log(1/p))^{k-1})}{m} = \exp(O(m \log \log(1/p)))$.

In the proof of Proposition 4.1 below, we give a more refined version of the above argument that allows us to recover the optimal power of the logarithm in the lower bound assumption on p . Instead of constructing cores in two steps, we build them element-by-element. This enables a finer control of the number of choices for each next element, given all the elements chosen so far. Roughly speaking, as we add more elements to I , the set I' from the previous paragraph is gradually increasing its size.

4.2. Estimating Φ_X . As mentioned above, we use the following extremal result about the largest number of k -APs in a set of integers of a given size, proved in the case $k = 3$ by Green–Sisask [37] and later extended in [10] to arbitrary $k \geq 3$; the corresponding statement in the case where $k \in \{1, 2\}$ is trivial. For a set $I \subseteq \mathbb{Z}$, we denote by $A_k(I)$ the number of k -APs in I . Recall that we only count k -APs with positive common difference.

Theorem 4.3 ([10, 37]). *For every positive integer k and $I \subseteq \mathbb{Z}$, we have $A_k(I) \leq A_k(\llbracket I \rrbracket)$.*

We reproduce the proof here for the sake of completeness.

Proof. We prove the statement by induction on k . The cases $k = 1$ and $k = 2$ are trivial as $A_k(I) = A_k(\llbracket I \rrbracket)$ for every set I , so we may assume that $k \geq 3$. Suppose that $|I| = n$ and let a_1, \dots, a_n be the elements of I listed in increasing order. We partition the set of k -APs in I into two parts depending on the location of the $(k - 1)$ st element. More precisely, we let $m = \lceil (k - 2)n/(k - 1) \rceil$, let

$$\mathcal{A}_1 = \{(i_1, \dots, i_k) \in \llbracket n \rrbracket^k : (a_{i_1}, \dots, a_{i_k}) \text{ is a } k\text{-AP and } i_{k-1} \leq m\},$$

and let \mathcal{A}_2 comprise the remaining k -APs (that is, ones with $i_{k-1} > m$). The removal of the k th term from a progression in \mathcal{A}_1 maps it to a $(k - 1)$ -AP contained the set $\{a_1, \dots, a_m\}$ and therefore $|\mathcal{A}_1| \leq A_{k-1}(\{a_1, \dots, a_m\}) \leq A_{k-1}(\llbracket m \rrbracket)$, by the induction hypothesis. On the other hand, we observe that for every $i > m$, there are at most $n - i$ arithmetic progression of length k

such that $i_{k-1} = i$ and thus

$$|\mathcal{A}_2| \leq \sum_{i=m+1}^n n - i.$$

In order to complete the proof, it is sufficient to verify that our choice of m ensures that

$$A_{k-1}(\llbracket m \rrbracket) + \sum_{i=m+1}^n n - i = A_k(\llbracket n \rrbracket).$$

Indeed, m satisfies the following two inequalities:

$$m + \left\lfloor \frac{m-1}{k-2} \right\rfloor \leq n \quad \text{and} \quad m+1 - (k-2)(n-m-1) \geq 1.$$

The first inequality implies that extending any arithmetic progression (i_1, \dots, i_{k-1}) contained in $\llbracket m \rrbracket$ by adjoining to it the element $i_k = 2i_{k-1} - i_{k-2}$ yields a k -AP contained in $\llbracket n \rrbracket$, whereas the second inequality implies that $\sum_{i=m+1}^n n - i$ is precisely the number of k -APs in $\llbracket n \rrbracket$ whose $(k-1)$ st term exceeds m . \square

For future reference, let us note that $A_k(\llbracket i \rrbracket) - A_k(\llbracket i-1 \rrbracket) = \lfloor \frac{i-1}{k-1} \rfloor$ for all positive integers i and $k \geq 2$ and, consequently,

$$A_k(\llbracket m \rrbracket) = \sum_{i=1}^m \left\lfloor \frac{i-1}{k-1} \right\rfloor = \frac{m^2}{2(k-1)} - \frac{(k-1)m}{2} \pm k^2. \quad (19)$$

Using Theorem 4.3, it is not difficult to compute the asymptotic value of $\Phi_X(\delta)$ and complete the proof of Theorem 1.3.

Proof of Proposition 4.2. Without loss of generality, we may assume that $\varepsilon \leq 1$. Given a subset $I \subseteq \llbracket N \rrbracket$, let $a_j(I)$ denote the number of k -APs in $\llbracket N \rrbracket$ that intersect I in exactly j elements. Note that

$$\mathbb{E}_I[X] = \sum_{j=0}^k a_j(I) p^{k-j} \quad \text{and that} \quad \mathbb{E}[X] = A_k(\llbracket N \rrbracket) p^k = \sum_{j=0}^k a_j(I) p^k. \quad (20)$$

It follows that $\mathbb{E}_I[X] - \mathbb{E}[X] \geq (1-p^k)a_k(I) = (1-p^k)A_k(I)$ for every $I \subseteq \llbracket N \rrbracket$. In particular, whenever $(1-p^k)A_k(\llbracket m \rrbracket) \geq \delta p^k A_k(\llbracket N \rrbracket)$, then $\mathbb{E}_{\llbracket m \rrbracket}[X] \geq (1+\delta)\mathbb{E}[X]$. Therefore,

$$\Phi_X(\delta) \leq \min \left\{ m \log(1/p) : A_k(\llbracket m \rrbracket) \geq \frac{\delta p^k A_k(\llbracket N \rrbracket)}{1-p^k} \right\} \leq (1+\varepsilon) \cdot \sqrt{\delta} \cdot N p^{k/2} \log(1/p),$$

where the last inequality follows from (19) and our assumption $N^2 p^k \geq C^2$ for a sufficiently large constant C . It remains to prove the matching lower bound.

Suppose that I is a smallest subset of $\llbracket N \rrbracket$ with $\mathbb{E}_I[X] \geq (1+\delta)\mathbb{E}[X]$. Then (20) implies

$$\delta A_k(\llbracket N \rrbracket) p^k = \delta \mathbb{E}[X] \leq \mathbb{E}_I[X] - \mathbb{E}[X] \leq \sum_{j=1}^k a_j(I) p^{k-j}. \quad (21)$$

Since every pair of distinct numbers in $\llbracket N \rrbracket$ is contained in at most $\binom{k}{2}$ arithmetic progressions of length k , it follows that $a_1(I) \leq |I| \cdot N k^2$ and $\sum_{j=2}^{k-1} a_j(I) \leq |I|^2 \cdot k^2$. Since we already know that $|I| \leq 2\sqrt{\delta} \cdot N p^{k/2}$, inequality (21) gives

$$\delta A_k(\llbracket N \rrbracket) p^k \leq 2\sqrt{\delta} k^2 N^2 p^{3k/2-1} + 4\delta k^2 N^2 p^{k+1} + A_k(I).$$

We now invoke Theorem 4.3 and (19) to obtain

$$(1-\varepsilon) \cdot \frac{\delta N^2 p^k}{2(k-1)} \leq \delta A_k(\llbracket N \rrbracket) p^k - 2\sqrt{\delta} k^2 N^2 p^{3k/2-1} - 4\delta k^2 N^2 p^{k+1} \leq A_k(\llbracket I \rrbracket) \leq \frac{|I|^2}{2(k-1)},$$

where we use the assumptions $p \leq C^{-2/k}$ and $N^2 p^k \geq C^2$ for a large enough C . Thus $\Phi_X(\delta) = |I| \log(1/p) \geq (1 - \varepsilon) \cdot \sqrt{\delta} \cdot N p^{k/2} \log(1/p)$, as required. \square

4.3. Janson's inequality. It remains to prove Proposition 4.1. The proof uses the following version of Janson's inequality for hypergeometric random variables. It follows from the (original version of) Janson's inequality for binomial distributions [43, Theorem 1] and the fact that the median of a binomial random variable whose mean is an integer is equal to its mean. Our argument is an adaptation of [5, Lemma 3.1].

Lemma 4.4. *Suppose that $\{B_\alpha\}_{\alpha \in A}$ is a family of subsets of a t -element set Ω . Let $s \in \{0, \dots, t\}$ and let*

$$\mu = \sum_{\alpha \in A} \binom{s}{t}^{|B_\alpha|} \quad \text{and} \quad \Delta = \sum_{\alpha \sim \beta} \binom{s}{t}^{|B_\alpha \cup B_\beta|},$$

where the second sum is over all ordered pairs $(\alpha, \beta) \in A^2$ such that $\alpha \neq \beta$ and $B_\alpha \cap B_\beta \neq \emptyset$. Let S be the uniformly chosen random s -element subset of Ω and let Z denote the number of $\alpha \in A$ such that $B_\alpha \subseteq S$. Then for every $\varepsilon \in (0, 1]$,

$$\mathbb{P}(Z \leq (1 - \varepsilon)\mu) \leq 2 \exp\left(-\frac{\varepsilon^2}{2} \cdot \frac{\mu^2}{\mu + \Delta}\right).$$

Proof. For every $k \in \{0, \dots, t\}$, let S_k be the uniformly chosen random k -element subset of Ω and let Z_k denote the number of $\alpha \in A$ such that $B_\alpha \subseteq S_k$, so that $Z = Z_s$, and note that there exists a natural coupling under which $Z_k \leq Z_{k+1}$ for every k . Let S' be the (s/t) -random subset of Ω , that is the random subset of Ω formed by keeping each of its elements with probability s/t , independently of others, and let Z' denote the number of $\alpha \in A$ such that $B_\alpha \subseteq S'$. Since $\mathbb{E}[Z' \mid |S'|] = Z_{|S'|}$, the stochastic ordering of the Z_k s implies that, for any number M , the function $k \mapsto \mathbb{P}(Z' \leq M \mid |S'| = k)$ is decreasing. Hence,

$$\begin{aligned} \mathbb{P}(Z' \leq (1 - \varepsilon)\mu) &= \sum_{k=0}^t \mathbb{P}(Z' \leq (1 - \varepsilon)\mu \mid |S'| = k) \cdot \mathbb{P}(|S'| = k) \\ &\geq \mathbb{P}(Z' \leq (1 - \varepsilon)\mu \mid |S'| = s) \cdot \mathbb{P}(|S'| \leq s) \\ &= \mathbb{P}(Z \leq (1 - \varepsilon)\mu) \cdot \mathbb{P}(|S'| \leq s) \geq \mathbb{P}(Z \leq (1 - \varepsilon)\mu)/2, \end{aligned}$$

where the last inequality follows from the well-known fact that if np is an integer, then it is the median of the binomial distribution with parameters n and p . We can now invoke the classical version of Janson's inequality and conclude that

$$\mathbb{P}(Z \leq (1 - \varepsilon)\mu) \leq 2\mathbb{P}(Z' \leq (1 - \varepsilon)\mu) \leq 2 \exp\left(-\frac{\varepsilon^2}{2} \cdot \frac{\mu^2}{\mu + \Delta}\right). \quad \square$$

4.4. Proof of Proposition 4.1. We may assume without loss of generality that ε is sufficiently small, say $\varepsilon < \min\{1/2, \delta/2\}$. Note also that the case $N \leq 2$ is trivial; indeed, in that case X is identically zero and thus $\log \mathbb{P}(X \geq (1 + \delta)\mathbb{E}[X]) = 0 = \Phi_X(\delta)$ for every $\delta \in \mathbb{R}$. We may therefore assume that $N \geq 3$, which, in turn, implies that $N^2 p^k \geq C^2$.

Denote by Y_i the indicator random variable of the event that $i \in \llbracket N \rrbracket_p$. Then $Y = (Y_1, \dots, Y_N)$ is a vector of independent $\text{Ber}(p)$ random variables and X is a nonzero polynomial with nonnegative coefficients and degree at most k in the coordinates of Y . Let $K = K(k, \varepsilon, \delta)$ be the constant given by Theorem 3.1. The proposition follows once we verify that X satisfies the various assumptions of the theorem.

First, our assumption on p implies that $p \leq 1 - \varepsilon$ whenever C is large enough. Second, it follows from Proposition 4.2 and the inequality $N^2 p^k \geq C^2$ that, whenever C is large enough, $\Phi_X(\delta - \varepsilon) \geq \Phi_X(\delta/2) \geq K \log(1/p)$. Recall that a subset $I \subseteq \llbracket N \rrbracket$ is called a core if

- (C1) $\mathbb{E}_I[X] \geq (1 + \delta - \varepsilon)\mathbb{E}[X]$,
- (C2) $|I| \leq K \cdot \Phi_X(\delta + \varepsilon)$, and

$$(C3) \quad \min_{i \in I} (\mathbb{E}_I[X] - \mathbb{E}_{I \setminus \{i\}}[X]) \geq \mathbb{E}[X] / (K \cdot \Phi_X(\delta + \varepsilon)).$$

The final assumption of Theorem 3.1 is that, for every integer m , there are at most $(1/p)^{\varepsilon m/2}$ cores of size m .

In order to count the cores, we must first unravel the meaning of (C1), (C2), and (C3), and show that each core enjoys a simple combinatorial property. Proposition 4.2 supplies a constant $K' = K'(K, k, \varepsilon, \delta)$ such that, whenever C is sufficiently large,

$$4kK \cdot \Phi_X(\delta + \varepsilon) \leq 4kK \cdot (1 + \varepsilon)\sqrt{\delta + \varepsilon} \cdot Np^{k/2} \log(1/p) \leq K' \cdot Np^{k/2} \log(1/p). \quad (22)$$

Given a set $I \subseteq \llbracket N \rrbracket$ and an $i \in \llbracket N \rrbracket$, we write $A_k(I; i)$ for the number of k -term arithmetic progressions in $I \cup \{i\}$ that contain the element i . The proof of the following claim is similar to the argument used to prove Proposition 4.2.

Claim 4.5. *For every core I of size m and all $i \in I$,*

$$A_k(I; i) \geq \frac{Np^{k/2}}{K' \log(1/p)}.$$

Proof. Given an $i \in I$, let $a_j(I; i)$ denote the number of k -APs in $\llbracket N \rrbracket$ that intersect I in exactly j elements, one of which is i . With this notation, $A_k(I; i) = a_k(I; i)$ and we may write $\mathbb{E}_I[X] - \mathbb{E}_{I \setminus \{i\}}[X] = \sum_{j=1}^k a_j(I; i) \cdot p^{k-j}(1-p)$. Since every pair of distinct numbers in $\llbracket N \rrbracket$ is contained in at most $\binom{k}{2}$ arithmetic progressions of length k , we have $a_1(I; i) \leq Nk^2$ and $\sum_{j=2}^{k-1} a_j(I; i) \leq mk^2$. In particular, as $m \leq K \cdot \Phi_X(\delta + \varepsilon)$ by (C2), we get

$$\mathbb{E}_I[X] - \mathbb{E}_{I \setminus \{i\}}[X] \leq k^2 Np^{k-1} + k^2 K \cdot \Phi_X(\delta + \varepsilon) \cdot p + A_k(I; i).$$

On the other hand, it follows from (C3) that

$$\mathbb{E}_I[X] - \mathbb{E}_{I \setminus \{i\}}[X] \geq \frac{\mathbb{E}[X]}{K \cdot \Phi_X(\delta + \varepsilon)}.$$

By (19), we have $\mathbb{E}[X] = A_k(\llbracket N \rrbracket)p^k \geq N^2 p^k / (2k)$, since $N \geq C$ and C is large. Combining the upper and lower bounds on $\mathbb{E}_I[X] - \mathbb{E}_{I \setminus \{i\}}[X]$ and using (22), we obtain

$$A_k(I; i) \geq \frac{2Np^{k/2}}{K' \log(1/p)} - k^2 Np^{k-1} - kK' Np^{k/2+1} \log(1/p).$$

Since $k \geq 3$ and $p \leq C^{-2/k}$ for a large enough C , we deduce the assertion of the claim. \square

For the remainder of the proof, fix some integer m satisfying $1 \leq m \leq K \cdot \Phi_X(\delta + \varepsilon)$ and let K'' be a sufficiently large positive constant depending on K' and k (but not on C). For a subset $I' \subseteq \llbracket N \rrbracket$ and an integer $i \in \llbracket N \rrbracket \setminus I'$, we shall say that i is *rich* with respect to I' if

$$A_k(I'; i) \geq \frac{Np^{k/2}}{K'' \log(1/p)} \cdot \left(\frac{|I'|}{m}\right)^{k-1}. \quad (23)$$

Moreover, given a sequence (i_1, \dots, i_m) of m distinct elements of $\llbracket N \rrbracket$, we shall say that an index $m' \in \llbracket m \rrbracket$ is *rich* if $i_{m'}$ is rich with respect to the set $\{i_1, \dots, i_{m'-1}\}$.

We first observe that for every $I' \subseteq \llbracket N \rrbracket$, there are relatively few integers $i \in \llbracket N \rrbracket \setminus I'$ that are rich with respect to I' . Indeed, there are at most $k|I'|^2$ arithmetic progressions P of length k in $\llbracket N \rrbracket$ for which $|I' \cap P| = k-1$, because any such progression is determined by its minimal and maximal element in I' and the position in the progression of the element in $P \setminus I'$. Then

$$|\{i \in \llbracket N \rrbracket \setminus I' : i \text{ is rich w.r.t. } I'\}| \cdot \frac{Np^{k/2}}{K'' \log(1/p)} \cdot \left(\frac{|I'|}{m}\right)^{k-1} \leq \sum_{i \in \llbracket N \rrbracket \setminus I'} A_k(I'; i) \leq k|I'|^2.$$

Consequently, as $m \leq K' \cdot Np^{k/2} \log(1/p)$ by (22),

$$|\{i \in \llbracket N \rrbracket \setminus I' : i \text{ is rich w.r.t. } I'\}| \leq kK'K'' \cdot m(\log(1/p))^2 \cdot \left(\frac{m}{|I'|}\right)^{k-3}. \quad (24)$$

The key property that allows us to control the number of cores I of size m is that, in a large proportion of orderings of the members of I , almost all indices are rich. This property implies that, if one builds an (ordered) core element by element, then, very often, one must choose the next element from the small set of integers that are rich with respect to the previously chosen ones. From this, it will be easy to obtain an upper bound on the number of cores of a given size.

Claim 4.6. *Suppose that I is a core of size m . Then there are at least $m!/2$ orderings (i_1, \dots, i_m) of the elements of I such that all but at most*

$$\left(\frac{K'' \log(1/p)}{Np^{k/2}} \right)^{\frac{1}{k-1}} \cdot m \quad (25)$$

indices $m' \in \llbracket m \rrbracket$ are rich.

Proof. Let (i_1, \dots, i_m) be a uniformly chosen random ordering of the elements of I . Fix integers $m' \in \llbracket m \rrbracket$ and $i \in I$ and condition on the event $\{i_{m'} = i\}$. Under this conditioning, the set $\{i_1, \dots, i_{m'-1}\}$ is a uniformly random $(m'-1)$ -element subset of $I \setminus \{i\}$. Therefore, we may use Janson's inequality for the hypergeometric distribution (Lemma 4.4) to get an upper bound for the probability that the given m' is not rich. It follows from the definition that $m' = 1$ is trivially rich, so assume $m' \geq 2$. Let \mathcal{B}_i be the collection of all $(k-1)$ -element subsets of $I \setminus \{i\}$ that form a k -AP with i . Define

$$\mu_{m'}(i) = \sum_{B \in \mathcal{B}_i} \left(\frac{m'-1}{m-1} \right)^{|B|} = A_k(I; i) \cdot \left(\frac{m'-1}{m-1} \right)^{k-1}$$

and, writing $B \sim B'$ to mean that $B \neq B'$ and $B \cap B' \neq \emptyset$,

$$\Delta_{m'}(i) = \sum_{\substack{B, B' \in \mathcal{B}_i \\ B \sim B'}} \left(\frac{m'-1}{m-1} \right)^{|B \cup B'|}.$$

Since for a given $B \in \mathcal{B}_i$, there are fewer than k^3 sets $B' \in \mathcal{B}_i$ such that $B \cap B' \neq \emptyset$, we have $\Delta_{m'}(i) \leq \mu_{m'}(i) \cdot k^3$. It follows from Claim 4.5 that

$$\mu_{m'}(i) \geq \frac{Np^{k/2}}{K' \log(1/p)} \cdot \left(\frac{m'-1}{m} \right)^{k-1},$$

which, provided that K'' is sufficiently large, is at least twice as large as the right-hand side of (23) with $|I'| = m' - 1$. Hence, by Lemma 4.4 with $\varepsilon = 1/2$,

$$\begin{aligned} \mathbb{P}(m' \text{ is not rich} \mid i_{m'} = i) &\leq 2 \exp \left(-\frac{\mu_{m'}(i)^2}{8(\mu_{m'}(i) + \Delta_{m'}(i))} \right) \\ &\leq 2 \exp \left(-\frac{\mu_{m'}(i)}{9k^3} \right) \\ &\leq 2 \exp \left(-\frac{Np^{k/2}}{9k^3 K' \log(1/p)} \cdot \left(\frac{m'-1}{m} \right)^{k-1} \right). \end{aligned}$$

Since this upper bound is independent of i , one may replace the conditional probability above with the unconditional one. Letting $\mathcal{X} \subseteq \llbracket m \rrbracket$ denote the (random) set of non-rich indices, we then find that

$$\mathbb{E} [|\mathcal{X}|] \leq 2 \sum_{m'=2}^m \exp \left(-\frac{Np^{k/2}}{9k^3 K' \log(1/p)} \cdot \left(\frac{m'-1}{m} \right)^{k-1} \right).$$

Since for every $\alpha > 0$, we have

$$\sum_{m'=2}^m \exp \left(-\left(\alpha \cdot \frac{m'-1}{m} \right)^{k-1} \right) \leq \int_0^\infty e^{-(\alpha x/m)^{k-1}} dx = \frac{m}{\alpha} \int_0^\infty e^{-y^{k-1}} dy \leq \frac{2m}{\alpha}$$

we obtain

$$\mathbb{E}[|\mathcal{X}|] \leq 4 \cdot \left(\frac{9k^3 K' \log(1/p)}{Np^{k/2}} \right)^{\frac{1}{k-1}} \cdot m.$$

The assertion of the claim now follows from Markov's inequality, provided that K'' is sufficiently large. \square

Equipped with the above facts, we can now prove the desired upper bound on the number of cores of size m . For a set $\mathcal{X} \subseteq \llbracket m \rrbracket$, let $\mathcal{S}_m(\mathcal{X})$ denote the family of all sequences of m distinct elements of $\llbracket N \rrbracket$ such that every index $m' \notin \mathcal{X}$ is rich. To control the number of sequences in $\mathcal{S}_m(\mathcal{X})$, note that we can pick the first element of the sequence arbitrarily and, for every subsequent index m' , bound the number of possible values for the m' th element of the sequence either by appealing to (24), if $m' \notin \mathcal{X}$, or simply by N , otherwise. Thus,

$$\begin{aligned} \frac{|\mathcal{S}_m(\mathcal{X})|}{m!} &\leq \frac{1}{m!} \cdot N \cdot N^{|\mathcal{X}|} \cdot \prod_{m'=2}^m \left(kK'K'' \cdot m(\log(1/p))^2 \cdot \left(\frac{m}{m'-1} \right)^{k-3} \right) \\ &\leq N \cdot N^{|\mathcal{X}|} \cdot \left(kK'K'' \cdot (\log(1/p))^2 \right)^m \cdot \prod_{m'=1}^m \left(\frac{m}{m'} \right)^{k-2}. \end{aligned}$$

Since $\prod_{m'=1}^m \left(\frac{m}{m'} \right)^{k-2} \leq e^{(k-2)m}$, we find that, whenever C is sufficiently large,

$$\frac{|\mathcal{S}_m(\mathcal{X})|}{m!} \leq e^{(|\mathcal{X}|+1)\log N} \cdot e^{3m \log \log(1/p)}.$$

Finally, denote by \mathcal{I}_m^* the set of all cores of size m . Claim 4.6 implies that

$$|\mathcal{I}_m^*| \leq \frac{2}{m!} \sum_{\mathcal{X}} |\mathcal{S}_m(\mathcal{X})| \leq 2^{m+1} \cdot \max_{\mathcal{X}} \frac{|\mathcal{S}_m(\mathcal{X})|}{m!},$$

where the sum and the maximum range over all $\mathcal{X} \subseteq \llbracket N \rrbracket$ of size at most $\left(\frac{K'' \log(1/p)}{Np^{k/2}} \right)^{\frac{1}{k-1}} m$. Hence,

$$|\mathcal{I}_m^*| \leq 2^{m+1} \cdot \exp \left(\left(\frac{K'' \log(1/p)}{Np^{k/2}} \right)^{\frac{1}{k-1}} \cdot m \cdot \log N + \log N \right) \cdot e^{3m \log \log(1/p)}.$$

Since we have assumed that $Np^{k/2} \geq C \log N$ and $p \leq C^{-2/k}$, then, whenever C is sufficiently large, the above inequality implies that

$$|\mathcal{I}_m^*| \leq (1/p)^{\varepsilon m/4} \cdot \exp \left(\left(\frac{K'' \log(1/p)}{Np^{k/2}} \right)^{\frac{1}{k-1}} \cdot m \cdot \log N \right) \leq (1/p)^{\varepsilon m/2},$$

where the last inequality can be seen, for example, by distinguishing between the cases $p \leq N^{-1/k}$ (in which case $\log N = \Theta(\log(1/p))$) and $p > N^{-1/k}$ (where $Np^{k/2} > N^{1/k}$). This completes the proof of Proposition 4.1.

5. COUNTING SMALL SUBGRAPHS—A GRAPH-THEORETIC INTERLUDE

As mentioned in the introduction, this section will collect some graph-theoretical results which will be required to analyze the localized regime of $X_{n,p}^{K_r}$ and $X_{n,p}^H$ for connected, regular graphs H in Sections 6 and 7, respectively.

The main goal is to bound the maximum number of embeddings of a given graph J into a larger graph G in terms of the number of vertices and edges in G , where we are interested both in bounding the number of such embeddings *globally* (i.e., without additional restrictions on the image) and *locally* (where we require that the image contain a particular edge of G). These estimates will play a crucial role in translating conditions (C1)–(C3) from Theorem 3.1 into structural restrictions on core graphs.

In the first two subsections, we collect results related to bounding the number of embeddings globally; the fractional independence number of a graph (defined below) plays an important role in this part. The next subsection contains bounds on the number of local embeddings, where the image of the embedding is required to contain a particular edge. In the final subsection, we establish several stability results (in the sense of extremal combinatorics) on graphs allowing a nearly maximal number of embeddings of stars and cliques; these results will be of use when establishing Theorem 1.8.

Recall that $\text{Emb}(J, G)$ denotes the set of embeddings of J into G and, for every edge uv of G , $\text{Emb}(J, G; uv)$ denotes the subset of $\text{Emb}(J, G)$ containing all embeddings that map an edge of J to uv .

5.1. Fractional graph theory. A *fractional independent set* in a graph J is an assignment $\alpha: V(J) \rightarrow [0, 1]$ that satisfies $\alpha_u + \alpha_v \leq 1$ for every edge uv of J . The *fractional independence number* of J , denoted by α_J^* , is the largest value of $\sum_{v \in V(J)} \alpha_v$ among all fractional independent sets α in J . The following result is folklore; we include a proof for completeness.

Lemma 5.1. *Every graph J admits a fractional independent set α with $\sum_{v \in V(J)} \alpha_v = \alpha_J^*$ such that $\alpha_v \in \{0, \frac{1}{2}, 1\}$ for every $v \in V(J)$. Moreover, there is a partition $V(J) = V_1 \cup V_2$ with $|V_1|/2 + |V_2| = \alpha_J^*$ such that V_1 can be covered by a collection of vertex-disjoint edges and cycles of J .*

Proof. Let J' be the bipartite double cover of J , that is, the graph with vertex set $V(J) \times \{1, 2\}$ whose edges are all pairs $\{(u, 1), (v, 2)\}$ such that $uv \in E(J)$. Moreover, let $\pi: V(J) \times \{1, 2\} \rightarrow V(J)$ be the projection onto the first coordinate. The König–Egerváry theorem (see, e.g., [13, Theorem 8.32] or [29, Theorem 2.1.1]) implies that J' contains a matching M' and an independent set I' such that $|I'| + |M'| = v_{J'}$. Define $\alpha: V(J) \rightarrow \{0, \frac{1}{2}, 1\}$ by letting $\alpha_v = |\pi^{-1}(v) \cap I'|/2$ for every $v \in V(J)$. Since I' is an independent set in J' , one can see that α is a fractional independent set with $\sum_{v \in V(J)} \alpha_v = |I'|/2$. In particular, we have

$$v_{J'} - |M'| = 2 \sum_{v \in V(J)} \alpha_v \leq 2\alpha_J^*. \quad (26)$$

Since π induces a projection of J' onto J , we can define $M = \pi(M')$ to be the image of the matching M' . Since $M \subseteq J$, we have $\alpha_J^* \leq \alpha_M^*$. Moreover, as M' is a matching in J' , we see that M has maximum degree at most two and thus each nontrivial connected component of M is either a cycle or a path. Let $V_2 \subseteq V(J)$ comprise all isolated vertices of M and one arbitrarily chosen endpoint of each path of even length; let $V_1 = V(J) \setminus V_2$. By construction, each connected component of $M[V_1]$ is either a cycle or a path of odd length. Since the fractional independence number of every cycle and every path of odd length is exactly half its number of vertices, it follows that $\alpha_M^* \leq \alpha_{M[V_1]}^* + |V_2| = |V_1|/2 + |V_2|$. It is clear that V_1 can be covered with vertex-disjoint edges and cycles of M and thus also of J . We now claim that $|V_1| \geq |M'|$. To see this, fix a connected component L of M and observe that $\pi^{-1}(L)$ has at most e_L edges unless L is a single edge, in which case $\pi^{-1}(L)$ has at most two edges. Therefore,

$$e_{\pi^{-1}(L)} \leq \begin{cases} v_L - 1 & \text{if } L \text{ is a path of length at least two,} \\ v_L & \text{otherwise.} \end{cases}$$

Let $\mathcal{C}(M)$ denote the nontrivial connected components of M . We have

$$\begin{aligned} |M'| &= \sum_{L \in \mathcal{C}(M)} e_{\pi^{-1}(L)} \leq \sum_{L \in \mathcal{C}(M)} v_L - \mathbb{1}[L \text{ is a path of length at least two}] \\ &\leq \sum_{L \in \mathcal{C}(M)} v_L - \mathbb{1}[L \text{ is a path of even length}] = |V_1|. \end{aligned}$$

Consequently, (26) shows that

$$|V_1| + 2|V_2| = 2v_J - |V_1| \leq 2v_J - |M'| = v_{J'} - |M'| = 2 \sum_{v \in V(J)} \alpha_v \leq 2\alpha_J^* \leq 2\alpha_M^* \leq |V_1| + 2|V_2|,$$

and so $\sum_{v \in V(J)} \alpha_v = \alpha_J^* = |V_1|/2 + |V_2|$. \square

The following lemma is implicit in [44, Appendix A].

Lemma 5.2. *Suppose that J is a nonempty subgraph of a connected, Δ -regular graph H . Then*

$$e_J \leq \Delta \cdot (v_J - \alpha_J^*) \leq \Delta \cdot \alpha_J^*.$$

If the first inequality is tight, then

(Q1) $J = H$ or

(Q2) J admits a bipartition $V(J) = A \cup B$ such that $\deg_J a = \Delta$ for all $a \in A$.

If both inequalities are tight, then $J = H$.

Remark 5.3. Since every graph J is a subgraph of the complete graph of v_J vertices, Lemma 5.2 implies that

$$e_J \leq (v_J - 1) \cdot (v_J - \alpha_J^*).$$

Moreover, equality holds if and only if J is complete, J is empty, or $J = K_{1, v_J - 1}$.

Proof of Lemma 5.2. By Lemma 5.1, J has a fractional independent set α such that $\alpha_v \in \{0, \frac{1}{2}, 1\}$. Then

$$e_J \leq \sum_{uv \in E(J)} (2 - \alpha_u - \alpha_v) = \sum_{v \in V(J)} (1 - \alpha_v) \deg_J v \leq \Delta \cdot \sum_{v \in V(J)} (1 - \alpha_v) = \Delta \cdot (v_J - \alpha_J^*), \quad (27)$$

which is the first inequality. For the second inequality, note that the function $\alpha: V(J) \rightarrow [0, 1]$ defined by $\alpha_v = 1/2$ is a fractional independent set, so $\alpha_J^* \geq v_J/2$.

Assume now that $e_J = \Delta \cdot (v_J - \alpha_J^*)$. Then both inequalities in (27) are equalities; this implies $\alpha_u + \alpha_v = 1$ for every edge $uv \in E(J)$ and $\deg_J(v) = \Delta$ whenever $\alpha_v \neq 1$. Let A , B , and C denote the sets of vertices that α maps to 0, 1, and 1/2, respectively. Each vertex in $A \cup C$ has degree Δ and each edge of J has either both endpoints in C or one endpoint in each of A and B . In particular, if C is not empty, then it induces a Δ -regular graph and hence $C = V(J)$ and $J = H$, as H is connected and Δ -regular. Otherwise, if C is empty, then $A \cup B$ is a bipartition of J and all vertices of A have degree Δ .

Lastly, suppose that $e_J = \Delta \cdot \alpha_J^*$, which implies $e_J = \Delta \cdot (v_J - \alpha_J^*)$. Let A, B, C be the same partition as above. If C is nonempty, then J is Δ -regular, and we are done. Otherwise,

$$|A| = e_J/\Delta = \alpha_J^* = v_J - e_J/\Delta = v_J - |A| = |B|.$$

Therefore, every vertex of B has degree Δ and $J = H$. \square

5.2. Global embedding bounds. The main result of this section is the following theorem of Janson, Oleszkiewicz, and Ruciński [44]. A closely related bound that does not depend on the number of vertices in G was obtained earlier by Alon [1] (see also [34] for a short proof). The dependence on the number of vertices will be essential in the case where J is a (double) star, see Figure 5.

Theorem 5.4 ([44]). *For every nonempty graph J without isolated vertices and every graph G with n vertices,*

$$|\text{Emb}(J, G)| \leq (2e_G)^{v_J - \alpha_J^*} \cdot \min\{2e_G, n\}^{2\alpha_J^* - v_J}$$

We derive Theorem 5.4 from Lemma 5.1 and the following result due to Alon [1], which establishes the theorem for the case where J is a cycle.

Lemma 5.5. *Let C_ℓ denote the cycle of length ℓ . For every $\ell \geq 3$ and every graph G ,*

$$|\text{Emb}(C_\ell, G)| \leq (2e_G)^{\ell/2}.$$

Remark 5.6. If ℓ is even, this follows immediately from the fact that C_ℓ contains a perfect matching of $\ell/2$ edges. If ℓ is prime, there is also a very short and pretty proof using the monotonicity of L^p norms; see [60] for this proof and more precise estimates. The proof presented below works for all $\ell \geq 3$.

Proof of Lemma 5.5. For each edge $e \in E(G)$, denote by c_e the number of copies of C_ℓ in G that contain the edge e . Since $\sum_{e \in E(G)} c_e = \ell \cdot N(C_\ell, G)$, where $N(C_\ell, G)$ is the number of copies of C_ℓ in G , it follows from the Cauchy–Schwarz inequality that

$$|\text{Emb}(C_\ell, G)|^2 = (2\ell \cdot N(C_\ell, G))^2 \leq 2e_G \cdot \sum_{e \in E(G)} 2c_e^2.$$

Let C_ℓ^* be the graph obtained from gluing two copies of C_ℓ along an edge. In other words, C_ℓ^* is obtained from the cycle of length $2\ell - 2$ by adding to it one longest chord. Observe that if (L_1, L_2) is an ordered pair of copies of C_ℓ in G , both containing e , then there are at exactly two homomorphisms $\varphi: V(C_\ell^*) \rightarrow V(G)$ that map the two vertices of degree three in C_ℓ^* onto the endpoints of e and the two copies of C_ℓ in C_ℓ^* onto L_1 and L_2 , respectively. Letting $\text{Hom}(C_\ell^*, G)$ be the collection of all homomorphisms from C_ℓ^* to G , we may conclude that

$$\sum_{e \in E(G)} 2c_e^2 \leq |\text{Hom}(C_\ell^*, G)| \leq |\text{Hom}((\ell - 1) \cdot K_2, G)| \leq (2e_G)^{\ell-1},$$

where the second inequality holds because C_ℓ^* contains a perfect matching of $\ell - 1$ edges. \square

Proof of Theorem 5.4. By Lemma 5.1, there is a partition of $V(J)$ into V_1 and V_2 such that $|V_1|/2 + |V_2| = \alpha_J^*$ and V_1 can be covered by a collection \mathcal{C} of vertex-disjoint edges and cycles of J . Let J' be the spanning subgraph comprising the edges and cycles of \mathcal{C} and one edge incident to every vertex in V_2 . We claim that

$$|\text{Emb}(J', G)| \leq \prod_{C \in \mathcal{C}} |\text{Emb}(C, G)| \cdot \min\{2e_G, n\}^{|V_2|}.$$

Indeed, every embedding of $J'[V_1]$ decomposes into embeddings of the graphs in \mathcal{C} , and there are at most $\min\{2e_G, n\}$ possible images for every vertex of V_2 . By Lemma 5.5, for every cycle $C \in \mathcal{C}$,

$$|\text{Emb}(C, G)| \leq (2e_G)^{v_C/2};$$

the same inequality holds when C is a single edge. Since every embedding of J into G is also an embedding of J' , we deduce that

$$|\text{Emb}(J, G)| \leq \prod_{C \in \mathcal{C}} (2e_G)^{v_C/2} \cdot \min\{2e_G, n\}^{|V_2|} = (2e_G)^{|V_1|/2} \cdot \min\{2e_G, n\}^{|V_2|}.$$

Since $|V_1|/2 = v_J - \alpha_J^*$ and $|V_2| = 2\alpha_J^* - v_J$, this completes the proof. \square

5.3. Local embedding bounds. We now state three lemmas that bound $|\text{Emb}(J, G; uv)|$ from above.

Lemma 5.7. *Suppose that H is a Δ -regular graph. For every graph G and each $uv \in E(G)$,*

$$|\text{Emb}(H, G; uv)| \leq 4e_H \cdot (2e_G)^{\frac{v_H}{2} - \frac{2\Delta-1}{\Delta}} \cdot (4 \deg_G u \cdot \deg_G v)^{\frac{\Delta-1}{\Delta}}.$$

Lemma 5.8. *Suppose that J is a nonempty, connected graph with maximum degree Δ that admits a bipartition $V(J) = A \cup B$ such that $|A| < |B|$ and $\deg_J a = \Delta$ for every $a \in A$. For every graph G and every $uv \in E(G)$,*

$$|\text{Emb}(J, G; uv)| \leq e_J \cdot (\deg_G u + \deg_G v) \cdot (2e_G)^{|A|-1} \cdot (\min\{e_G, v_G\})^{|B|-|A|-1}.$$

Lemma 5.9. *Suppose that H is a Δ -regular graph. For every graph G and every $G' \subseteq G$,*

$$\sum_{uv \in E(G')} |\text{Emb}(H, G; uv)| \leq e_H \cdot (2e_G)^{v_H/2} \cdot \left(\frac{e_{G'}}{e_G}\right)^{1/\Delta}.$$

Our proofs of Lemmas 5.7 and 5.9 are relatively straightforward adaptations of the elegant entropy argument of Friedgut and Kahn [34] (see also the excellent survey of Galvin [35]). They will be derived from the following somewhat abstract form of the main result of [34]. The proof of Lemma 5.8 is elementary.

Lemma 5.10. *Suppose that H is a Δ -regular graph. Let \mathcal{E} be a family of embeddings of H into a graph G and, for every edge ab of H , let*

$$\mathcal{E}_{ab} = \{\varphi(ab) : \varphi \in \mathcal{E}\}.$$

Then

$$|\mathcal{E}| \leq \prod_{ab \in E(H)} (2|\mathcal{E}_{ab}|)^{1/\Delta}.$$

Proof. Let $\bar{\varphi}: V(H) \rightarrow V(G)$ be a uniformly chosen random element of \mathcal{E} . Write $H(Z)$ for the entropy of a discrete random variable Z and observe that $H(\bar{\varphi}) = \log |\mathcal{E}|$. Since H is Δ -regular, Shearer's inequality [22] implies that

$$H(\bar{\varphi}) \leq \frac{1}{\Delta} \cdot \sum_{ab \in E(H)} H(\bar{\varphi}(a), \bar{\varphi}(b)).$$

The random variable $(\bar{\varphi}(a), \bar{\varphi}(b))$ can take on at most $2|\mathcal{E}_{ab}|$ values, as it is an ordered pair of vertices that make up the edge $\varphi(ab)$. Using the fact that the entropy of any distribution on a set is at most that of the uniform distribution on that set, it follows that $H(\bar{\varphi}(a), \bar{\varphi}(b)) \leq \log(2|\mathcal{E}_{ab}|)$. This implies the assertion of the lemma. \square

Proof of Lemma 5.7. Given an ordered pair (c, d) of adjacent vertices of H , let $\mathcal{E}^{(c,d)}$ be the family of embeddings φ of H into G such that $\varphi(c) = u$ and $\varphi(d) = v$. For a given edge ab of H , define $\mathcal{E}_{ab}^{(c,d)} = \{\varphi(ab) : \varphi \in \mathcal{E}^{(c,d)}\}$ as in the statement of Lemma 5.10. Observe that

$$|\mathcal{E}_{ab}^{(c,d)}| \leq \begin{cases} e_G & \text{if } \{a, b\} \cap \{c, d\} = \emptyset, \\ \deg_G u & \text{if } \{a, b\} \cap \{c, d\} = \{c\}, \\ \deg_G v & \text{if } \{a, b\} \cap \{c, d\} = \{d\}, \\ 1 & \text{if } \{a, b\} = \{c, d\}. \end{cases}$$

Invoking Lemma 5.10 to bound $|\mathcal{E}^{(c,d)}|$ from above and summing over all $2e_H$ pairs (c, d) of adjacent vertices of H , the claimed upper bound on the number of embeddings of H into G that use the edge uv follows. \square

Proof of Lemma 5.8. We fix an edge ab of J , where $a \in A$ and $b \in B$, and count the embeddings φ of J into G such that $\varphi(ab) = uv$. To this end, we first show that $J - b$ contains a matching M that saturates A . To see this, note that for every nonempty $S \subseteq A$,

$$|S| \cdot \Delta = \sum_{c \in S} \deg_J c \leq \sum_{d \in N_J(S)} \deg_J d \leq |N_J(S)| \cdot \Delta,$$

yielding $|N_J(S)| \geq |S|$. Moreover, this inequality is strict unless the subgraph of J induced by $S \cup N_J(S)$ is Δ -regular. However, the latter is impossible since, as J is connected, the only Δ -regular subgraph of J could be J itself, but our assumption $|A| < |B|$ implies that J is not regular. Hence $|N_{J-b}(S)| \geq |N_J(S)| - 1 \geq |S|$, verifying Hall's condition. Now, given a matching $M \subseteq J - b$ that saturates A , we may bound the number of embeddings φ as above in the following way. Let c be the neighbour of a in M . There are at most $\deg_G u + \deg_G v$ embeddings of

$J[\{a, b, c\}]$ into G that map ab to uv . Each of them admits at most $(2e_G)^{|M|-1}$ extensions to an embedding of $J[V(M) \cup \{b\}]$. Each of those embeddings can be extended to an embedding of J in at most $\min\{e_G, v_G\}^{|B \setminus (\{b\} \cup V(M))|}$ ways. Since $|M| = |A|$ and $|B \setminus (\{b\} \cup V(M))| = |B| - |A| - 1$, summing over all $ab \in E(J)$ gives the claimed bound on $|\text{Emb}(J, G; uv)|$. \square

Proof of Lemma 5.9. Given an edge cd of H , let \mathcal{E}^{cd} be the family of embeddings of H into G that map cd onto an edge of G' . Define $\mathcal{E}_{ab}^{cd} = \{\varphi(ab) : \varphi \in \mathcal{E}^{cd}\}$ as in the statement of Lemma 5.10 and observe that

$$|\mathcal{E}_{ab}^{cd}| \leq \begin{cases} e_{G'} & \text{if } \{a, b\} = \{c, d\}, \\ e_G & \text{otherwise.} \end{cases}$$

Invoking Lemma 5.10 to bound $|\mathcal{E}^{cd}|$ from above and summing over all edges cd of H gives the claimed upper bound. \square

5.4. Stability results. Observe that, when J is the complete graph, then Theorem 5.4 yields the upper bound $|\text{Emb}(K_r, G)| \leq (2e_G)^{r/2}$. This is a weak version of a more precise result due to Erdős–Hanani [32] and also follows from the Kruskal–Katona theorem [50, 55]. One can see that the upper bound $(2e_G)^{r/2}$ is asymptotically optimal if G contains a clique comprising all of its edges. Our next theorem states that, when the upper bound given in Theorem 5.4 is nearly tight, then G resembles such a graph, in the sense that it must contain a subgraph of density $1 - o(1)$ covering nearly all of its edges. This could be proved by appealing to a stability version of the Kruskal–Katona theorem due to Keevash [51]. The proof we present below is elementary.

Theorem 5.11. *Suppose that $r \geq 3$. If a graph G satisfies*

$$|\text{Emb}(K_r, G)| \geq (1 - \varepsilon) \cdot (2e_G)^{r/2},$$

for some $\varepsilon \geq e_G^{-1/2}$, then G has a subgraph G' with minimum degree at least $(1 - 4\varepsilon^{1/2}) \cdot (2e_G)^{1/2}$.

Proof. The assertion of the theorem follows once we establish the case $r = 3$ and an analogous property for the path with four vertices (and three edges), which we denote by P_4 . Indeed, if r is odd, then K_r contains a spanning subgraph that is the disjoint union of K_3 and a matching of size $(r - 3)/2$. Thus,

$$|\text{Emb}(K_r, G)| \leq |\text{Emb}(K_3, G)| \cdot |\text{Emb}(K_2, G)|^{(r-3)/2} = |\text{Emb}(K_3, G)| \cdot (2e_G)^{(r-3)/2}$$

and hence $|\text{Emb}(K_3, G)| \geq (1 - \varepsilon) \cdot (2e_G)^{3/2}$. Analogously, if r is even, then K_r contains a subgraph that is the disjoint union of P_4 and a matching of size $(r - 4)/2$. Thus,

$$|\text{Emb}(K_r, G)| \leq |\text{Emb}(P_4, G)| \cdot |\text{Emb}(K_2, G)|^{(r-4)/2} = |\text{Emb}(P_4, G)| \cdot (2e_G)^{(r-4)/2},$$

which implies that $|\text{Emb}(P_4, G)| \geq (1 - \varepsilon) \cdot (2e_G)^2$. Therefore, it suffices to prove the following two claims.

Claim 5.12. *If $|\text{Emb}(P_4, G)| \geq (1 - \varepsilon) \cdot (2e_G)^2$, for some positive ε , then G has a subgraph G' with minimum degree at least $(1 - 2\varepsilon^{1/2})(2e_G)^{1/2}$.*

Claim 5.13. *If $|\text{Emb}(K_3, G)| \geq (1 - \varepsilon) \cdot (2e_G)^{3/2}$, for some $\varepsilon \geq e_G^{-1/2}$, then G has a subgraph G' with minimum degree at least $(1 - 4\varepsilon^{1/2})(2e_G)^{1/2}$.*

Proof of Claim 5.12. We may assume that $\varepsilon < 1/4$, as otherwise the assertion of the claim is trivially satisfied. Let F be the graph with vertex set $E(G)$ whose edges are all pairs $\{uv, xy\}$ such that the set $\{u, v, x, y\}$ induces a K_4 in G . Let $\mathbb{1}_G$ be the indicator function of the edge set

of G and note that

$$\begin{aligned} |\text{Emb}(P_4, G)| &= \sum_{\substack{\{uv, xy\} \subseteq E(G) \\ \{u, v\} \cap \{x, y\} = \emptyset}} 2 \cdot (\mathbb{1}_G(uv) + \mathbb{1}_G(xy) + \mathbb{1}_G(vx) + \mathbb{1}_G(vy)) \\ &\leq 8e_F + 6 \binom{e_G}{2} \leq 3e_G^2 + 2e_F. \end{aligned}$$

In particular, our assumption implies that $e_F \geq (1/2 - 2\varepsilon) \cdot e_G^2$.

Let F' be the subgraph obtained from F by iteratively removing vertices whose degree is smaller than $(1 - 2\varepsilon^{1/2}) \cdot e_G$. We claim that fewer than $2\varepsilon^{1/2} \cdot e_G$ vertices are removed this way and, consequently, the graph F' is nonempty and its minimum degree is at least $(1 - 2\varepsilon^{1/2}) \cdot e_G$. Suppose that this were not true. We would then have

$$e_F \leq \binom{(1 - 2\varepsilon^{1/2}) \cdot e_G}{2} + 2\varepsilon^{1/2} \cdot e_G \cdot (1 - 2\varepsilon^{1/2}) \cdot e_G < \left(\frac{1}{2} - 2\varepsilon\right) \cdot e_G^2 \leq e_F,$$

a contradiction.

Finally, let G' be the subgraph of G induced by the set of endpoints of the edges from $V(F')$. Let u be an arbitrary vertex of G' . There must be another vertex v of G' such that $uv \in V(F')$. Since $\deg_{F'} uv \geq (1 - 2\varepsilon^{1/2}) \cdot e_G$, the common neighbourhood of u and v in G' induces a subgraph with at least $(1 - 2\varepsilon^{1/2}) \cdot e_G$ edges in G . In particular,

$$\deg_{G'} u \geq \deg_{G'}(u, v) \geq \sqrt{2 \cdot (1 - 2\varepsilon^{1/2}) \cdot e_G} \geq (1 - 2\varepsilon^{1/2}) \cdot (2e_G)^{1/2}.$$

Since u was arbitrary, we obtain the desired lower bound on the minimum degree of G' . \square

Proof of Claim 5.13. We may assume that $\varepsilon < 1/16$, as otherwise the assertion of the claim is trivially satisfied. For every edge $e \in E(G)$, let t_e denote the number of copies of K_3 in G that contain the edge e . Observe that, for each $e \in E(G)$, there are at least $2t_e(t_e - 1)$ embeddings of P_4 into G that map the middle edge of P_4 onto e . Since $\sum_{e \in E(G)} t_e = |\text{Emb}(K_3, G)|/2$ and the function $t \mapsto 2t(t - 1)$ is convex, we conclude that

$$|\text{Emb}(P_4, G)| \geq \sum_{e \in E(G)} 2t_e(t_e - 1) \geq e_G \cdot \frac{|\text{Emb}(K_3, G)|}{e_G} \cdot \left(\frac{|\text{Emb}(K_3, G)|}{2e_G} - 1 \right).$$

Our assumptions imply that

$$1 \leq \varepsilon \cdot e_G^{1/2} \leq \varepsilon \cdot \frac{(1 - \varepsilon) \cdot (2e_G)^{3/2}}{2e_G} \leq \varepsilon \cdot \frac{|\text{Emb}(K_3, G)|}{2e_G}$$

and consequently,

$$|\text{Emb}(P_4, G)| \geq (1 - \varepsilon) \cdot \frac{|\text{Emb}(K_3, G)|^2}{2e_G} \geq (1 - \varepsilon)^3 \cdot (2e_G)^2 \geq (1 - 3\varepsilon) \cdot (2e_G)^2,$$

It now follows from Claim 5.12 that G contains a subgraph G' with minimum degree at least

$$(1 - 2 \cdot (3\varepsilon)^{1/2}) \cdot (2e_G)^{1/2} \geq (1 - 4\varepsilon^{1/2}) \cdot (2e_G)^{1/2},$$

as claimed. $\square \square$

Our next lemma gives a tight upper bound on the number of stars $K_{1,s}$ in a given bipartite graph, as well as a structural characterisation of the bipartite graphs that are close to achieving this bound. This lemma and Theorem 5.11 above constitute the main combinatorial ingredient in the proof of Theorem 1.8. Given a graph G and a set U of vertices of G , we let $\text{Emb}_U(K_{1,s}, G)$ denote the set of embeddings of $K_{1,s}$ into G that map the centre vertex to a vertex of U .

Lemma 5.14. *Let $s \geq 2$ be an integer and suppose that G is a bipartite graph with parts U and V and at most $q|V|$ edges, for some $q \in (0, |U|]$. Then the following holds:*

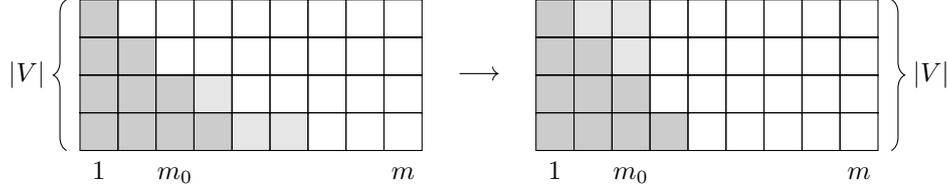


FIGURE 4. The degree sequence on the left can be turned into the degree sequence on the right by moving mass from columns $i \geq m_0 + 1$ to columns $j \leq m_0 + 1$.

(i)

$$|\text{Emb}_U(K_{1,s}, G)| \leq (\lfloor q \rfloor + \{q\}^s) |V|^s.$$

(ii) For every positive ε , there exists a positive η such that, if

$$|\text{Emb}_U(K_{1,s}, G)| \geq (1 - \eta) \cdot (\lfloor q \rfloor + \{q\}^s) |V|^s,$$

then there is a subset $W \subseteq U$ of size $\lceil q \rceil$ such that $e_G(W, V) \geq (1 - \varepsilon)q|V|$, and a further subset $W' \subseteq W$ of size at least $\lfloor (1 - \varepsilon)|W| \rfloor$ such that $\deg_G u \geq (1 - \varepsilon)|V|$ for every $u \in W'$.

Proof. We will use the following inequality, valid for any two numbers x and y with $x \geq y \geq 1$:

$$(x + 1)^s + (y - 1)^s - x^s - y^s \geq (x + 1 - y) \cdot (x + 1)^{s-2}. \quad (28)$$

It is clear that

$$|\text{Emb}_U(K_{1,s}, G)| = \sum_{u \in U} (\deg_G u)(\deg_G u - 1) \cdots (\deg_G u - s + 1) \leq \sum_{u \in U} (\deg_G u)^s. \quad (29)$$

Let $m = |U|$ and, given a sequence $\mathbf{d} = (d_1, \dots, d_m)$, define

$$S(\mathbf{d}) = \sum_{i=1}^m d_i^s.$$

By the degree sequence of a bipartite graph with parts U and V , we will mean the sequence of degrees of the vertices in U , listed in a nonincreasing order. Thus (29) implies that $|\text{Emb}_U(K_{1,s}, G)| \leq S(\mathbf{d}_G)$, where \mathbf{d}_G is the degree sequence of G . Let $m_0 = \lfloor e_G/|V| \rfloor$ and define $\mathbf{d}_{\max} = (d_1^*, \dots, d_m^*)$ by

$$d_i^* = \begin{cases} |V| & \text{if } i \leq m_0, \\ \{e_G/|V|\} \cdot |V| & \text{if } i = m_0 + 1, \\ 0 & \text{otherwise.} \end{cases}$$

Note that $\sum_{i=1}^m d_i^* = e_G$; in particular, $\{e_G/|V|\} \cdot |V|$ is an integer. We claim that \mathbf{d}_{\max} maximises S over all degree sequences whose sum is e_G . Indeed, for any other such degree sequence $\mathbf{d}' = (d'_1, \dots, d'_m)$, there must be two distinct indices i and j such that $0 < d'_i \leq d'_j < |V|$. Let \mathbf{d}'' be the degree sequence obtained from \mathbf{d}' by decreasing d'_i by one and increasing d'_j by one (and reordering the degrees, if necessary). It follows from (28) that

$$S(\mathbf{d}'') - S(\mathbf{d}') = (d'_j + 1)^s + (d'_i - 1)^s - (d'_j)^s - (d'_i)^s \geq (d'_j - d'_i + 1) \cdot (d'_j + 1)^{s-2} \geq 1.$$

Therefore,

$$|\text{Emb}_U(K_{1,s}, G)| \leq S(\mathbf{d}_G) \leq S(\mathbf{d}_{\max}) = (\lfloor e_G/|V| \rfloor + \{e_G/|V|\}^s) |V|^s.$$

Since $e_G \leq q|V|$, this completes the proof of the first part of the lemma.

For the second stipulation of the lemma, fix a positive ε . Let Ω be the set of degree sequences of bipartite graphs G with parts U and V for which there is a subset $W \subseteq U$ of size $\lceil e_G/|V| \rceil$ such that $e_G(W, V) \geq (1 - \varepsilon)e_G$ and, additionally, a further subset $W' \subseteq W$ of size $\lfloor (1 - \varepsilon)|W| \rfloor$ such that $\deg_G u \geq (1 - \varepsilon)|V|$ for each $u \in W'$.

It suffices to show that any degree sequence $\mathbf{d} \notin \Omega$ whose sum is e_G satisfies $S(\mathbf{d}) < (1 - \eta)S(\mathbf{d}_{\max})$, for some positive η that depends only on ε . Let $\mathbf{d} = (d_1, \dots, d_m)$. The crucial observation is that we may obtain \mathbf{d}_{\max} from \mathbf{d} by successively increasing some d_i with $i \leq m_0 + 1$ by one and, simultaneously, decreasing some d_j with $j \geq m_0 + 1$ by one (see Figure 4). Note that, when doing so, we perform at least

$$\sum_{i=1}^{m_0} (|V| - d_i) + \max \{d_{m_0+1}^* - d_{m_0+1}, 0\}$$

such operations. We split further analysis into two cases.

First, assume that $\sum_{i=1}^{m_0} (|V| - d_i) \geq \varepsilon^2 e_G$; this implies that $m_0 \geq 1$. In this case, in at least $\varepsilon^2 e_G / 2$ steps of the above procedure, we will be increasing some d_i with $i \leq m_0$ which is, at this time, already at least $\lfloor (|V| + d_i) / 2 \rfloor \geq (|V| + d_{m_0}) / 2 - 1$, while decreasing some d_j with $j \geq m_0 + 1$ which is at most d_{m_0} . Inequality (28) implies that

$$S(\mathbf{d}_{\max}) - S(\mathbf{d}) \geq \frac{\varepsilon^2 e_G}{2} \cdot \left(\frac{|V| + d_{m_0}}{2} - d_{m_0} \right) \cdot \left(\frac{|V| + d_{m_0}}{2} \right)^{s-2}.$$

Since our assumption implies that $|V| - d_{m_0} \geq \varepsilon^2 e_G / m_0$, it follows that

$$S(\mathbf{d}_{\max}) - S(\mathbf{d}) \geq \frac{\varepsilon^2 e_G}{2} \cdot \frac{\varepsilon^2 e_G}{2m_0} \cdot \left(\frac{|V|}{2} \right)^{s-2} = \frac{\varepsilon^4}{2^{s+1}} \cdot \left(\frac{e_G}{m_0 |V|} \right)^2 \cdot 2m_0 |V|^s.$$

Finally, we note that $e_G \geq m_0 |V|$ and that $S(\mathbf{d}_{\max}) \leq (m_0 + 1) |V|^s \leq 2m_0 |V|^s$, and thus

$$S(\mathbf{d}_{\max}) - S(\mathbf{d}) \geq \frac{\varepsilon^4}{2^{s+1}} \cdot S(\mathbf{d}_{\max}),$$

proving $S(\mathbf{d}) < (1 - \eta)S(\mathbf{d}_{\max})$ for some positive η .

Assume now that $\sum_{i=1}^{m_0} (|V| - d_i) < \varepsilon^2 e_G$. Let W be the set comprising the $\lceil e_G / |V| \rceil$ vertices with largest degrees in U . Suppose first that $e_G(W, V) \geq (1 - \varepsilon)e_G$, but every set of $\lfloor (1 - \varepsilon)|W| \rfloor$ vertices of W contains a vertex of degree smaller than $(1 - \varepsilon)|V|$. In this case, $(1 - \varepsilon)|W| \geq 1$, as otherwise the latter condition is vacuously false, and $d_i < (1 - \varepsilon)|V|$ whenever $i \geq \lfloor (1 - \varepsilon)|W| \rfloor$. Then

$$\sum_{i=1}^{m_0} (|V| - d_i) > (m_0 + 1 - \lfloor (1 - \varepsilon)|W| \rfloor) \cdot \varepsilon |V| \geq (|W| - \lfloor (1 - \varepsilon)|W| \rfloor) \cdot \varepsilon |V| \geq \varepsilon^2 e_G,$$

contradicting our assumption. Thus, since $\mathbf{d} \notin \Omega$, we may assume that $e_G(W, V) < (1 - \varepsilon)e_G$. Therefore,

$$(1 - \varepsilon)e_G > \sum_{i=1}^{|W|} d_i \geq \sum_{i=1}^{m_0} d_i = m_0 |V| - \sum_{i=1}^{m_0} (|V| - d_i) > m_0 |V| - \varepsilon^2 e_G.$$

This means, in particular, that $m_0 < e_G / |V|$ and hence $|W| = m_0 + 1$. Moreover,

$$m_0 |V| + d_{m_0+1}^* = e_G > \sum_{i=1}^{|W|} d_i + \varepsilon e_G > m_0 |V| + d_{m_0+1} + (\varepsilon - \varepsilon^2)e_G,$$

which implies that $d_{m_0+1}^* - d_{m_0+1} > (\varepsilon - \varepsilon^2)e_G$. Therefore, there exist at least $(\varepsilon - \varepsilon^2)e_G / 2$ steps in which we increase d_{m_0+1} at a time where it is already at least $\lfloor (\varepsilon - \varepsilon^2)e_G / 2 \rfloor$. Inequality (28) implies that

$$S(\mathbf{d}_{\max}) - S(\mathbf{d}) \geq \left(\frac{(\varepsilon - \varepsilon^2)e_G}{2} \right)^s.$$

However, we trivially have $S(\mathbf{d}_{\max}) \leq e_G^s$ and thus

$$S(\mathbf{d}_{\max}) - S(\mathbf{d}) \geq \frac{(\varepsilon - \varepsilon^2)^s}{2^s} \cdot S(\mathbf{d}_{\max}),$$

completing the proof. \square

We remark that the extremal structures given by Theorem 5.11 and Lemma 5.14(ii) are quite different and, in a sense, incompatible. This has the following technically important consequence: if a graph G simultaneously contains many copies of K_r and many copies of $K_{1,r-1}$, then it can be split into two edge-disjoint graphs, one containing nearly all the copies of K_r and the other containing nearly all the copies of $K_{1,r-1}$. The following lemma formalises this statement; its proof is similar to an argument of Lubetzky and Zhao [57]. We write $G[U, V]$ for the bipartite induced subgraph of G with parts U and V .

Lemma 5.15. *For every integer $r \geq 3$ and positive real number ε , there is a positive η such that the following holds. Let G be a graph on n vertices with $e_G \leq \eta n^2$. Then there exists a partition $V(G) = U \cup V$ satisfying $|U| \leq \varepsilon n$,*

$$|\text{Emb}(K_r, G[V])| \geq |\text{Emb}(K_r, G)| - \varepsilon e_G^{r/2},$$

and

$$|\text{Emb}_U(K_{1,r-1}, G[U, V])| \geq |\text{Emb}(K_{1,r-1}, G)| - \varepsilon e_G n^{r-2}.$$

Proof. Assume that η is sufficiently small, let U be the set of vertices in G with degree at least $\eta^{1/3}n$, and let $V = V(G) \setminus U$. Note that $|U| \leq 2e_G/(\eta^{1/3}n) \leq 2\eta^{2/3}n \leq \varepsilon n$.

Every embedding of K_r into G that maps a vertex of K_r to a vertex of U can be specified by first choosing a vertex v of K_r , then a vertex of U that v will be mapped to, and finally an embedding of K_{r-1} into G . Using Theorem 5.4, we thus obtain

$$|\text{Emb}(K_r, G)| - |\text{Emb}(K_r, G[V])| \leq r \cdot |U| \cdot |\text{Emb}(K_{r-1}, G)| \leq r \cdot \frac{2e_G}{\eta^{1/3}n} \cdot (2e_G)^{(r-1)/2}.$$

Since $e_G^{1/2} \leq \eta^{1/2}n$, this implies the first assertion of the lemma.

Next, note that

$$|\text{Emb}(K_{1,r-1}, G)| = |\text{Emb}_U(K_{1,r-1}, G[U, V])| + t_1 + t_2, \quad (30)$$

where t_1 is the number of embeddings of $K_{1,r-1}$ into G that map the centre vertex and at least one leaf of $K_{1,r-1}$ to U and t_2 is the number of embeddings that map the centre vertex of $K_{1,r-1}$ to V . We have

$$t_1 \leq (r-1) \cdot |U|^2 \cdot n^{r-2} \leq (r-1) \left(\frac{2e_G}{\eta^{1/3}n} \right)^2 n^{r-2} \leq \varepsilon e_G n^{r-2}/2,$$

as $e_G \leq \eta n^2$. Finally, in order to bound t_2 , observe that every embedding counted by t_2 can be specified by first choosing a leaf a of $K_{1,r-1}$, then choosing the image e of the edge of $K_{1,r-1}$ incident with a , then choosing the endpoint $v \in V$ of e that is the image of the centre vertex of $K_{1,r-1}$, and finally choosing the images of the remaining $r-2$ leaves of $K_{1,r-1}$ among the neighbours of v in G . Since every vertex $v \in V$ has degree at most $\eta^{1/3}n$, it follows that

$$t_2 \leq (r-1) \cdot e_G \cdot 2 \cdot (\eta^{1/3}n)^{r-2} \leq \varepsilon e_G n^{r-2}/2.$$

Together with (30), these bounds on t_1 and t_2 imply the second assertion of the lemma. \square

6. CLIQUES IN RANDOM GRAPHS

Fix an integer $r \geq 3$ and let $X = X_{n,p}^{K_r}$ be the number of r -vertex cliques in the random graph $G_{n,p}$. In this section, we shall use Theorem 3.1 not only to determine the logarithmic upper tail probability of X but also to provide a detailed description of the upper tail event. Before we restate the two theorems that will be proved in this section, we discuss the combinatorial constructions that are responsible for the localisation phenomenon in more detail.

As was shown in [57], when $n^{-1} \ll p^{(r-1)/2} \ll 1$, there are essentially two optimal strategies for planting a subgraph inside $G_{n,p}$ that increases the expected number of copies of K_r by the required $\delta \mathbb{E}[X]$. The first, and most straightforward, involves planting a clique with $\delta^{1/r} np^{(r-1)/2}$

vertices; note that our assumption on p implies that this expression is unbounded and thus we may implicitly assume that it is an integer. Note that such a clique has close to

$$\frac{\delta^{2/r} n^2 p^{r-1}}{2}$$

edges and contains approximately $\delta \binom{n}{r} p^{\binom{r}{2}} = \delta \mathbb{E}[X]$ copies of K_r . If np^{r-1} is bounded from below, then there is an alternative strategy that competes with planting a clique. By a *hub* of order $\lfloor \delta np^{r-1}/r \rfloor + 1$, we mean a subgraph of $G_{n,p}$ constructed as follows. Let U be a set of $\lfloor \delta np^{r-1}/r \rfloor$ vertices of $G_{n,p}$ and let u be another vertex that lies outside of U . Connect every vertex in U to every vertex outside of $U \cup \{u\}$ and connect u to some $\{\delta np^{r-1}/r\}^{1/(r-1)} \cdot n$ such vertices. Note that every hub has close to

$$\left(\lfloor \delta np^{r-1}/r \rfloor + \{\delta np^{r-1}/r\}^{1/(r-1)} \right) \cdot n$$

edges, which is $\Theta(n^2 p^{r-1})$, as np^{r-1} is bounded from below. Unlike in the previous construction, the hub itself contains *no* copies of K_r . However, as $p \ll 1$, planting a hub creates approximately $\lfloor \delta np^{r-1}/r \rfloor \cdot \binom{n}{r-1}$ copies of the star graph $K_{1,r-1}$ whose centre vertex lies in U and approximately $\{\delta np^{r-1}/r\} \cdot \binom{n}{r-1}$ copies of $K_{1,r-1}$ whose centre vertex is u . The total number of planted copies of $K_{1,r-1}$ is thus approximately $\delta p^{r-1} \cdot \binom{n}{r} = \delta \mathbb{E}[X] \cdot p^{-\binom{r-1}{2}}$. Since each of the planted copies of $K_{1,r-1}$ lies in a copy of K_r that now appears in $G_{n,p}$ with probability $p^{\binom{r-1}{2}}$, one expects to see approximately $\delta \mathbb{E}[X]$ such extra copies of K_r . We remark that if np^{r-1} is large, then the contribution of the single vertex u becomes negligible and the hub construction can be described more concisely as connecting some $\delta np^{r-1}/r$ vertices to all the others, using $\delta n^2 p^{r-1}/r$ edges.

We prove that, for a vast majority of values of p in the range of interest, the logarithmic upper tail probability of X corresponds to one of the two strategies described above. In particular, we show that

- (i) If $np^{r-1} \rightarrow 0$, then the logarithmic upper tail probability is asymptotically equal to the ‘cost’ of planting the smallest clique that has $\delta \mathbb{E}[X]$ copies of K_r .
- (ii) If $np^{r-1} \rightarrow \infty$, then the logarithmic upper tail probability is asymptotically equal to the ‘cost’ of planting either a clique as above (when $\delta^{2/r} \leq \delta/r$) or the smallest hub that has $\delta \mathbb{E}[X] \cdot p^{-\binom{r-1}{2}}$ copies of $K_{1,r-1}$ (when $\delta^{2/r} \geq \delta/r$).

Note that in the regime $np^{r-1} \rightarrow \infty$, we may approximate the number of edges planted in the hub construction by $\delta n^2 p^{r-1}/r$. However, when $np^{r-1} \rightarrow c$ for some constant $c \in (0, \infty)$, this approximation is no longer valid and we are forced to account for the lack of smoothness that stems from the integral and fractional parts of $\delta c/r$. As a result, we find that, for certain values of the parameters δ and c , the logarithmic upper tail probability corresponds to a mixture of the first and second strategies: it is equal to the cost of planting a graph comprising both a hub and a clique, each contributing a nonnegligible proportion of the (expected) extra $\delta \mathbb{E}[X]$ copies of K_r .

Finally, suppose that one conditions $G_{n,p}$ on the upper tail event $\{X \geq (1 + \delta) \mathbb{E}[X]\}$. We prove that, with probability close to one, the conditioned random graph contains a subgraph that very closely resembles the graph described by the optimal strategy (for the particular values of n , p , and δ). For example, in cases where the logarithmic upper tail probability corresponds to planting a clique, we show that $G_{n,p}$ conditioned on the event $\{X \geq (1 + \delta) \mathbb{E}[X]\}$ contains a set of $\delta^{1/r} n p^{(r-1)/2}$ vertices that induces an ‘almost-clique’, that is, a subgraph of density $1 - o(1)$.

We now turn to the details. As in the introduction, we define continuous functions $\psi_r: (0, \infty)^2 \times [0, 1] \rightarrow (0, \infty)$ and $\varphi_r: (0, \infty)^2 \rightarrow (0, \infty)$ by

$$\psi_r(\delta, c, x) = \frac{(\delta(1-x))^{2/r}}{2} + \frac{\lfloor x\delta c/r \rfloor + \{x\delta c/r\}^{\frac{1}{r-1}}}{c} \quad \text{and} \quad \varphi_r(\delta, c) = \min_{x \in [0,1]} \psi_r(\delta, c, x).$$

Note that $\psi_r(\delta, np^{r-1}, x) \cdot n^2 p^{r-1}$ is approximately the number of edges in the disjoint union of a clique with $\delta(1-x)np^{(r-1)/2}$ vertices and a hub of order $\lfloor \delta x np^{r-1}/r \rfloor + 1$. Recalling the discussion above, $\varphi_r(\delta, np^{r-1}) \cdot n^2 p^{r-1}$ then represents the smallest number of edges among all combinations of clique and hub that yield an expected $\delta \mathbb{E}[X]$ copies of K_{r-1} .

In order to handle the three cases $np^{r-1} \rightarrow 0$, $np^{r-1} \rightarrow c \in (0, \infty)$, and $np^{r-1} \rightarrow \infty$ in a unified manner, it will be convenient to extend φ_r to a continuous function $\varphi_r: (0, \infty) \times [0, \infty] \rightarrow (0, \infty)$. This extension may be defined by noting that

$$\lim_{c \rightarrow 0} \varphi_r(\delta, c) = \delta^{2/r}/2 \quad \text{and} \quad \lim_{c \rightarrow \infty} \varphi_r(\delta, c) = \min\{\delta^{2/r}/2, \delta/r\}$$

uniformly as functions of δ . For every $\delta > 0$ and $c \in [0, \infty]$, we then define the set

$$\bar{X}_r(\delta, c) = \{x \in [0, 1] : \varphi_r(\delta, c) = \lim_{c' \rightarrow c} \psi_r(\delta, c', x)\}$$

of (asymptotic) minimisers to $x \mapsto \psi_r(\delta, c, x)$. One can check that this set is nonempty for any δ and c , though it might contain more than one element. The following lemma describes the set of possible minimisers in more detail.

Lemma 6.1. *Let $r \geq 3$ be an integer and let δ and c be positive real numbers. Let $x^* = r \lfloor \delta c/r \rfloor / (\delta c)$. Then*

$$\bar{X}_r(\delta, 0) = \{0\}, \quad \bar{X}_r(\delta, c) \subseteq \{0, x^*, 1\}, \quad \text{and} \quad \bar{X}_r(\delta, \infty) \subseteq \{0, 1\}.$$

Proof. The statement for $\bar{X}_r(\delta, 0)$ holds because $\lim_{c \rightarrow 0} \psi_r(\delta, c, x) = \infty$ whenever $x > 0$. The statement for $\bar{X}_r(\delta, \infty)$ follows because

$$\lim_{c \rightarrow \infty} \psi_r(\delta, c, x) = (\delta(1-x))^{2/r}/2 + x\delta/r$$

and the right-hand side is strictly concave in $x \in [0, 1]$. In particular, since

$$\varphi_r(\delta, \infty) = \min\{\delta^{2/r}, \delta/r\} = \min_{x \in [0, 1]} (\delta(1-x))^{2/r}/2 + x\delta/r,$$

the equality $\varphi_r(\delta, \infty) = \lim_{c \rightarrow \infty} \psi_r(\delta, c, x)$ implies that x is one of the endpoints of the interval $[0, 1]$. Assume now that $c \in (0, \infty)$. Treating r , δ , and c as fixed, consider the function $f: [0, 1] \rightarrow (0, \infty)$ defined by $f(x) = \psi_r(\delta, c, x)$. Since ψ_r is continuous, we have $x \in \bar{X}_r(\delta, c)$ if and only if x is a minimiser of $f(x)$ in $[0, 1]$. We cover the domain of f with essentially disjoint intervals as follows:

$$[0, 1] = \frac{r}{\delta c} [0, 1] \cup \frac{r}{\delta c} [1, 2] \cup \dots \cup \frac{r}{\delta c} [\lfloor \delta c/r \rfloor - 1, \lfloor \delta c/r \rfloor] \cup \frac{r}{\delta c} [\lfloor \delta c/r \rfloor, \delta c/r].$$

Observe that f is strictly concave on each of these intervals and thus f can achieve its minimum value only at an endpoint of one of the intervals, that is, either at some x for which $\delta x c/r \in \mathbb{Z}$ or at $x = 1$. Define the function $h: [0, 1] \rightarrow (0, \infty)$ by

$$h(x) = \frac{(\delta(1-x))^{2/r}}{2} + \frac{x\delta}{r}$$

and observe that h is strictly concave and that $h(x) = f(x)$ whenever $\delta x c/r \in \mathbb{Z}$. It follows that the minimum of $f(x)$ is achieved either at $x = 1$ or at the smallest or largest x satisfying $\delta x c/r \in \mathbb{Z}$. Since the latter two points are $x = 0$ and $x = r \lfloor \delta c/r \rfloor / (\delta c) = x^*$, respectively, we may conclude that the minimiser is in the set $\{0, x^*, 1\}$, as desired. \square

Let us now state the two main results of this section. The following is a straightforward reformulation of Theorem 1.7.

Theorem 6.2. *Let $r \geq 3$ be an integer and let $X = X_{n,p}^{K_r}$ denote the number of r -vertex cliques in the random graph $G_{n,p}$. Then, for every fixed positive constant δ and all $p = p(n)$ such that $n^{-1}(\log n)^{\frac{1}{r-2}} \ll p^{\frac{r-1}{2}} \ll 1$ and $\lim_{n \rightarrow \infty} np^{r-1} = c \in [0, \infty]$,*

$$\lim_{n \rightarrow \infty} \frac{-\log \mathbb{P}(X \geq (1 + \delta) \mathbb{E}[X])}{n^2 p^{r-1} \log(1/p)} = \varphi_r(\delta, c).$$

Recall the following three events:

- (i) Let $\text{UT}(\delta)$ be the upper tail event $\{X \geq (1 + \delta) \mathbb{E}[X]\}$.
- (ii) Let $\text{Clique}_\varepsilon(x)$ be the event that $G_{n,p}$ contains a set $U \subseteq [n]$ of size at least $(1 - \varepsilon)x^{1/r} np^{(r-1)/2}$ such that $G_{n,p}[U]$ has minimum degree at least $(1 - \varepsilon)|U|$.
- (iii) Let $\text{Hub}_\varepsilon(x)$ be the event that $G_{n,p}$ contains a set $U \subseteq [n]$ such that at least $\lfloor (1 - \varepsilon)|U| \rfloor$ vertices in U have degree at least $(1 - \varepsilon)n$ and

$$e(U, [n] \setminus U) \geq (1 - \varepsilon)n(\lfloor xnp^{r-1}/r \rfloor + \{xnp^{r-1}/r\}^{\frac{1}{r-1}}).$$

Together with Lemma 6.1, the following theorem directly implies Theorem 1.8.

Theorem 6.3. *Let $r \geq 3$ be an integer and let $X = X_{n,p}^{K_r}$ denote the number of r -vertex cliques in the random graph $G_{n,p}$. Then, for every fixed positive constant δ and all $p = p(n)$ such that $n^{-1}(\log n)^{\frac{1}{r-2}} \ll p^{\frac{r-1}{2}} \ll 1$ and $\lim_{n \rightarrow \infty} np^{r-1} = c \in [0, \infty]$,*

$$\lim_{n \rightarrow \infty} \mathbb{P}\left(\bigcup_{x \in \bar{X}(\delta, c)} \text{Clique}_\varepsilon(\delta(1-x)) \cap \text{Hub}_\varepsilon(\delta x) \mid \text{UT}(\delta)\right) = 1.$$

In order to prove Theorems 6.2 and 6.3, we will first relate $-\log \mathbb{P}(X \geq (1 + \delta) \mathbb{E}[X])$ to the solutions of the optimisation problem

$$\Phi_X(\delta) = \min \{e_G \log(1/p) : G \subseteq K_n \text{ and } \mathbb{E}_G[X] \geq (1 + \delta) \mathbb{E}[X]\},$$

where $\mathbb{E}_G[X] = \mathbb{E}[X \mid G \subseteq G_{n,p}]$. For every $\varepsilon > 0$, we define the event

$$\text{Near-min}(\varepsilon) = \{G_{n,p} \text{ contains a subgraph } G \text{ such that } e_G \log(1/p) \leq (1 + \varepsilon)\Phi_X(\delta + \varepsilon) \text{ and } \mathbb{E}_G[X] \geq (1 + \delta - \varepsilon) \mathbb{E}[X]\}.$$

Using Theorem 3.1, we shall prove the following result.

Proposition 6.4. *For every integer $r \geq 3$ and all positive reals ε and δ , there exists a positive constant C such that the following holds. Suppose that an integer n and $p \in (0, 1)$ satisfy $Cn^{-1}(\log n)^{1/(r-2)} \leq p^{(r-1)/2} \leq 1/C$. Then $X = X_{n,p}^{K_r}$ satisfies*

$$(1 - \varepsilon)\Phi_X(\delta - \varepsilon) \leq -\log \mathbb{P}(X \geq (1 + \delta) \mathbb{E}[X]) \leq (1 + \varepsilon)\Phi_X(\delta + \varepsilon).$$

Furthermore,

$$\mathbb{P}(\text{Near-min}(\varepsilon) \mid X \geq (1 + \delta) \mathbb{E}[X]) \geq 1 - \varepsilon.$$

In order to complete the proof of Theorem 6.2, we must also evaluate the asymptotic value of the function Φ_X . This is the content of our second proposition.

Proposition 6.5. *Let $r \geq 3$ be an integer and let $X = X_{n,p}^{K_r}$. Then, for every fixed positive real δ and all $p = p(n)$ such that $n^{-1} \ll p^{(r-1)/2} \ll 1$ and $\lim_{n \rightarrow \infty} np^{r-1} = c \in [0, \infty]$,*

$$\lim_{n \rightarrow \infty} \frac{\Phi_X(\delta)}{n^2 p^{r-1} \log(1/p)} = \varphi_r(\delta, c).$$

We recall that, when $c \in \{0, \infty\}$, this result was already established in [57]. The proof we provide is not unsimilar to that work, although it is slightly easier, as Φ_X is a discrete restriction of the variational problem in their work (see the discussion following Proposition 4.2, which discusses the same issue in the case of arithmetic progressions). However, the extension to the case $c \in (0, \infty)$ is delicate, and requires the use of the more precise result of Lemma 5.14.

We would like to emphasise that, while the proofs of Propositions 6.4 and 6.5 have certain similarities, they are conceptually very different. The analysis of the variational problem in the proof of Proposition 6.5 relies on bounding the maximal number of embeddings of K_r and $K_{1,r-1}$ in a graph with given numbers of vertices and edges, using the results of Section 5.2. In contrast, the proof of Proposition 6.4 exploits the bounds on the number of embeddings of K_r and $K_{1,r-1}$ (as well as several other key graphs) containing a given edge of the core, proved in Section 5.3. The defining properties of a core allow us to translate these upper bounds on the number of local embeddings into lower bounds on the degrees of the endpoints of edges of a core. These degree restrictions are sufficient to prove that the family of cores is entropically stable.

Finally, to prove Theorem 6.3, we characterise the near-minimisers of the optimisation problem for $\Phi_X(\delta)$. This is what we do in the final proposition of this section.

Proposition 6.6. *Let $r \geq 3$ be an integer and let $X = X_{n,p}^{K_r}$. For all fixed $\varepsilon, \delta > 0$ and $c \in [0, \infty]$, there exists some positive constant η such that the following holds. Assume $p = p(n)$ is such that $n^{-1} \ll p^{(r-1)/2} \ll 1$ and $\lim_{n \rightarrow \infty} np^{r-1} = c$. Then*

$$\text{Near-min}(\eta) \subseteq \bigcup_{x \in \bar{X}(\delta, c)} \text{Clique}_\varepsilon(\delta(1-x)) \cap \text{Hub}_\varepsilon(\delta x)$$

whenever n is sufficiently large.

The propositions above readily imply Theorems 6.2 and 6.3.

6.1. Proof of Proposition 6.4. We may assume without loss of generality that ε is sufficiently small, say $\varepsilon < \min\{1/2, \delta/2\}$. Note also that the case $n < r$ is trivial; indeed, in that case X is identically zero and thus $-\log \mathbb{P}(X \geq (1+\delta)\mathbb{E}[X]) = 0 = \Phi_X(\delta)$ for every $\delta \in \mathbb{R}$, and $\text{Near-min}(\varepsilon)$ holds vacuously. We may therefore assume that $n \geq r \geq 3$, which, in turn, implies that $n \geq C$.

Set $N = \binom{n}{2}$ and let $Y = (Y_1, \dots, Y_N)$ be the sequence of indicator random variables of the events that $e \in E(G_{n,p})$, where e ranges over $\binom{[n]}{2}$ in some arbitrary order. Observe that Y is a vector of independent $\text{Ber}(p)$ random variables and that X is a nonzero polynomial with nonnegative coefficients and degree at most $\binom{r}{2}$ in the coordinates of Y . Let $K = K(r, \varepsilon, \delta)$ be the constant given by Theorem 3.1. We shall show that X satisfies the various assumptions of the theorem; the theorem then implies both assertions of the proposition.

First, our assumption on p implies that $p \leq 1 - \varepsilon$, provided that C is sufficiently large.

Recall that $N(J, G)$ denotes the number of copies of J in G . Note that for all $J \subseteq K_r$ without isolated vertices and all $G \subseteq K_n$, we can trivially bound $N(J, G)$ from above by $e_G^{e_J}$. It follows that

$$\begin{aligned} \mathbb{E}_G[X] - \mathbb{E}[X] &\leq \sum_{\emptyset \neq J \subseteq K_r} N(J, G) \cdot n^{r-v_J} p^{\binom{r}{2} - e_J} \leq 2^{\binom{r}{2}} \cdot \max_{\emptyset \neq J \subseteq K_r} e_G^{e_J} \cdot n^{r-v_J} p^{\binom{r}{2} - \binom{v_J}{2}} \\ &\leq (2e_G)^{\binom{r}{2}} \cdot \frac{n^r p^{\binom{r}{2}}}{\min_{2 \leq k \leq r} n^k p^{\binom{k}{2}}}, \end{aligned}$$

where the sum ranges over all nonempty subgraphs $J \subseteq K_r$ without isolated vertices. Since $\mathbb{E}[X] = \Theta(n^r p^{\binom{r}{2}})$ and our assumption on p implies that $n^k p^{\binom{k}{2}} \geq C^{2/(r-1)}$ for each $k \in \{2, \dots, r\}$, the right-hand side above is at most $(\delta/2)\mathbb{E}[X]$ unless $e_G \geq K$, provided that C is sufficiently large. Therefore, $\Phi_X(\delta - \varepsilon) \geq \Phi_X(\delta/2) \geq K \log(1/p)$. Furthermore, since a clique with $\lceil (1+2\delta)^{1/r} np^{(r-1)/2} \rceil$ vertices contains at least $(1+\delta+\varepsilon)\mathbb{E}[X]$ copies of K_r and has fewer than $K'n^2 p^{r-1}$ edges, for some constant $K' = K'(\delta)$, we deduce that $\Phi_X(\delta + \varepsilon) \leq K'n^2 p^{r-1} \log(1/p)$.

Recall that a graph $G^* \subseteq K_n$ is a *core* if it satisfies the following three conditions:

- (C1) $\mathbb{E}_{G^*}[X] \geq (1+\delta-\varepsilon)\mathbb{E}[X]$,
- (C2) $e_{G^*} \leq K \cdot \Phi_X(\delta + \varepsilon)$, and

$$(C3) \quad \min_{e \in E(G^*)} (\mathbb{E}_{G^*}[X] - \mathbb{E}_{G^* \setminus e}[X]) \geq \mathbb{E}[X] / (K \cdot \Phi_X(\delta + \varepsilon)).$$

Our goal is to show that, for every integer m , there are at most $(1/p)^{\varepsilon m/2}$ cores with m edges. The key observation is that, for any core G^* and any edge uv of G^* , either u and v have many common neighbours or the sum of the degrees of u and v is large. More precisely, letting $\deg_{G^*}(u, v)$ denote the number of common neighbours of u and v in G^* , we shall establish the following statement.

Claim 6.7. *There exists a positive constant $\eta = \eta(\delta, r, K)$ such that, for every core G^* and each edge $uv \in E(G^*)$, either*

$$\deg_{G^*}(u, v) \geq \frac{\eta n p^{(r-1)/2}}{(\log(1/p))^{1/(r-2)}} \quad \text{or} \quad \deg_{G^*} u + \deg_{G^*} v \geq \frac{\eta n}{\log(1/p)}.$$

The above claim readily implies the desired bound on the number of cores with m edges. Note first that this number is zero whenever $m > KK'n^2 p^{r-1} \log(1/p)$, see (C2), so we may assume that $m \leq KK'n^2 p^{r-1} \log(1/p)$. Given a core G^* , we denote by A_{G^*} the set of vertices of G^* with degree at least $\eta n p^{(r-1)/2} / (\log(1/p))^{1/(r-2)}$ and by $B_{G^*} \subseteq A_{G^*}$ the set of vertices of G^* with degree at least $\eta n / \log(1/p)$. Since G^* has m edges,

$$|A_{G^*}| \leq a := \frac{2m (\log(1/p))^{1/(r-2)}}{\eta n p^{(r-1)/2}} \quad \text{and} \quad |B_{G^*}| \leq b := \frac{4m \log(1/p)}{\eta n}$$

Claim 6.7 states that every edge of G^* is either fully contained in A_{G^*} or has at least one endpoint in B_{G^*} . In particular, for fixed sets $B \subseteq A \subseteq \llbracket n \rrbracket$ with $|A| = a$ and $|B| = b$, there are at most $\binom{a^2/2 + bn}{m}$ cores G^* with m edges that satisfy $A_{G^*} \subseteq A$ and $B_{G^*} \subseteq B$. We can thus (generously) upper bound the number of cores with m edges by

$$\binom{n}{a} \binom{n}{b} \binom{a^2/2 + bn}{m}.$$

Recalling the inequality $\binom{x}{y} \leq (ex/y)^y$, valid for all nonnegative integers x and y , we may conclude that the number of cores with m edges is at most

$$n^{\frac{6m(\log(1/p))^{1/(r-2)}}{\eta n p^{(r-1)/2}}} \cdot \left(\frac{2em (\log(1/p))^{2/(r-2)}}{\eta^2 n^2 p^{r-1}} + \frac{4e \log(1/p)}{\eta} \right)^m.$$

Since $p^{(r-1)/2} \geq Cn^{-1} (\log n)^{1/(r-2)}$, the first factor is at most $e^{\varepsilon m \log(1/p)/4}$. Since we have assumed that $m \leq KK'n^2 p^{r-1} \log(1/p)$ and $1/p \geq C^{2/(r-1)}$, the second factor is at most $e^{O(m \log \log(1/p))} \leq e^{\varepsilon m \log(1/p)/4}$. This shows that the number of cores with m edges is indeed at most $(1/p)^{\varepsilon m/2}$, as claimed.

Proof of Claim 6.7. For any nonempty $J \subseteq K_r$, we shall let $N(J, G^*; uv)$ denote the number of copies of J in G^* that contain the edge uv . For the sake of brevity, we set $m_{\max} = KK' \cdot n^2 p^{r-1} \log(1/p)$. Observe that

$$\mathbb{E}_{G^*}[X] - \mathbb{E}_{G^* \setminus uv}[X] \leq \sum_{\emptyset \neq J \subseteq K_r} N(J, G^*; uv) \cdot n^{r-v_J} p^{\binom{r}{2} - e_J},$$

where J ranges over all nonempty subgraphs of K_r that have no isolated vertices. Since $\mathbb{E}[X] = \binom{n}{r} p^{\binom{r}{2}}$ and $\Phi_X(\delta + \varepsilon) \leq K'n^2 p^{r-1} \log(1/p) = m_{\max}/K$, it follows from (C3) in the definition of the core that

$$\sum_{\emptyset \neq J \subseteq K_r} \frac{N(J, G^*; uv)}{n^{v_J} p^{e_J}} \geq \frac{\mathbb{E}[X]}{K \cdot \Phi_X(\delta + \varepsilon)} \geq \frac{\gamma}{m_{\max}},$$

where $\gamma = \gamma(r)$ is a constant that depends only on r .

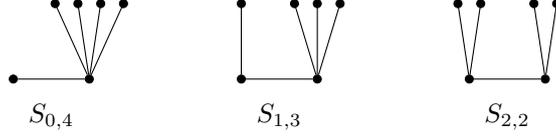


FIGURE 5. The double stars with six vertices.

For every edge $ab \in E(J)$, let $\text{Emb}(J, G^*; ab, uv)$ denote the set of embeddings of J into G^* that map ab to uv . Then the above inequality implies that there is a nonempty $J \subseteq K_r$ with no isolated vertices, an edge $ab \in E(J)$, and a constant $\gamma' = \gamma'(r)$ such that

$$\frac{|\text{Emb}(J, G^*; ab, uv)|}{n^{v_J} p^{e_J}} \geq \frac{\gamma'}{m_{\max}}. \quad (31)$$

For nonnegative integers i and j , let $S_{i,j}$ denote the graph obtained from a copy of $K_{1,i}$ and a copy of $K_{1,j}$ by joining their centres (vertices of degrees i and j , respectively) by an edge; see Figure 5 for an illustration. As the graphs $K_{1,i}$ are often called *stars*, we shall refer to the $S_{i,j}$ as *double stars*. Moreover, we shall call an edge of $S_{i,j}$ whose endpoints have degrees $i+1$ and $j+1$ a *centre edge*. Note that if $i, j > 0$, then $S_{i,j}$ has only one centre edge; otherwise, $S_{i,j}$ is just a star graph and each of its edges is a centre edge.

We shall now show that, unless $\deg_{G^*}(u, v) \geq np^{(r-1)/2}$, any graph J that satisfies (31) for some $ab \in E(J)$ must be either K_r or a double star with r vertices. We first show how this fact implies the assertion of the claim. Assume first that

$$|\text{Emb}(K_r, G^*; ab, uv)| \geq \frac{\gamma' n^r p^{\binom{r}{2}}}{m_{\max}} = \frac{\gamma'}{KK'} \cdot \frac{n^{r-2} p^{\binom{r-1}{2}}}{\log(1/p)}.$$

Since $|\text{Emb}(K_r, G^*; ab, uv)| \leq 2 \deg_{G^*}(u, v)^{r-2}$, we conclude that

$$\deg_{G^*}(u, v) \geq \left(\frac{\gamma'}{2KK'} \right)^{1/(r-2)} \cdot \frac{np^{(r-1)/2}}{(\log(1/p))^{1/(r-2)}},$$

as claimed. Next, assume that, for some i and j with $i+j = r-2$,

$$|\text{Emb}(S_{i,j}, G^*; ab, uv)| \geq \frac{\gamma' n^r p^{i+j+1}}{m_{\max}} = \frac{\gamma'}{KK'} \cdot \frac{n^{r-2}}{\log(1/p)}.$$

If ab is a centre edge of $S_{i,j}$, then $|\text{Emb}(S_{i,j}, G^*; ab, uv)| \leq (\deg_{G^*} u + \deg_{G^*} v)^{i+j} \leq (\deg_{G^*} u + \deg_{G^*} v) \cdot (2n)^{r-3}$. Otherwise, we have $i, j > 0$ and so $|\text{Emb}(S_{i,j}, G^*; ab, uv)| \leq (\deg_{G^*} u + \deg_{G^*} v)^{\min\{i,j\}} \cdot n^{\max\{i,j\}} \leq (\deg_{G^*} u + \deg_{G^*} v) \cdot (2n)^{r-3}$ as well. Thus, in both cases,

$$\deg_{G^*} u + \deg_{G^*} v \geq \frac{\gamma'}{2^{r-3} KK'} \cdot \frac{n}{\log(1/p)},$$

which completes the proof of the claim.

It remains to prove our assertion. We first consider the special case $r = 3$. The only nonempty subgraph of K_3 with no isolated vertices that is not isomorphic to a double star with three vertices is K_2 . However, as $p \leq C^{-2/(r-1)} = C^{-1}$, we have

$$\frac{|\text{Emb}(K_2, G^*; ab, uv)|}{n^2 p} = \frac{2}{n^2 p} \leq \frac{2 \log C}{C n^2 p^2 \log(1/p)} < \frac{\gamma'}{m_{\max}}$$

whenever C is sufficiently large, which contradicts (31). We henceforth assume that $r \geq 4$.

By way of contradiction, suppose that J is neither K_r nor a double star with r vertices and that $\deg_{G^*}(u, v) < np^{(r-1)/2}$. Let ab be an arbitrary edge of J , let J_{ab} be the subgraph of J induced by $V(J) \setminus \{a, b\}$, and let α_{ab}^* be the fractional independence number of J_{ab} . By Lemma 5.1, there is a partition of $V(J_{ab})$ into V_1 and V_2 such that

$$(P1) \quad |V_1|/2 + |V_2| = \alpha_{ab}^*,$$

(P2) V_1 can be covered by a collection of vertex-disjoint edges and cycles of J_{ab} .

Among all partitions satisfying (P1) and (P2), choose one that maximises the number of common neighbours of a and b in V_2 , that is, the cardinality of the set

$$X = \{c \in V_2 : ac, bc \in E(J)\}.$$

Finally, let $v_1 = |V_1|$, $v_2 = |V_2|$, and $x = |X|$.

We now observe that for every partition satisfying (P1) and (P2), we have

$$|\text{Emb}(J, G^*; ab, uv)| \leq 2 \cdot \deg_{G^*}(u, v)^x \cdot (2e_{G^*})^{v_1/2} \cdot \min\{2e_{G^*}, n\}^{v_2-x}. \quad (32)$$

To see this, note first that there are two embeddings of ab onto uv . Since X lies in the common neighbourhood of a and b , each such embedding can be extended to an embedding of $J[\{a, b\} \cup X]$ in at most $\deg_{G^*}(u, v)^x$ ways. Next, since V_1 can be covered by cycles and edges of J (by property (P2)), Lemma 5.5 implies that every embedding of $J[\{a, b\} \cup X]$ can be extended in at most $(2e_{G^*})^{v_1/2}$ ways to an embedding of $J[\{a, b\} \cup X \cup V_1]$. Finally, since J contains no isolated vertices, any embedding of $J[\{a, b\} \cup X \cup V_1]$ can be extended in at most $\min\{2e_{G^*}, n\}^{v_2-x}$ ways to an embedding of J .

Combining (31), (32), the inequality $e_{G^*} \leq m_{\max}$, which follows from (C2), and the assumption $\deg_{G^*}(u, v) < np^{(r-1)/2}$, we deduce that

$$\begin{aligned} \gamma' n^{v_J} p^{e_J} &\leq \left(np^{(r-1)/2}\right)^x \cdot (2m_{\max})^{v_1/2+1} \cdot \min\{2m_{\max}, n\}^{v_2-x} \\ &= \frac{1}{(KK' \log(1/p))^{x/2}} \cdot (2m_{\max})^{v_1/2+x/2+1} \cdot \min\{2m_{\max}, n\}^{v_2-x}. \end{aligned}$$

On the other hand,

$$n^{v_J} p^{e_J} = (n^2 p^{r-1})^{e_J/(r-1)} \cdot n^{v_J - 2e_J/(r-1)} = \left(\frac{m_{\max}}{KK' \log(1/p)}\right)^{e_J/(r-1)} \cdot n^{v_J - 2e_J/(r-1)}.$$

Therefore, for some constant $K'' = K''(K, K', r)$, we must have

$$(m_{\max})^{v_1/2+x/2-e_J/(r-1)+1} \cdot n^{2e_J/(r-1)-v_J} \cdot \min\{m_{\max}, n\}^{v_2-x} \geq (K'' \log(1/p))^{-K''}. \quad (33)$$

In order to reach the desired contradiction, it suffices to prove that there is some positive $\sigma = \sigma(r)$ such that the left-hand side of (33) is bounded from above by $\max\{n^{-1}, KK' \cdot p^{r-1} \log(1/p)\}^\sigma$. We first observe that such an upper bound is implied by the following inequality:

$$\max\{v_1/2 + v_2 - x/2 - e_J/(r-1) + 1, 0\} < v_J - 2e_J/(r-1). \quad (34)$$

Indeed, if $m_{\max} \leq n$, then the left-hand side of (33) is upper bounded by

$$\begin{aligned} (m_{\max})^{v_1/2+v_2-x/2-e_J/(r-1)+1} \cdot n^{2e_J/(r-1)-v_J} \\ \leq \max\{n^{2e_J/(r-1)-v_J}, n^{v_1/2+v_2-x/2+1-(v_J-e_J/(r-1))}\}. \end{aligned}$$

On the other hand, if $n < m_{\max} = KK' \cdot n^2 p^{r-1} \log(1/p)$, then the left-hand side of (33) becomes

$$\begin{aligned} (KK' \cdot n^2 p^{r-1} \log(1/p))^{v_1/2+x/2-e_J/(r-1)+1} \cdot n^{2e_J/(r-1)-v_J+v_2-x} \\ = (KK' \cdot p^{r-1} \log(1/p))^{v_J-e_J/(r-1)-(v_1/2+v_2-x/2+1)}, \end{aligned}$$

using $v_1 + v_2 = v_{J_{ab}} = v_J - 2$. Note that (34) guarantees that, in both cases, the left-hand side of (33) is at most $\max\{n^{-1}, KK' \cdot p^{r-1} \log(1/p)\}^\sigma$, for a suitable positive $\sigma = \sigma(r)$.

In order to complete the proof of Claim 6.7, we now prove inequality (34). Recall that $V(J_{ab}) = V_1 \cup V_2$ is a partition that satisfies (P1) and (P2) that maximises the cardinality of the set $X = \{c \in V_2 : ac, bc \in E(J)\}$. Since, by the definition of X , each vertex in $V_2 \setminus X$ has at most one neighbour in $\{a, b\}$,

$$e_J \leq e_{J_{ab}} + 2v_1 + v_2 + x + 1 = e_{J_{ab}} + 3v_J - 2\alpha_{ab}^* + x - 5 \quad (35)$$

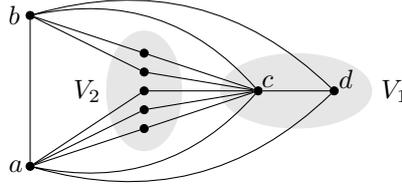


FIGURE 6. Illustration for the case $J_{ab} = K_{1,r-3}$ and $r > 4$. Any partition $V(J_{ab}) = V'_1 \cup V'_2$ obtained by exchanging d with a vertex from V_2 satisfies (P1) and (P2) but violates (D3).

and equality holds only if every vertex in V_1 is adjacent to both a and b and every vertex in $V_2 \setminus X$ is adjacent to exactly one of a and b . Moreover, Lemma 5.2 and Remark 5.3 give

$$e_{J_{ab}} \leq (v_{J_{ab}} - 1)(v_{J_{ab}} - \alpha_{ab}^*) = (v_J - 3)(v_J - \alpha_{ab}^* - 2), \quad (36)$$

where equality holds only if J_{ab} is complete, empty, or isomorphic to K_{1,v_J-3} . Putting (35) and (36) together yields the inequality

$$e_J \leq (v_J - 1)(v_J - \alpha_{ab}^* - 1) + x. \quad (37)$$

Moreover, inequality (37) is strict unless both (35) and (36) hold with equality. Rearranging (37) gives the inequality

$$\frac{e_J}{r-1} \leq \frac{v_J - 1}{r-1} \cdot (v_J - \alpha_{ab}^* - 1) + \frac{x}{r-1} \leq v_J - \alpha_{ab}^* - 1 + \frac{x}{2}. \quad (38)$$

Since $\alpha_{ab}^* \leq v_{J_{ab}} = v_J - 2$ and $r \geq 4$, the second inequality in (38) is strict unless $v_J = r$ and $x = 0$. Consequently, the left-hand and the right-hand sides of (38) can be equal only if the following conditions are satisfied for every partition $V(J_{ab}) = V_1 \cup V_2$ with properties (P1) and (P2):

- (D1) J_{ab} is either K_{r-2} , E_{r-2} (the empty graph with $r-2$ vertices), or $K_{1,r-3}$;
- (D2) every vertex in V_1 is adjacent to both a and b ;
- (D3) every vertex of V_2 is adjacent to exactly one of a and b .

We now show that our assumptions preclude (D1)–(D3) holding simultaneously. Indeed, note first that, if $J_{ab} = K_{r-2}$ (which also includes the case $J_{ab} = K_{1,r-3}$ and $r = 4$), then $\alpha_{ab}^* = v_{J_{ab}}/2$. Since $v_1/2 + v_2 = \alpha_{ab}^*$ and $v_1 + v_2 = v_{J_{ab}}$, this implies $V_1 = V(J_{ab})$. Then (D2) shows that $J = K_r$, a contradiction. Second, if $J_{ab} = E_{r-2}$, then $\alpha_{ab}^* = r-2$. Since $v_1/2 + v_2 = \alpha_{ab}^*$ and $v_1 + v_2 = r-2$, this implies $V_2 = V(J_{ab})$ and it follows from (D3) that J is a double star whose centre edge is ab , another contradiction. Finally, suppose that $J_{ab} = K_{1,r-3}$ and $r > 4$. Since $\alpha_{ab}^* = v_1/2 + v_2 = r-3$ and $v_1 + v_2 = r-2$, we see that $v_1 = 2$ and $v_2 = r-4$. Property (P2) implies that $V_1 = \{c, d\}$, where c is the vertex of degree $r-3$ in J_{ab} and d is one of its neighbours. Since $d \in V_1$, it must be adjacent to both a and b (see Figure 6 for an illustration). Let e be an arbitrary vertex of V_2 ; there is at least one such vertex as $v_2 = r-4 \geq 1$. The partition of $V(J_{ab})$ into $V'_1 = \{c, e\}$ and $V'_2 = (V_2 \setminus \{e\}) \cup \{d\}$ satisfies both conditions (P1) and (P2) but the set $X' = \{c' \in V'_2 : ac', vc' \in E(J_{ab})\}$ is nonempty (as it contains the vertex d); this contradicts property (D3) for the partition $V(J) = V'_1 \cup V'_2$.

To summarise, at least one of the inequalities in (38) is strict. Since $J \neq K_r$, not every vertex of J has degree $r-1$ and thus $2e_J < (r-1) \cdot v_J$. We conclude that

$$\max \{ \alpha_{ab}^* - e_J/(r-1) - x/2 + 1, 0 \} < v_J - 2e_J/(r-1).$$

Since $\alpha_{ab}^* = v_1/2 + v_2$, this is exactly (34). \square

6.2. Proof of Proposition 6.5. We begin by showing that

$$\limsup_{n \rightarrow \infty} \frac{\Phi_X(\delta)}{n^2 p^{r-1} \log(1/p)} \leq \varphi_r(\delta, c). \quad (39)$$

For every small $\varepsilon > 0$ and sufficiently large n , we shall construct a graph G with vertex set $\llbracket n \rrbracket$ and at most $(\varphi_r(\delta + \varepsilon, c) + \varepsilon) \cdot n^2 p^{r-1}$ edges that satisfies $\mathbb{E}_G[X] \geq (1 + \delta) \mathbb{E}[X]$. The existence of such a graph and the continuity of φ_r will imply that

$$\limsup_{n \rightarrow \infty} \frac{\Phi_X(\delta)}{n^2 p^{r-1} \log(1/p)} \leq \lim_{\varepsilon \rightarrow 0} \varphi_r(\delta + \varepsilon, c) + \varepsilon = \varphi_r(\delta, c),$$

as required.

Let $x \in \bar{X}(\delta + \varepsilon, c)$ and note that it follows from Lemma 6.1 that $x = 0$ if $c = 0$. Set

$$\ell_1 = ((\delta + \varepsilon)(1 - x))^{1/r} n p^{(r-1)/2} \quad \text{and} \quad \ell_2 = x(\delta + \varepsilon) n p^{r-1} / r.$$

and fix an arbitrary partition $\llbracket n \rrbracket = \{u\} \cup U_1 \cup U_2 \cup U_3$, where $|U_1| = \lfloor \ell_1 \rfloor$ and $|U_2| = \lfloor \ell_2 \rfloor$. Let G be the union of the clique on U_1 , the complete bipartite graph between U_2 and U_3 , and an arbitrary star, centred at u , with $\lfloor \{\ell_2\}^{1/(r-1)} |U_3| \rfloor$ edges whose non- u endpoints are in U_3 . We have

$$\begin{aligned} e_G &\leq \binom{\ell_1}{2} + \lfloor \ell_2 \rfloor \cdot |U_3| + \{\ell_2\}^{1/(r-1)} \cdot |U_3| \\ &\leq \frac{((\delta + \varepsilon)(1 - x))^{2/r} n^2 p^{r-1}}{2} + (\lfloor \ell_2 \rfloor + \{\ell_2\}^{1/(r-1)}) \cdot n \\ &\leq \left(\frac{((\delta + \varepsilon)(1 - x))^{2/r}}{2} + \frac{\lfloor \ell_2 \rfloor + \{\ell_2\}^{1/(r-1)}}{n p^{r-1}} \right) \cdot n^2 p^{r-1} \\ &= \psi_r(\delta + \varepsilon, n p^{r-1}, x) \cdot n^2 p^{r-1}, \end{aligned}$$

and our choice of x ensures that $\psi_r(\delta + \varepsilon, n p^{r-1}, x) \leq \varphi_r(\delta + \varepsilon, c) + \varepsilon$ for all large enough n .

It remains to show that $\mathbb{E}_G[X] \geq (1 + \delta) \mathbb{E}[X]$. To this end, observe first that:

- (i) The complete graph $G[U_1]$ contains $\binom{\lfloor \ell_1 \rfloor}{r}$ copies of K_r .
- (ii) The complete bipartite graph $G[U_2, U_3]$ contains $\lfloor \ell_2 \rfloor \cdot \binom{|U_3|}{r-1}$ copies of $K_{1, r-1}$ whose centre vertex lies in U_2 .
- (iii) The star $G[u, U_3]$ contains $\binom{\lfloor \{\ell_2\}^{1/(r-1)} |U_3| \rfloor}{r-1}$ copies of $K_{1, r-1}$.

In particular,

$$\begin{aligned} \mathbb{E}_G[X] - \mathbb{E}[X] &= \sum_{\emptyset \neq J \subseteq K_r} N(J, G) \cdot \binom{n - v_J}{r - v_J} \cdot p^{\binom{r}{2}} (p^{-e_J} - 1) \\ &\geq (1 - p) \cdot \left[\binom{\lfloor \ell_1 \rfloor}{r} + \left[\lfloor \ell_2 \rfloor \cdot \binom{|U_3|}{r-1} + \binom{\lfloor \{\ell_2\}^{1/(r-1)} |U_3| \rfloor}{r-1} \right] \cdot p^{\binom{r}{2} - (r-1)} \right]. \end{aligned}$$

We now estimate the right-hand side of the above inequality. Since $n p^{(r-1)/2} \rightarrow \infty$, we have

$$\binom{\lfloor \ell_1 \rfloor}{r} \geq \frac{(\ell_1 - r)^r}{r!} \geq ((\delta + \varepsilon)(1 - x) - \varepsilon/2) \cdot \frac{n^r p^{\binom{r}{2}}}{r!},$$

provided that n is sufficiently large. In the case $c = 0$, this is already sufficient, as $x = 0$ and

$$\mathbb{E}_G[X] - \mathbb{E}[X] \geq (1 - p) \cdot \binom{\lfloor \ell_1 \rfloor}{r} \geq (1 - p) \cdot (\delta + \varepsilon/2) \cdot \frac{n^r p^{\binom{r}{2}}}{r!} \geq \delta \mathbb{E}[X].$$

We may therefore assume that $c \in (0, \infty]$. Since $p \rightarrow 0$, and thus $|U_3|/n \rightarrow 1$, we have

$$\begin{aligned} [\ell_2] \cdot \binom{|U_3|}{r-1} + \binom{[\{\ell_2\}^{1/(r-1)}|U_3|]}{r-1} &\geq [\ell_2] \cdot \frac{(1-\varepsilon^2) \cdot n^{r-1}}{(r-1)!} + \frac{(\{\ell_2\}^{1/(r-1)}|U_3| - r)^{r-1}}{(r-1)!} \\ &\geq \frac{(1-\varepsilon^2) \cdot \ell_2 \cdot n^{r-1}}{(r-1)!} - \frac{\varepsilon^2 n^{r-1}}{(r-1)!} \\ &= \left((1-\varepsilon^2)x(\delta + \varepsilon) - \frac{\varepsilon^2 r}{np^{r-1}} \right) \cdot \frac{n^r p^{r-1}}{r!}. \end{aligned}$$

Consequently,

$$\begin{aligned} \mathbb{E}_G[X] - \mathbb{E}[X] &\geq (1-p) \cdot \left((\delta + \varepsilon)(1-x) - \varepsilon/2 + (1-\varepsilon^2)x(\delta + \varepsilon) - \frac{\varepsilon^2 r}{np^{r-1}} \right) \cdot \frac{n^r p^{\binom{r}{2}}}{r!} \\ &\geq (1-p) \cdot \left(\delta + \varepsilon/2 - \varepsilon^2(\delta + \varepsilon) - \frac{\varepsilon^2 r}{np^{r-1}} \right) \cdot \mathbb{E}[X] \geq \delta \mathbb{E}[X], \end{aligned}$$

where the last inequality holds for all sufficiently small ε since $np^{r-1} \rightarrow c > 0$. This completes the proof of (39).

It remains to prove the matching lower bound

$$\liminf_{n \rightarrow \infty} \frac{\Phi_X(\delta)}{n^2 p^{r-1} \log(1/p)} \geq \varphi_r(\delta, c). \quad (40)$$

Fix $\varepsilon > 0$ small enough and suppose that n is sufficiently large. By the continuity of φ_r , it is enough to show that any graph G on n vertices satisfying

$$\mathbb{E}_G[X] \geq (1 + \delta) \mathbb{E}[X]$$

has at least $(1 - \varepsilon) \cdot \varphi_r(\delta - \varepsilon, c) \cdot n^2 p^{r-1}$ edges. By way of contradiction, assume that $e_G < (1 - \varepsilon) \cdot \varphi_r(\delta - \varepsilon, c) \cdot n^2 p^{r-1}$. Note that, for all large enough n ,

$$(\delta - \varepsilon/4) \cdot \frac{n^r p^{\binom{r}{2}}}{r!} \leq \delta \mathbb{E}[X] \leq \mathbb{E}_G[X] - \mathbb{E}[X] \leq \sum_{\emptyset \neq J \subseteq K_r} N(J, G) \cdot n^{r-v_J} \cdot p^{\binom{r}{2} - e_J},$$

where the sum ranges over the nonempty subgraphs J of K_r without isolated vertices, so

$$\sum_{\emptyset \neq J \subseteq K_r} \frac{N(J, G)}{n^{v_J} p^{e_J}} \geq \frac{\delta - \varepsilon/3}{r!}. \quad (41)$$

Using our assumed upper bound on e_G , Theorem 5.4 implies that

$$\begin{aligned} \frac{N(J, G)}{\delta n^{v_J} p^{e_J}} &\leq \frac{(2e_G)^{v_J - \alpha_J^*} \cdot \min\{2e_G, n\}^{2\alpha_J^* - v_J}}{n^{v_J} p^{e_J}} \\ &\leq C \cdot \frac{\min\{(n^2 p^{r-1})^{\alpha_J^*}, n^{v_J} p^{(r-1)(v_J - \alpha_J^*)}\}}{n^{v_J} p^{e_J}} \end{aligned}$$

for a suitable constant C . If the minimum is achieved by the first term, then $np^{r-1} \leq 1$, and thus

$$\frac{(n^2 p^{r-1})^{\alpha_J^*}}{n^{v_J} p^{e_J}} = n^{\alpha_J^* - v_J + e_J/(r-1)} \cdot (np^{r-1})^{\alpha_J^* - e_J/(r-1)} \leq n^{\alpha_J^* - v_J + e_J/(r-1)},$$

as $\alpha_J^* \geq v_J/2 \geq e_J/(r-1)$ for every graph J with maximum degree at most $r-1$. A straightforward algebraic manipulation in the case where the minimum is achieved by the second term then shows that, in both cases,

$$\frac{N(J, G)}{n^{v_J} p^{e_J}} \leq C \cdot \min\{n^{-1}, p^{r-1}\}^{v_J - \alpha_J^* - e_J/(r-1)}.$$

If $J \subseteq K_r$ is not equal to either K_r or $K_{1,r-1}$, then Lemma 5.2 and the remark that follows it imply that $v_J - \alpha_J^* - e_J/(r-1) > 0$, and then the right-hand side goes to zero as $n \rightarrow \infty$. It thus follows from (41) that

$$\frac{N(K_r, G)}{n^r p^{\binom{r}{2}}} + \frac{N(K_{1,r-1}, G)}{n^r p^{r-1}} \geq \frac{\delta - \varepsilon/2}{r!},$$

or, equivalently, since $|\text{Aut}(K_r)| = r!$ and $|\text{Aut}(K_{1,r-1})| = (r-1)!$,

$$|\text{Emb}(K_r, G)| + r \cdot |\text{Emb}(K_{1,r-1}, G)| \cdot p^{\binom{r}{2}-r+1} \geq (\delta - \varepsilon/2) \cdot n^r p^{\binom{r}{2}}. \quad (42)$$

Recalling our assumption $e_G < (1 - \varepsilon) \cdot \varphi_r(\delta - \varepsilon, c) \cdot n^2 p^{r-1} \ll n^2$, it follows from Lemma 5.15 that there is a partition $V(G) = U \cup V$ such that $|U| \leq \varepsilon^2 n$,

$$|\text{Emb}(K_r, G[V])| \geq |\text{Emb}(K_r, G)| - \varepsilon^2 e_G^{r/2} \geq |\text{Emb}(K_r, G)| - \varepsilon n^r p^{\binom{r}{2}}/4,$$

and, recalling that $\text{Emb}_U(K_{1,r-1}, G)$ denotes the set of embeddings of $K_{1,r-1}$ into G that map the centre vertex of $K_{1,r-1}$ to a vertex of U ,

$$\begin{aligned} r \cdot |\text{Emb}_U(K_{1,r-1}, G[U, V])| \cdot p^{\binom{r}{2}-r+1} &\geq r \cdot |\text{Emb}(K_{1,r-1}, G)| \cdot p^{\binom{r}{2}-r+1} - \varepsilon^2 e_G n^{r-2} \cdot p^{\binom{r}{2}-r+1} \\ &\geq r \cdot |\text{Emb}(K_{1,r-1}, G)| \cdot p^{\binom{r}{2}-r+1} - \varepsilon n^r p^{\binom{r}{2}}/4, \end{aligned}$$

where the stated inequalities are valid if ε is sufficiently small. From this and (42), we obtain

$$|\text{Emb}(K_r, G[V])| + r \cdot |\text{Emb}_U(K_{1,r-1}, G[U, V])| \cdot p^{\binom{r}{2}-r+1} \geq (\delta - \varepsilon) \cdot n^r p^{\binom{r}{2}};$$

consequently, there exists an $x \in [0, 1]$ such that

$$|\text{Emb}(K_r, G[V])| \geq (1 - x) \cdot (\delta - \varepsilon) \cdot n^r p^{\binom{r}{2}}$$

and

$$r \cdot |\text{Emb}_U(K_{1,r-1}, G[U, V])| \cdot p^{\binom{r}{2}-r+1} \geq x \cdot (\delta - \varepsilon) \cdot n^r p^{\binom{r}{2}}.$$

By Theorem 5.4 and Lemma 5.14, we thus obtain the bounds

$$\begin{cases} (2e_{G[V]})^{r/2} \geq (1 - x) \cdot (\delta - \varepsilon) \cdot n^r p^{\binom{r}{2}} \\ \left(\lfloor e_{G[U, V]} / |V| \rfloor + \{e_{G[U, V]} / |V|\}^{r-1} \right) \cdot n^{r-1} \geq x \cdot (\delta - \varepsilon) \cdot n^r p^{r-1} / r, \end{cases}$$

and solving for $e_{G[V]}$ and $e_{G[U, V]}$, we get

$$\begin{cases} e_{G[V]} \geq ((1 - x) \cdot (\delta - \varepsilon))^{2/r} \cdot \frac{n^2 p^{r-1}}{2} \\ e_{G[U, V]} \geq |V| \cdot \left(\lfloor x \cdot (\delta - \varepsilon) n p^{r-1} / r \rfloor + \{x \cdot (\delta - \varepsilon) n p^{r-1} / r\}^{1/(r-1)} \right). \end{cases}$$

Finally, since $|V| = n - |U| \geq (1 - \varepsilon^2)n$, the definition of ψ_r shows that

$$e_G \geq e_{G[V]} + e_{G[U, V]} \geq (1 - \varepsilon^2) \cdot \psi_r(\delta - \varepsilon, n p^{r-1}, x) \cdot n^2 p^{r-1}.$$

As $\psi_r(\delta - \varepsilon, n p^{r-1}, x) \geq \varphi_r(\delta - \varepsilon, n p^{r-1}) \rightarrow \varphi_r(\delta - \varepsilon, c)$, this contradicts our assumption that $e_G < (1 - \varepsilon) \cdot \varphi_r(\delta - \varepsilon, c) \cdot n^2 p^{r-1}$, provided that ε is sufficiently small and n is large enough.

6.3. Proof of Proposition 6.6. Fix $\varepsilon, \delta > 0$ and $c \in [0, \infty]$ and assume that $n p^{r-1} \rightarrow c$. We fix three additional positive constants η, η' , and γ , where γ is sufficiently small given the parameters of the proposition, η' is sufficiently small given γ , and finally η is sufficiently small given both η' and γ .

If Near-min(η) occurs, then $G_{n,p}$ contains a subgraph G such that $e_G \log(1/p) \leq (1 + \eta)\Phi_X(\delta + \eta)$ and $\mathbb{E}_G[X] \geq (1 + \delta - \eta)\mathbb{E}[X]$. We claim that, if n is sufficiently large, then every such graph admits a partition $V(G) = U \cup V$ such that, for some $x \in \bar{X}(\delta, c)$,

- (i) V contains a subset V' of size at least $(1 - \varepsilon)(\delta(1 - x))^{1/r} n p^{(r-1)/2}$ such that $G[V']$ has minimum degree at least $(1 - \varepsilon)|V'|$.

- (ii) U contains a set $W \subseteq U$ such that at least $\lfloor (1 - \varepsilon)|W| \rfloor$ vertices in W have degree at least $(1 - \varepsilon)n$ and

$$e(W, V) \geq (1 - \varepsilon)n(\lfloor \delta x n p^{r-1}/r \rfloor + \{\delta x n p^{r-1}/r\}^{1/(r-1)}).$$

Note that these properties imply $\text{Clique}_\varepsilon(\delta(1 - x)) \cap \text{Hub}_\varepsilon(\delta x)$.

By repeating the argument found in the proof of the lower bound in Proposition 6.5, we can find a partition $V(G) = U \cup V$ and some $x' \in [0, 1]$ such that

$$\begin{cases} |\text{Emb}(K_r, G[V])| \geq (1 - x') \cdot (\delta - 2\eta) \cdot n^r p^{\binom{r}{2}} \\ r \cdot |\text{Emb}_U(K_{1,r-1}, G[U, V])| \cdot p^{\binom{r}{2}-r+1} \geq x' \cdot (\delta - 2\eta) \cdot n^r p^{\binom{r}{2}} \end{cases} \quad (43)$$

and

$$\begin{cases} e_{G[V]} \geq ((1 - x') \cdot (\delta - 2\eta))^{2/r} \cdot \frac{n^2 p^{r-1}}{2} \\ e_{G[U, V]} \geq (\lfloor x'(\delta - 2\eta)n p^{r-1}/r \rfloor + \{x'(\delta - 2\eta)n p^{r-1}/r\}^{1/(r-1)}) \cdot n. \end{cases} \quad (44)$$

It follows that $e_G \geq e_{G[V]} + e_{G[U, V]} \geq \psi_r(\delta - 2\eta, n p^{r-1}, x') \cdot n^2 p^{r-1}$. Thus, by Proposition 6.5,

$$\psi_r(\delta - 2\eta, n p^{r-1}, x') \leq \frac{e_G}{n^2 p^{r-1}} \leq \frac{(1 + \eta)\Phi_X(\delta + \eta)}{n^2 p^{r-1} \log(1/p)} \leq (1 + 2\eta)\varphi_r(\delta + \eta, c). \quad (45)$$

Our next claim tells us how to choose the constant η . Given a set $S \subseteq \mathbb{R}$, we write $B_{\eta'}(S) = \{x \in \mathbb{R} : \inf_{s \in S} |x - s| < \varepsilon\}$ for the η' -neighbourhood of S .

Claim 6.8. *We may choose $\eta = \eta(\eta') > 0$ such that $x' \in B_{\eta'}(\bar{X}_r(\delta, c))$ whenever n is sufficiently large.*

Proof. Suppose there is no such choice. Then, by invoking (45) with successively smaller values of η , we find that there is a subsequence (η_m, c_m, x_m) of points in $(0, \infty)^2 \times [0, 1]$ such that $\eta_m \rightarrow 0$ and $c_m \rightarrow c$ as $m \rightarrow \infty$, and, for all m , we have $x_m \notin B_{\eta'}(\bar{X}_r(\delta, c))$ and

$$\varphi_r(\delta - 2\eta_m, c_m) \leq \psi_r(\delta - 2\eta_m, c_m, x_m) \leq (1 + 2\eta_m)\varphi_r(\delta + \eta_m, c).$$

Since both the left-hand and the right-hand sides above converge to $\varphi_r(\delta, c)$ as $m \rightarrow \infty$, so must $\psi_r(\delta - 2\eta_m, c_m, x_m)$. Denote by $(\eta_\ell, c_\ell, x_\ell)$ a subsequence on which x_ℓ converges to some $x_\infty \in [0, 1]$. We claim that $x_\infty \in \bar{X}_r(\delta, c)$, which contradicts the fact that $x_\ell \notin B_{\eta'}(\bar{X}_r(\delta, c))$ for all ℓ .

If $c = 0$, then we have $x_\infty = 0 \in \bar{X}_r(\delta, c)$, since otherwise the definition of ψ_r would imply that $\psi_r(\delta - 2\eta_\ell, c_\ell, x_\ell) \rightarrow \infty$. If $c \in (0, \infty)$, then continuity of ψ_r implies that $\varphi_r(\delta, c) = \lim_{\ell \rightarrow \infty} \psi_r(\delta - 2\eta_\ell, c_\ell, x_\ell) = \psi_r(\delta, c, x_\infty)$ and thus $x_\infty \in \bar{X}_r(\delta, c)$. Finally, if $c = \infty$, then $(\delta, x) \mapsto \psi_r(\delta, c_m, x)$ converges uniformly to the continuous function $(\delta, x) \mapsto (\delta x)^{2/r}/2 + \delta x/r$, which implies that

$$\lim_{\ell \rightarrow \infty} \psi_r(\delta - 2\eta_\ell, c_\ell, x_\ell) = \frac{(\delta x_\infty)^{2/r}}{2} + \frac{\delta x_\infty}{r} = \lim_{c' \rightarrow \infty} \psi(\delta, c', x_\infty),$$

so $x_\infty \in \bar{X}_r(\delta, c)$ in this case as well. \square

Suppose now that $x' \in B_{\eta'}(\bar{X}(\delta, c))$. Since the right-hand sides of (43) and (44) are continuous in x' , we may choose η and η' sufficiently small, as a function of γ , so that there is some $x \in \bar{X}_r(\delta, c)$ such that

$$\begin{cases} |\text{Emb}(K_r, G[V])| \geq (1 - x) \cdot (\delta - \gamma) \cdot n^r p^{\binom{r}{2}} \\ r \cdot |\text{Emb}_U(K_{1,r-1}, G[U, V])| \cdot p^{\binom{r}{2}-r+1} \geq x \cdot (\delta - \gamma) \cdot n^r p^{\binom{r}{2}}, \\ e_{G[V]} \geq ((1 - x) \cdot (\delta - \gamma))^{2/r} \cdot \frac{n^2 p^{r-1}}{2} \\ e_{G[U, V]} \geq (\lfloor x(\delta - \gamma)n p^{r-1}/r \rfloor + \{x(\delta - \gamma)n p^{r-1}/r\}^{1/(r-1)}) \cdot n. \end{cases}$$

Since $e_{G[V]} + e_{G[U,V]} \leq e_G \leq (1 + 2\eta)\psi_r(\delta + \eta, c, x') \cdot n^2 p^{r-1}$ and since the lower bounds on $e_{G[V]}$ and $e_{G[U,V]}$ in (44) sum to $\psi_r(\delta - 2\eta, np^{r-1}, x') \cdot n^2 p^{r-1}$, the two lower bounds on $e_{G[V]}$ and $e_{G[U,V]}$ stated above must be nearly tight. More precisely, the continuity of ψ_r implies that we may choose η and η' sufficiently small so that

$$\begin{cases} e_{G[V]} \leq ((1-x) \cdot (\delta + \gamma))^{2/r} \cdot \frac{n^2 p^{r-1}}{2} \\ e_{G[U,V]} \leq (\lfloor x(\delta + \gamma)np^{r-1}/r \rfloor + \{x(\delta + \gamma)np^{r-1}/r\}^{1/(r-1)}) \cdot n. \end{cases}$$

Finally, if we choose γ sufficiently small, then the above statements for $G[V]$ and Theorem 5.11 yield a set $V' \subseteq V$ satisfying the conditions in (i). Similarly, the statements for $G[U, V]$ and Lemma 5.14(ii) yield a subset $W \subseteq U$ satisfying (ii).

7. EXTENSIONS TO REGULAR GRAPHS

Fix a connected and Δ -regular graph H . In this section, we apply Theorem 3.1 to study the upper tail of the random variable $X = X_{n,p}^H$. In this setting, (11) may be rewritten as

$$\Phi_X(\delta) = \min \{e_G \log(1/p) : G \subseteq K_n \text{ and } \mathbb{E}_G[X] \geq (1 + \delta) \mathbb{E}[X]\},$$

where we use the notation $\mathbb{E}_G[X] = \mathbb{E}[X \mid G \subseteq G_{n,p}]$. Our main result in this section is the following:

Proposition 7.1. *For every $\Delta \geq 2$, every connected, nonbipartite, Δ -regular graph H , and all positive real numbers ε and δ , there exists a positive constant C such that the following holds. Suppose that an integer n and $p \in (0, 1)$ satisfy $Cn^{-1}(\log n)^{\Delta v_H^2} \leq p^{\Delta/2} \leq 1/C$. Then $X = X_{n,p}^H$ satisfies*

$$(1 - \varepsilon)\Phi_X(\delta - \varepsilon) \leq -\log \mathbb{P}(X \geq (1 + \delta) \mathbb{E}[X]) \leq (1 + \varepsilon)\Phi_X(\delta + \varepsilon).$$

Additionally, there is a positive constant $\xi = \xi(\Delta, \varepsilon)$ such that the assumption that H is nonbipartite is not necessary when $p^{\Delta/2} \geq n^{-1/2-\xi}$.

Remark 7.2. As was mentioned in the introduction, the lower bound on the density p is suboptimal by a polylogarithmic factor. We require this slightly stronger bound on the density to prove an analogue of Claim 6.7 for Δ -regular graphs H . In this more general setup, the lower bound on the number of common neighbours of the endpoints of core edges may not hold, and will be replaced by a lower bound on the product of degrees of the endpoints of core edges. We will prove this bound via an application of Lemma 5.7; unfortunately, this will incur a polylogarithmic loss, and thus necessitates the suboptimal bound on the density p .

The optimisation problem $\Phi_X(\delta)$ is a discretisation of the variational problem considered by Bhattacharya, Ganguly, Lubetzky, and Zhao [9]. As was the case in the context of arithmetic progressions and cliques, their variational problem was more general (being optimised over a larger set), but is asymptotically equivalent to the one considered in this paper. Their results imply

$$\lim_{n \rightarrow \infty} \frac{\Phi_X(\delta)}{n^2 p^\Delta \log(1/p)} = \begin{cases} \delta^{2/v_H} / 2 & \text{if } np^\Delta \rightarrow 0, \\ \min \{ \delta^{2/v_H} / 2, \theta \} & \text{if } np^\Delta \rightarrow \infty, \end{cases}$$

where θ is the unique positive solution to $P_H(\theta) = 1 + \delta$ and P_H is the independence polynomial of H . Thus, Proposition 7.1 implies Theorem 1.5. The proof of Proposition 7.1 does not require these precise estimates, but only that $\Phi_X(\delta)$ is of order $n^2 p^\Delta \log(1/p)$, see Lemma 7.3 below. We include a short proof of this weaker statement for the sake of completeness, and to emphasize it may be proved combinatorially; the original proof given in [9] is conceptually equivalent but more analytically-flavoured.

Lemma 7.3. *For every $\Delta \geq 2$, every connected, Δ -regular graph H , and every positive real number δ , there exists a positive constant C such that the following holds. Assume $n \in \mathbb{N}$ and $p \in (0, 1)$ are such that $Cn^{-1} \leq p^{\Delta/2} \leq 1/C$. Then $X = X_{n,p}^H$ satisfies*

$$1/C \leq \frac{\Phi_X(\delta)}{n^2 p^\Delta \log(1/p)} \leq C.$$

Proof. The upper bound follows by noting that, if C is large enough, then a clique with $\lceil (1 + 2\delta)^{1/v_H} n p^{\Delta/2} \rceil$ vertices contains at least $(1 + \delta) \mathbb{E}[X]$ copies of H and has fewer than $Cn^2 p^\Delta$ edges. For the lower bound, suppose that G is a graph with $\mathbb{E}_G[X] \geq (1 + \delta) \mathbb{E}[X]$ with fewer than $C^{-1} n^2 p^\Delta$ edges. Then

$$\frac{\delta}{2} \cdot \frac{n^{v_H} p^{e_H}}{|\text{Aut}(H)|} \leq \delta \mathbb{E}[X] \leq \mathbb{E}_G[X] - \mathbb{E}[X] \leq \sum_{\emptyset \neq J \subseteq H} N(J, G) \cdot n^{v_H - v_J} p^{e_H - e_J},$$

where the sum ranges over all nonempty subgraphs J of H without isolated vertices. This implies that there is a nonempty subgraph $J \subseteq H$ without isolated vertices and a positive constant $\gamma = \gamma(H, \delta)$ such that

$$|\text{Emb}(J, G)| \geq \gamma \cdot n^{v_J} p^{e_J}. \quad (46)$$

Theorem 5.4 implies that

$$|\text{Emb}(J, G)| \leq (2e_G)^{v_J - \alpha_J^*} \cdot \min\{2e_G, n\}^{2\alpha_J^* - v_J}$$

and Lemma 5.2 yields $\alpha_J^* \leq v_J - e_J/\Delta$. Therefore, if $2e_G \leq n$, then $|\text{Emb}(J, G)|$ is bounded from above by

$$(2e_G)^{\alpha_J^*} \leq (2e_G)^{\alpha_J^* - v_J + 2e_J/\Delta} \cdot n^{v_J - 2e_J/\Delta} \leq \left(\frac{2n^2 p^\Delta}{C}\right)^{e_J/\Delta} n^{v_J - 2e_J/\Delta} = \frac{n^{v_J} p^{e_J}}{(C/2)^{e_J/\Delta}},$$

where the first inequality holds since $2e_J/\Delta \leq v_J$, as the maximum degree of J is at most Δ . If $n < 2e_G$, then $|\text{Emb}(J, G)|$ is bounded from above by

$$\left(\frac{2n^2 p^\Delta}{C}\right)^{v_J - \alpha_J^*} \cdot n^{2\alpha_J^* - v_J} = \frac{n^{v_J} p^{e_J}}{(C/2)^{v_J - \alpha_J^*}} \cdot p^{\Delta(v_J - \alpha_J^*) - e_J} \leq \frac{n^{v_J} p^{e_J}}{(C/2)^{v_J - \alpha_J^*}}.$$

In both cases, the obtained upper bound on $|\text{Emb}(J, G)|$ contradicts (46) whenever C is large; indeed e_J/Δ and $v_J - \alpha_J^*$ are both positive, as J is nonempty. \square

7.1. Proof of Proposition 7.1. Fix $\Delta \geq 2$, a nonempty, connected, Δ -regular graph H , and positive reals ε and δ . Without loss of generality, we may assume that $\varepsilon \leq \min\{1/3, \delta/2\}$. Let $X = X_{n,p}^H$ and assume $Cn^{-1}(\log n)^{\Delta v_H^2} \leq p^{\Delta/2} \leq 1/C$. Note that the case $n < v_H$ is trivial as then X is identically zero. We may therefore assume that $n \geq v_H \geq \Delta + 1 \geq 3$, which, in turn, implies that $n \geq C$.

Set $N = \binom{n}{2}$ and let $Y = (Y_1, \dots, Y_N)$ be the sequence of indicator random variables of the events that $e \in E(G_{n,p})$, where e ranges over $\binom{[n]}{2}$ in some arbitrary order. Then Y is a vector of independent $\text{Ber}(p)$ random variables and X is a nonzero polynomial with nonnegative coefficients and degree at most e_H in the coordinates of Y . Let $K = K(e_H, \varepsilon, \delta)$ be the constant whose existence is asserted by Theorem 3.1. To prove the proposition, it suffices to verify the assumptions of the theorem.

It follows from Lemma 7.3 and our assumptions on p that $p \leq 1 - \varepsilon$ and $\Phi_X(\delta - \varepsilon) \geq K \log(1/p)$ for a large enough choice of C . It thus only remains to bound the number of cores of a given size. To this end, let \mathcal{I}_m^* be the set of cores with m edges, that is, subgraphs $G^* \subseteq K_n$ such that

- (C1) $\mathbb{E}_{G^*}[X] \geq (1 + \delta - \varepsilon) \mathbb{E}[X]$,
- (C2) $e_{G^*} = m \leq K \cdot \Phi_X(\delta + \varepsilon)$, and
- (C3) $\min_{e \in E(G^*)} (\mathbb{E}_{G^*}[X] - \mathbb{E}_{G^* \setminus e}[X]) \geq \mathbb{E}[X]/(K \cdot \Phi_X(\delta + \varepsilon))$.

Proposition 7.1 will follow once we prove

$$|\mathcal{I}_m^*| \leq (1/p)^{\varepsilon m/2} \quad \text{for all } m. \quad (47)$$

Observe that due to (C1), (C2), and the definition of Φ_X , it suffices to verify (47) for integers m such that $m_{\min} \leq m \leq m_{\max}$ with

$$\begin{aligned} m_{\min} &:= \Phi_X(\delta - \varepsilon) / \log(1/p) \geq n^2 p^\Delta / K' \quad \text{and} \\ m_{\max} &:= K \cdot \Phi_X(\delta + \varepsilon) \leq K' \cdot n^2 p^\Delta \log(1/p), \end{aligned}$$

where the stated inequalities follow, for a suitable constant $K' = K'(H, \varepsilon, \delta)$, from Lemma 7.3, our bounds on p , and the assumption that C is sufficiently large.

The first step towards establishing (47) is to understand the combinatorial meaning of (C3). Suppose that $m_{\min} \leq m \leq m_{\max}$ and $G^* \in \mathcal{I}_m^*$. Recall that $N(J, G^*; e)$ denotes the number of copies of J in G^* that contain the edge e . Note that (C3) implies that, for every $e \in E(G^*)$,

$$\frac{\mathbb{E}[X]}{m_{\max}} \leq \mathbb{E}_{G^*}[X] - \mathbb{E}_{G^* \setminus e}[X] \leq \sum_{\emptyset \neq J \subseteq H} N(J, G^*; e) \cdot n^{v_H - v_J} p^{e_H - e_J}, \quad (48)$$

where the sum ranges over the nonempty subgraphs J of H without isolated vertices. Since $n \geq C$ for a large enough constant C , we can bound $\mathbb{E}[X] \geq \binom{n}{v_H} p^{e_H} \geq \frac{n^{v_H} p^{e_H}}{2^{v_H}!}$ and, consequently, (48) implies that there is a positive constant $\gamma = \gamma(H)$ such that

$$\sum_{\emptyset \neq J \subseteq H} \frac{N(J, G^*; e)}{n^{v_J} p^{e_J}} \geq \frac{2\gamma}{m_{\max}} \quad \text{for every } e \in E(G^*). \quad (49)$$

Definition 7.4. Let \mathcal{Q}_H denote the family of all nonempty subgraphs $J \subseteq H$ without isolated vertices satisfying

- (Q1) $J = H$ or
- (Q2) J admits a bipartition $V(J) = A \cup B$ such that $\deg_J a = \Delta$ for all $a \in A$.

Our first claim is that, for the vast majority of $e \in E(G^*)$, the sum on the left-hand side of (49) is dominated by subgraphs $J \in \mathcal{Q}_H$. Let G_{exc}^* comprise all edges e of G^* such that

$$\sum_{J \in \mathcal{Q}_H} \frac{N(J, G^*; e)}{n^{v_J} p^{e_J}} < \frac{\gamma}{m_{\max}}.$$

Claim 7.5. *There is a positive constant $\sigma = \sigma(H)$ such that $e_{G_{\text{exc}}^*} \leq p^\sigma \cdot m_{\min}$.*

Proof. Let m_{exc} denote the number of edges of G_{exc}^* . The definition of G_{exc}^* and (49) imply that

$$\sum_{e \in E(G_{\text{exc}}^*)} \sum_{\substack{\emptyset \neq J \subseteq H \\ J \notin \mathcal{Q}_H}} \frac{N(J, G^*; e)}{n^{v_J} p^{e_J}} \geq \frac{\gamma \cdot m_{\text{exc}}}{m_{\max}} \geq \frac{\gamma \cdot m_{\text{exc}}}{m_{\min} \cdot (K')^2 \cdot \log(1/p)}.$$

On the other hand, since $\sum_{e \in E(G^*)} N(J, G^*; e) \leq e_J \cdot |\text{Emb}(J, G^*)|$ for every graph J , there must exist a nonempty subgraph $J \subseteq H$ without isolated vertices such that $J \notin \mathcal{Q}_H$ and

$$\frac{m_{\text{exc}}}{m_{\min}} \leq \frac{e_J \cdot (K')^2 \cdot \log(1/p) \cdot 2^{e_H + v_H}}{\gamma} \cdot \frac{|\text{Emb}(J, G^*)|}{n^{v_J} p^{e_J}}. \quad (50)$$

We now show that the right-hand side of (50) is at most p^σ , for some positive constant $\sigma = \sigma(H)$. To this end, recall that Theorem 5.4 states that

$$|\text{Emb}(J, G^*)| \leq (2m)^{v_J - \alpha_J^*} \cdot \min\{2m, n\}^{2\alpha_J^* - v_J}. \quad (51)$$

Since J is a proper subgraph of a connected, Δ -regular graph, it must contain a vertex of degree smaller than Δ and hence $v_J > 2e_J/\Delta$. Moreover, since $J \notin \mathcal{Q}_H$, Lemma 5.2 implies that

$\alpha_J^* < v_J - e_J/\Delta$. Therefore, if $2m \leq n$, then there is a positive $\sigma = \sigma(H)$ such that the right-hand side of (51) can be bounded from above as follows:

$$\begin{aligned} (2m)^{\alpha_J^*} &\leq (2m)^{\alpha_J^* - v_J + 2e_J/\Delta + 2\sigma} \cdot n^{v_J - 2e_J/\Delta - 2\sigma} \\ &\leq (2m)^{e_J/\Delta} \cdot n^{v_J - 2e_J/\Delta - 2\sigma} \\ &\leq (2K' \cdot n^2 p^\Delta \log(1/p))^{e_J/\Delta} n^{v_J - 2e_J/\Delta - 2\sigma} \\ &= n^{-2\sigma} \cdot (2K' \cdot \log(1/p))^{e_J/\Delta} \cdot n^{v_J} p^{e_J}. \end{aligned}$$

Similarly, if $n < 2m$, there is a positive $\sigma = \sigma(H)$ such that the right-hand side of (51) is bounded from above by

$$\begin{aligned} (2m)^{v_J - \alpha_J^*} \cdot n^{2\alpha_J^* - v_J} &\leq (2K' \cdot n^2 p^\Delta \log(1/p))^{v_J - \alpha_J^*} \cdot n^{2\alpha_J^* - v_J} \\ &= p^{\Delta(v_J - \alpha_J^*) - e_J} \cdot (2K' \cdot \log(1/p))^{v_J - \alpha_J^*} \cdot n^{v_J} p^{e_J} \\ &\leq p^{2\sigma} \cdot (2K' \cdot \log(1/p))^{v_J - \alpha_J^*} \cdot n^{v_J} p^{e_J}. \end{aligned}$$

Since $p \leq C^{-2/\Delta}$ and $Cn^{-1} \leq p$, it follows that, in both cases,

$$|\text{Emb}(J, G^*)| \leq p^\sigma \cdot n^{v_J} p^{e_J} \cdot \frac{\gamma}{e_J \cdot (K')^2 \cdot \log(1/p) \cdot 2^{e_H + v_H}},$$

provided that C is sufficiently large. Substituting this inequality into (50) proves the claim. \square

The next claim shows that the endpoints of every edge of $G^* \setminus G_{\text{exc}}^*$ satisfy a certain degree restriction.

Claim 7.6. *There is a positive $\gamma' = \gamma'(H, K')$ such that, for every edge uv of $G^* \setminus G_{\text{exc}}^*$, either*

$$\deg_{G^*} u \cdot \deg_{G^*} v \geq \frac{\gamma'}{(\log(1/p))^{v_H}} \cdot m \quad \text{or} \quad \deg_{G^*} u + \deg_{G^*} v \geq \frac{\gamma'}{(\log(1/p))^{v_H/2}} \cdot n.$$

Moreover, if the second inequality fails, then uv is contained in at least $\gamma' n^{v_H} p^{e_H} / m_{\text{max}}$ copies of H in G^* .

Proof. Suppose that uv is an edge of $G^* \setminus G_{\text{exc}}^*$. It follows from (49) and the definition of G_{exc}^* that there is some $J \in \mathcal{Q}_H$ such that

$$|\text{Emb}(J, G; uv)| \geq N(J, G; uv) \geq \frac{\gamma}{|\mathcal{Q}_H|} \cdot \frac{n^{v_J} p^{e_J}}{m_{\text{max}}}. \quad (52)$$

Let $\mu = n^2 p^\Delta / m$ and observe that $\mu \geq (K' \cdot \log(1/p))^{-1} \cdot m_{\text{max}} / m$. We split the remainder of the proof into two cases, depending on whether or not $J = H$.

Assume first $J = H$. Then, since $n^{v_H} p^{e_H} = (n^2 p^\Delta)^{v_H/2} = (\mu m)^{v_H/2}$, Lemma 5.7 and (52) imply that

$$\begin{aligned} (4 \deg_{G^*} u \cdot \deg_{G^*} v)^{\frac{\Delta-1}{\Delta}} &\geq |\text{Emb}(J, G; uv)| \cdot \frac{(2m)^{\frac{2\Delta-1}{\Delta} - \frac{v_H}{2}}}{4e_H} \\ &\geq \frac{\gamma}{|\mathcal{Q}_H|} \cdot \frac{(\mu m)^{v_H/2}}{m_{\text{max}}} \cdot \frac{(2m)^{\frac{2\Delta-1}{\Delta} - \frac{v_H}{2}}}{4e_H} \\ &= \frac{2^{\frac{2\Delta-1}{\Delta} - \frac{v_H}{2}} \gamma}{4e_H |\mathcal{Q}_H|} \cdot \mu^{v_H/2} \cdot \frac{m}{m_{\text{max}}} \cdot m^{\frac{\Delta-1}{\Delta}} \\ &\geq (4\gamma')^{\frac{\Delta-1}{\Delta}} \cdot (\log(1/p))^{-v_H/2} \cdot m^{\frac{\Delta-1}{\Delta}}, \end{aligned}$$

for a suitable positive constant $\gamma' = \gamma'(H, K')$. Since $\Delta \geq 2$, this implies the claimed lower bound on $\deg_{G^*} u \cdot \deg_{G^*} v$.

Assume now that $J \neq H$. In this case, J admits a bipartition $V(J) = A \cup B$ such that every vertex in A has degree Δ . In particular, $e_J = |A| \cdot \Delta$; moreover, as $J \neq H$ and H is connected

and Δ -regular, we also have $|B| > |A|$. Since $n^{v_J} p^{e_J} = (n^2 p^\Delta)^{|A|} \cdot n^{|B|-|A|} = (\mu m)^{|A|} \cdot n^{|B|-|A|}$, it follows from Lemma 5.8 and (52) that

$$\begin{aligned} \deg_{G^*} u + \deg_{G^*} v &\geq \frac{|\text{Emb}(J, G; uv)|}{e_J \cdot (2m)^{|A|-1} \cdot (\min\{m, n\})^{|B|-|A|-1}} \\ &\geq \frac{\gamma}{|\mathcal{Q}_H|} \cdot \frac{(\mu m)^{|A|} \cdot n^{|B|-|A|}}{m_{\max}} \cdot \frac{1}{e_J \cdot (2m)^{|A|-1} \cdot (\min\{m, n\})^{|B|-|A|-1}} \\ &\geq \frac{\gamma}{|\mathcal{Q}_H|} \cdot \frac{\mu^{|A|} m}{m_{\max}} \cdot \frac{n}{e_J \cdot 2^{|A|-1}} \\ &\geq \gamma' \cdot (\log(1/p))^{-|A|} \cdot n \end{aligned}$$

for a suitable positive constant γ' . Since $|A| < v_H/2$, this implies the claimed lower bound on $\deg_{G^*} u + \deg_{G^*} v$. In particular, if the second inequality in the statement of the claim fails, then the above shows that $J = H$, and thus the second assertion of the claim follows from (52). \square

To prove (47), we will further distinguish between the cases $p^{\Delta/2} \geq n^{-1/2-\xi}$ and $p^{\Delta/2} \leq n^{-1/2-\xi}$, for a small constant ξ . The first case is easier and is handled by our next claim.

Claim 7.7. *There is a positive constant $\xi = \xi(\Delta, \varepsilon)$ such that (47) holds if $p^{\Delta/2} \geq n^{-1/2-\xi}$.*

Proof. To simplify the presentation, we shall prove (47) with 7ε instead of $\varepsilon/2$. Suppose that G^* is a core with m edges, where $m_{\min} \leq m \leq m_{\max}$. For each $k \in \mathbb{Z}$, we define the set

$$B_k = \{v \in V(G^*) : \deg_{G^*} v \geq e^k \sqrt{mp}^\varepsilon\}$$

and let k_{\max} be the largest integer such that $e^{k_{\max}} \sqrt{m} \leq np^{2\varepsilon}$. Note that the bounds $n^2 p^\Delta / K' \leq m \leq K' \cdot n^2 p^\Delta \log(1/p)$ and our assumption $p \leq C^{-2/\Delta}$ imply that $0 \leq k_{\max} \leq (\Delta/2) \cdot \log(1/p)$ whenever C is large enough.

We will first prove that Claim 7.6 implies that, for each edge uv of $G^* \setminus G_{\text{exc}}^*$, either

- (i) uv has an endpoint in $B_{k_{\max}}$ or
- (ii) $u \in B_k$ and $v \in B_{-k}$ for some $k \in \{-k_{\max}, \dots, k_{\max}\}$.

Indeed, if we suppose that

$$\deg_{G^*} u + \deg_{G^*} v \geq \frac{\gamma' \cdot n}{(\log(1/p))^{v_H/2}},$$

then (i) follows immediately (with room to spare) because $\frac{\gamma' \cdot n}{(\log(1/p))^{v_H/2}} \geq np^{3\varepsilon} \geq 2e^{k_{\max}} \sqrt{mp}^\varepsilon$. On the other hand, if (i) does not hold (i.e., neither u nor v is in $B_{k_{\max}}$), then the above lower bound on $\deg_{G^*} u + \deg_{G^*} v$ does not hold, and Claim 7.6 guarantees

$$\deg_{G^*} u \cdot \deg_{G^*} v \geq \frac{\gamma' \cdot m}{(\log(1/p))^{v_H}} \geq emp^{2\varepsilon}.$$

Observe that the lower bound above and the assumption on u and v imply that

$$\min\{\deg_{G^*} u, \deg_{G^*} v\} = \frac{\deg_{G^*} u \cdot \deg_{G^*} v}{\max\{\deg_{G^*} u, \deg_{G^*} v\}} > \frac{emp^{2\varepsilon}}{e^{k_{\max}} \sqrt{mp}^\varepsilon} = e^{-k_{\max}+1} \sqrt{mp}^\varepsilon.$$

Let $k \in \{-k_{\max} + 1, \dots, k_{\max} - 1\}$ be the largest index such that $\{u, v\} \subseteq B_k$. Without loss of generality, $u \notin B_{k+1}$, so

$$emp^{2\varepsilon} \leq \deg_{G^*} u \cdot \deg_{G^*} v < e^{k+1} \sqrt{mp}^\varepsilon \cdot \deg_{G^*} v,$$

which implies (ii).

To make use of this property of the edges of $G^* \setminus G_{\text{exc}}^*$, we also require upper bounds on the cardinalities of the sets B_k . To that end, observe that, for all $k \in \mathbb{Z}$,

$$2m \geq \sum_{v \in B_k} \deg_{G^*} v \geq |B_k| \cdot e^k \sqrt{mp}^\varepsilon$$

and hence

$$|B_k| \leq 2e^{-k} \sqrt{mp}^{-\varepsilon}. \quad (53)$$

We claim that there is a positive constant $\xi = \xi(\Delta, \varepsilon)$ such that $n^2 p^\Delta \geq 2K' n p^\varepsilon \log n$ whenever $p^{\Delta/2} \geq n^{-1/2-\xi}$ and C is sufficiently large. Indeed, if $p^{\Delta/2} \geq n^{-1/3}$, then $n^2 p^\Delta \geq n^{4/3} \geq 2K' n p^\varepsilon \log n$, provided that C is sufficiently large; otherwise, $p^\varepsilon \leq n^{-2\varepsilon/(3\Delta)}$ and our assumption implies that $n^2 p^\Delta \geq n^{1-2\xi} \geq K' n p^\varepsilon \log n$, provided that ξ is sufficiently small and C is sufficiently large. Now inequality (53) and the choice of k_{\max} imply

$$|B_{-k_{\max}}| \leq 2e^{k_{\max}} \sqrt{mp}^{-\varepsilon} \leq 2np^\varepsilon \leq \frac{m}{\log n} \quad (54)$$

and (since $e^{k_{\max}+1} > np^{2\varepsilon}$)

$$|B_{k_{\max}}| \cdot n \leq 2e^{-k_{\max}} \sqrt{mp}^{-\varepsilon} \cdot n < 2emp^{-3\varepsilon}. \quad (55)$$

Recall from Claim 7.5 that $e_{G_{\text{exc}}^*} \leq p^\sigma \cdot m_{\min}$ for a positive constant σ depending only on H . It follows that we may construct each $G^* \in \mathcal{I}_m^*$ as follows:

- (1) Choose some $m_{\text{exc}} \leq p^\sigma \cdot m_{\min}$ and then choose m_{exc} edges of K_n to form G_{exc}^* .
- (2) Choose the sets $B_{-k_{\max}}, \dots, B_{k_{\max}}$ and then choose $m - m_{\text{exc}}$ edges from

$$\mathcal{B} = \{uv \in E(K_n) : u \in B_{k_{\max}}\} \cup \bigcup_{k=0}^{k_{\max}} \{uv \in E(K_n) : u \in B_k, v \in B_{-k}\}$$

to form $G^* \setminus G_{\text{exc}}^*$.

Since $B_{k_{\max}} \subseteq \dots \subseteq B_{-k_{\max}}$ and $|B_{-k_{\max}}| \leq m/\log n$ by (54), the number of ways to choose the sets $B_{-k_{\max}}, \dots, B_{k_{\max}}$ is at most

$$((2k_{\max} + 2) \cdot n)^{m/\log n} \leq e^{2m},$$

using the (very crude) bound $k_{\max} \leq n/4$. Moreover, inequalities (53) and (55) imply that

$$|\mathcal{B}| \leq |B_{k_{\max}}| \cdot n + \sum_{k=0}^{k_{\max}} |B_k| \cdot |B_{-k}| \leq 2emp^{-3\varepsilon} + (k_{\max} + 1) \cdot 4mp^{-2\varepsilon} \leq mp^{-4\varepsilon},$$

where we use $k_{\max} \leq (\Delta/2) \cdot \log(1/p)$ and $p^{\Delta/2} \leq 1/C$ for a large enough C . We conclude that

$$|\mathcal{I}_m^*| \leq e^{2m} \cdot \sum_{m_{\text{exc}}=0}^{p^\sigma \cdot m_{\min}} \binom{n^2}{m_{\text{exc}}} \cdot \binom{mp^{-4\varepsilon}}{m - m_{\text{exc}}}.$$

In order to bound the right-hand side above, we note that

$$\binom{mp^{-4\varepsilon}}{m - m_{\text{exc}}} \leq \binom{mp^{-4\varepsilon}}{m} \leq \left(\frac{emp^{-4\varepsilon}}{m}\right)^m \leq p^{-5\varepsilon m}.$$

Moreover, using the inequalities $\sum_{i=0}^k \binom{n}{i} \leq (en/k)^k$ and $m \geq m_{\min} \geq n^2 p^\Delta / K'$,

$$\sum_{m_{\text{exc}}=0}^{p^\sigma \cdot m_{\min}} \binom{n^2}{m_{\text{exc}}} \leq \left(\frac{en^2}{p^\sigma \cdot m_{\min}}\right)^{p^\sigma \cdot m_{\min}} \leq \left(\frac{eK'}{p^{\Delta+\sigma}}\right)^{mp^\sigma} \leq p^{-\varepsilon m}.$$

We may conclude that

$$|\mathcal{I}_m^*| \leq e^{2m} \cdot p^{-6\varepsilon m} \leq p^{-7\varepsilon m},$$

which completes the proof of (47) (with 7ε instead of $\varepsilon/2$). \square

The argument above shows that we do not need the assumption that H is nonbipartite when $p^{\Delta/2} \geq n^{-1/2-\xi}$. In the following, we will assume that $p^{\Delta/2} \leq n^{-1/2-\xi}$ and that H is not bipartite.

Let γ' be the constant from Claim 7.6. Our assumption on p implies that

$$m_{\max} \leq K' \cdot n^2 p^\Delta \log(1/p) \leq K' \cdot n^{1-2\xi} \log(1/p) < \frac{\gamma' \cdot n}{(\log(1/p))^{v_H/2}} - 1,$$

where the last inequality holds because $n \geq C$ and C is large. In particular, if $G^* \in \mathcal{I}_m^*$ for some $m_{\min} \leq m \leq m_{\max}$, then for any two vertices $u, v \in V(G^*)$, we have $\deg_{G^*}(u) + \deg_{G^*}(v) \leq m_{\max} + 1 < \gamma' n / (\log(1/p))^{v_H/2}$, so it follows from Claim 7.6 that

$$\deg_{G^*} u \cdot \deg_{G^*} v \geq \frac{\gamma'}{(\log(1/p))^{v_H}} \cdot m \quad \text{for every } uv \in E(G^* \setminus G_{\text{exc}}^*) \quad (56)$$

and that every edge uv of $G^* \setminus G_{\text{exc}}^*$ belongs to at least $\gamma' n^{v_H} p^{e_H} / m_{\max}$ copies of H in G^* . Set

$$\beta = \Delta(1 + v_H/2) + 1/2,$$

let G_{high}^* comprise all edges uv of G^* such that

$$\deg_{G^*} u \cdot \deg_{G^*} v \geq (\log(1/p))^\beta \cdot m, \quad (57)$$

and denote by m_{high} the number of edges in G_{high}^* . We claim that

$$m_{\text{high}} \leq \frac{8m}{(\log(1/p))^\beta}. \quad (58)$$

In order to show it, we estimate the number of copies of P_4 , the path with four vertices (and three edges), in G^* in two different ways. On the one hand,

$$2N(P_4, G^*) = |\text{Emb}(P_4, G^*)| \leq (2m)^2,$$

since every embedding of P_4 into G^* is determined by the images of its two nonincident edges. On the other hand,

$$\begin{aligned} N(P_4, G^*) &\geq \sum_{uv \in E(G^*)} (\deg_{G^*} u - 1) \cdot (\deg_{G^*} v - 2) \\ &= \sum_{uv \in E(G^*)} \deg_{G^*} u \cdot \deg_{G^*} v - 3 \sum_{v \in V(G^*)} (\deg_{G^*} v)^2 + 2m \\ &\geq m_{\text{high}} \cdot (\log(1/p))^\beta \cdot m - 3m \sum_{v \in V(G^*)} \deg_{G^*} v \\ &= m_{\text{high}} \cdot (\log(1/p))^\beta \cdot m - 6m^2. \end{aligned}$$

These two lower and upper bounds on $N(P_4, G^*)$ imply (58).

Claim 7.8. *Suppose that φ is an embedding of H into $G^* \setminus (G_{\text{exc}}^* \cup G_{\text{high}}^*)$. Then for every $a \in V(H)$,*

$$\deg_{G^*} \varphi(a) \geq \frac{m^{1/2}}{(\log(1/p))^{\beta \cdot v_H}}.$$

Proof. Define $f: V(H) \rightarrow \mathbb{R}$ by

$$f(a) = \log \left(\frac{\deg_{G^*} \varphi(a)}{m^{1/2}} \right)$$

and let

$$f^* = \beta \log \log(1/p).$$

It suffices to show that $f(a) \geq -v_H f^*$ for every $a \in V(H)$. To this end, note that (56) and our definition of G_{high}^* (see (57)) imply that

$$-f^* \leq f(a) + f(b) \leq f^* \quad \text{for every } ab \in E(H), \quad (59)$$

since $\beta \geq v_H + 1$ and $p \leq C^{-2/\Delta}$ for a large constant C . Since H is not bipartite, it contains an odd cycle. Let Z be one such cycle and suppose that $a_0, \dots, a_{2\ell}$ are its vertices (listed in an arbitrarily chosen cyclic ordering). It follows from (59), applied to all $2\ell + 1$ edges of Z , that

$$2f(a_0) = f(a_0) + f(a_{2\ell}) + \sum_{i=0}^{2\ell-1} (-1)^i (f(a_i) + f(a_{i+1})) \in [-(2\ell + 1)f^*, (2\ell + 1)f^*].$$

Since the particular choice of a_0 among all vertices of Z was arbitrary, we may conclude that

$$-(2\ell + 1)f^* \leq f(a) \leq (2\ell + 1)f^* \quad \text{for every } a \in V(Z),$$

with room to spare. Since $2\ell + 1 = v_Z \leq v_H$, this proves the desired inequality for all $a \in V(Z)$. Suppose now that $b \in V(H) \setminus V(Z)$. Since H is connected, it contains a path from b to Z . Let $b_0, b_1, \dots, b_{\ell'}$, where $b_0 = b$ and $b_{\ell'} \in V(Z)$, be the vertices of a shortest such path (listed in their natural order) and note that $\ell' + 2\ell + 1 \leq v_H$. It follows from (59), applied to all ℓ' edges of the path, that

$$f(b) + (-1)^{\ell'-1} f(b_{\ell'}) = \sum_{i=0}^{\ell'-1} (-1)^i (f(b_i) + f(b_{i+1})) \in [-\ell' f^*, \ell' f^*]$$

and consequently, as $b_{\ell'} \in V(Z)$, that

$$f(b) \geq -\ell' f^* - |f(b_{\ell'})| \geq -(\ell' + 2\ell + 1)f^* \geq -v_H f^*,$$

as claimed. \square

Let G_{bad}^* comprise all edges of G^* that do not belong to a copy of H in the graph $G^* \setminus (G_{\text{exc}}^* \cup G_{\text{high}}^*)$ and let m_{bad} be the number of such edges; note that $G_{\text{bad}}^* \supseteq G_{\text{exc}}^* \cup G_{\text{high}}^*$. Since each edge of $G_{\text{bad}}^* \setminus G_{\text{exc}}^*$ belongs to at least $\gamma' n^{v_H} p^{e_H} / m_{\text{max}}$ copies of H in G^* , none of which are in $G^* \setminus G_{\text{exc}}^* \cup G_{\text{high}}^*$, we have

$$(m_{\text{bad}} - m_{\text{exc}}) \cdot \frac{\gamma' n^{v_H} p^{e_H}}{m_{\text{max}}} \leq \sum_{e \in E(G_{\text{exc}}^* \cup G_{\text{high}}^*)} |\text{Emb}(H, G^*; e)|.$$

It follows from Lemma 5.9, Claim 7.5, and inequality (58) that

$$\sum_{e \in E(G_{\text{exc}}^* \cup G_{\text{high}}^*)} |\text{Emb}(H, G^*; e)| \leq e_H \cdot (2m)^{v_H/2} \cdot \left(\frac{m_{\text{exc}} + m_{\text{high}}}{m} \right)^{1/\Delta} \leq 8e_H \cdot \frac{(2m)^{v_H/2}}{(\log(1/p))^{\beta/\Delta}}.$$

Consequently,

$$\begin{aligned} m_{\text{bad}} &\leq m_{\text{exc}} + \frac{m_{\text{max}}}{\gamma' n^{v_H} p^{e_H}} \cdot 8e_H \cdot \frac{(2m)^{v_H/2}}{(\log(1/p))^{\beta/\Delta}} \\ &\leq p^\sigma \cdot m_{\text{min}} + \frac{8 \cdot 2^{v_H/2} \cdot e_H}{\gamma'} \cdot \frac{m_{\text{max}}^{v_H/2}}{n^{v_H} p^{e_H} (\log(1/p))^{\beta/\Delta}} \cdot m \\ &\leq \left(p^\sigma + \frac{8 \cdot 2^{v_H/2} \cdot e_H}{\gamma'} \cdot \frac{(K' \cdot n^2 p^\Delta \log(1/p))^{v_H/2}}{n^{v_H} p^{e_H} (\log(1/p))^{\beta/\Delta}} \right) \cdot m \\ &= \left(p^\sigma + \frac{8 \cdot e_H \cdot (2K')^{v_H/2}}{\gamma' \cdot (\log(1/p))^{\beta/\Delta - v_H/2}} \right) \cdot m \leq \frac{m}{\log(1/p)}, \end{aligned}$$

since $\beta/\Delta = 1 + v_H/2 + 1/(2\Delta)$ and $p \leq C^{-2/\Delta}$ for a large constant C .

Let U be the set of nonisolated vertices of $G^* \setminus G_{\text{bad}}^*$. It follows from Claim 7.8 that

$$|U| \cdot \frac{m^{1/2}}{(\log(1/p))^{\beta \cdot v_H}} \leq \sum_{v \in U} \deg_{G^*} v \leq 2m$$

and thus

$$|U| \leq 2m^{1/2} \cdot (\log(1/p))^{\beta \cdot v_H} \leq m/\log n,$$

as $p^{\Delta/2} \geq Cn^{-1}(\log n)^{\Delta v_H^2}$ and C is large enough.

To summarise, we may construct each $G^* \in \mathcal{I}_m^*$ as follows:

- (1) Choose $m_{\text{bad}} \leq m/\log(1/p)$ and the m_{bad} edges of K_n to form G_{bad}^* .
- (2) Choose the vertices of U and the edges of $G^* \setminus G_{\text{bad}}^*$ from the set $\binom{U}{2}$.

Using the above bounds on the size of U , we conclude that

$$|\mathcal{I}_m^*| \leq \sum_{m_{\text{bad}}=0}^{m/\log(1/p)} \binom{n^2}{m_{\text{bad}}} \cdot n^{m/\log n} \cdot \binom{4m \cdot (\log(1/p))^{2\beta \cdot v_H}}{m - m_{\text{bad}}}.$$

In order to bound the right-hand side from above, we note that, for sufficiently large C ,

$$\binom{4m \cdot (\log(1/p))^{2\beta \cdot v_H}}{m - m_{\text{bad}}} \leq \binom{4m \cdot (\log(1/p))^{2\beta \cdot v_H}}{m} \leq \left(\frac{4em \cdot (\log(1/p))^{2\beta \cdot v_H}}{m} \right)^m \leq p^{-\varepsilon m/6}$$

and, since $m \geq n^2 p^\Delta / K'$,

$$\sum_{m_{\text{bad}}=0}^{m/\log(1/p)} \binom{n^2}{m_{\text{bad}}} \leq \left(\frac{en^2 \log(1/p)}{m} \right)^{m/\log(1/p)} \leq \left(\frac{eK' \cdot \log(1/p)}{p^\Delta} \right)^{m/\log(1/p)} \leq p^{-\varepsilon m/6}.$$

Since $n^{m/\log n} = e^m \leq p^{-\varepsilon m/6}$, we may conclude that $|\mathcal{I}_m^*| \leq p^{-\varepsilon m/2}$, which completes the proof of Proposition 7.1.

8. THE POISSON REGIME

Given a nonnegative real μ , we shall denote by $\text{Po}(\mu)$ the Poisson distribution with mean μ . Suppose that $X \sim \text{Po}(\mu)$. A classical result in large deviation theory is that, for every fixed $\delta > 0$,

$$-\log \mathbb{P}(X \geq (1 + \delta)\mu) = ((1 + \delta) \log(1 + \delta) - \delta)\mu + o(\mu),$$

as $\mu \rightarrow \infty$. Motivated by this estimate, for any random variable X with positive expectation and any $\delta > 0$, we define

$$\Psi_X(\delta) = ((1 + \delta) \log(1 + \delta) - \delta) \mathbb{E}[X].$$

Theorems 1.4 and 1.6 follow immediately from the following two propositions.

Proposition 8.1. *For every integer $k \geq 3$ and all positive real numbers ε and δ , there exists a positive constant C such that the following holds. Suppose that $N \in \mathbb{N}$ and $p \in (0, 1)$ satisfy $CN^{-1} \leq p^{k/2} \leq C^{-1}N^{-1} \log N$. Then $X = X_{N,p}^{k\text{-AP}}$ satisfies*

$$(1 - \varepsilon)\Psi_X(\delta) \leq -\log \mathbb{P}(X \geq (1 + \delta) \mathbb{E}[X]) \leq (1 + \varepsilon)\Psi_X(\delta).$$

Proposition 8.2. *For every $\Delta \geq 2$, every connected, Δ -regular graph H , and all positive real numbers ε and δ , there exists a positive constant C such that the following holds. Suppose that $n \in \mathbb{N}$ and $p \in (0, 1)$ satisfy $Cn^{-1} \leq p^{\Delta/2} \leq C^{-1}n^{-1}(\log n)^{\frac{1}{v_H - 2}}$. Then $X = X_{n,p}^H$ satisfies*

$$(1 - \varepsilon)\Psi_X(\delta) \leq -\log \mathbb{P}(X \geq (1 + \delta) \mathbb{E}[X]) \leq (1 + \varepsilon)\Psi_X(\delta).$$

It is not difficult to show that the requirements on p in these results are optimal up to the choice of the constant C . Indeed, if $X = X_{N,p}^{k\text{-AP}}$, then planting an interval of length $C_\delta N p^{k/2}$ creates $(1 + \delta) \mathbb{E}[X]$ k -term arithmetic progressions (for a sufficiently large C_δ), which shows that

$$-\log \mathbb{P}(X \geq (1 + \delta) \mathbb{E}[X]) = O(N p^{k/2} \log(1/p)).$$

Similarly, if $X = X_{n,p}^H$, then planting a clique of size $C_\delta n p^{\Delta/2}$ results in $(1 + \delta) \mathbb{E}[X]$ copies of H , which proves

$$-\log \mathbb{P}(X \geq (1 + \delta) \mathbb{E}[X]) = O(n^2 p^\Delta \log(1/p)).$$

The upper bounds are $o(\Psi_X(\delta))$, and therefore dominate the Poisson bounds, whenever $p^{k/2} \gg N^{-1} \log N$ or $p^{\Delta/2} \gg n^{-1}(\log n)^{\frac{1}{\nu_H-2}}$, respectively.

8.1. Poisson approximation via factorial moments. For a real number m and a nonnegative integer t , write $m^{\underline{t}} = m(m-1)\cdots(m-t+1)$ for the t -th falling factorial of m . For a random variable X , let $M_t(X) = \mathbb{E}[X^{\underline{t}}]$ be the t -th factorial moment of X . It is straightforward to verify that, if $X \sim \text{Po}(\mu)$, then $M_t(X) = \mu^t$ for all $t \geq 0$.

A classical application of the method of moments is that, if (X_n) is a sequence of random variables whose t -th factorial moments converge to μ^t for some fixed μ , then X_n converges in distribution to $\text{Po}(\mu)$. The lemma below can be viewed as an extension of this result to the case when $\mu \rightarrow \infty$. It states that, if the t -th factorial moment of some random variable X is approximately μ^t , for each t around $\delta\mu$, then the logarithmic upper tail probability $-\log \mathbb{P}(X \geq (1+\delta)\mathbb{E}[X])$ is well approximated by $\Psi_X(\delta)$.

Proposition 8.3. *For all positive real numbers ε and δ , there exists a positive constant η such that the following holds. Let X be a nonnegative integer-valued random variable with mean $\mu \geq 1/\eta$ such that $|M_t(X) - \mu^t| \leq \eta\mu^t$ for every integer t satisfying $(\delta - \varepsilon)\mu \leq t \leq (\delta + \varepsilon)\mu$. Then*

$$(1 - \varepsilon)\Psi_X(\delta) \leq -\log \mathbb{P}(X \geq (1 + \delta)\mathbb{E}[X]) \leq (1 + \varepsilon)\Psi_X(\delta).$$

Define the continuous function $I: [0, \infty) \rightarrow [0, \infty)$ by

$$I(\delta) = (1 + \delta) \log(1 + \delta) - \delta,$$

so that $\Psi_X(\delta) = I(\delta) \cdot \mu$. Note that $I(\delta) > 0$ whenever $\delta > 0$.

Lemma 8.4. *For every nonnegative integer t and every positive real x ,*

$$\log(x+t)^{\underline{t}} = I(t/x) \cdot x + t \log x + \lambda(x, t),$$

where $0 \leq \lambda(x, t) \leq (t+1)/x$.

Proof. Observe that $\log(x+t)^{\underline{t}} = \sum_{s=1}^t \log(x+s)$. Since \log is an increasing function,

$$\int_x^{x+t} \log y \, dy \leq \sum_{s=1}^t \log(x+s) \leq \int_{x+1}^{x+t+1} \log y \, dy.$$

Recalling that $\int \log y \, dy = y(\log y - 1) + C$, we have

$$\int_x^{x+t} \log y \, dy = ((1+t/x) \log(1+t/x) - t/x) \cdot x + t \log x = I(t/x) \cdot x + t \log x.$$

On the other hand,

$$\int_{x+1}^{x+t+1} \log y \, dy - \int_x^{x+t} \log y \, dy \leq \log(x+t+1) - \log x = \log(1+(t+1)/x) \leq (t+1)/x.$$

This proves the claimed estimate. \square

Proof of Proposition 8.3. Since, for every positive integer t , the function $x \mapsto x^{\underline{t}}$ is increasing on $[t-1, \infty)$ and nonnegative on $\mathbb{Z}_{\geq 0}$, Markov's inequality implies that

$$\mathbb{P}(X \geq (1+\delta)\mathbb{E}[X]) \leq \mathbb{P}(X^{\underline{t}} \geq ((1+\delta)\mathbb{E}[X])^{\underline{t}}) \leq \frac{M_t(X)}{((1+\delta)\mathbb{E}[X])^{\underline{t}}}$$

for every positive integer $t \leq (1+\delta)\mathbb{E}[X]$. This implies that, for every such t ,

$$-\log \mathbb{P}(X \geq (1+\delta)\mathbb{E}[X]) \geq \log((1+\delta)\mu)^{\underline{t}} - \log M_t(X). \quad (60)$$

Let $t = \lfloor \delta\mu \rfloor$. Since $(1+\delta)\mu - t \geq \mu$, it follows from Lemma 8.4 that

$$\log((1+\delta)\mu)^{\underline{t}} \geq \log(\mu+t)^{\underline{t}} \geq I(t/\mu) \cdot \mu + t \log \mu.$$

On the other hand, our assumption implies that

$$\log M_t(X) \leq \log((1 + \eta)\mu^t) = t \log \mu + \log(1 + \eta) \leq t \log \mu + \eta.$$

Finally, since I is continuous, $I(\delta) > 0$, and $|t/\mu - \delta| \leq 1/\mu \leq \eta$, substituting the above two inequalities into (60) yields

$$-\log \mathbb{P}(X \geq (1 + \delta) \mathbb{E}[X]) \geq I(\delta) \cdot \mu - \varepsilon \Psi_X(\delta) = (1 - \varepsilon) \Psi_X(\delta),$$

provided that η is sufficiently small (as a function of ε and δ).

For the upper bound, we will use the tilting argument, which is a standard trick of large deviation theory. For the sake of brevity, let $m_t = M_t(X)$ for every nonnegative integer t . Let $t = (\delta + \gamma)\mu$, where γ is a small positive constant that depends on ε and δ (but not on η). Since η is allowed to depend on γ and $\mu \geq 1/\eta$, we may assume that t is an integer. The idea is to consider a ‘tilted’ random variable \tilde{X} defined by the relation

$$\mathbb{P}(\tilde{X} = x) = \frac{\mathbb{P}(X = x) \cdot x^t}{m_t} \quad \text{for every } x \in \mathbb{Z}.$$

The definition of m_t ensures that \tilde{X} is a well-normalised random variable. In particular,

$$\mathbb{E}[g(\tilde{X})] = \frac{\mathbb{E}[g(X) \cdot X^t]}{m_t} \quad \text{for every } g: \mathbb{Z} \rightarrow \mathbb{Z}. \quad (61)$$

Using the identities

$$x \cdot x^t = x^{t+1} + t \cdot x^t \quad \text{and} \quad x^2 \cdot x^t = x^{t+2} + (2t + 1) \cdot x^{t+1} + t^2 \cdot x^t,$$

we have

$$\mathbb{E}[\tilde{X}] = \frac{\mathbb{E}[X \cdot X^t]}{m_t} = \frac{m_{t+1}}{m_t} + t$$

and

$$\mathbb{E}[\tilde{X}^2] = \frac{\mathbb{E}[X^2 \cdot X^t]}{m_t} = \frac{m_{t+2} + (2t + 1)m_{t+1}}{m_t} + t^2,$$

and so

$$\text{Var}[\tilde{X}] = \mathbb{E}[\tilde{X}^2] - \mathbb{E}[\tilde{X}]^2 = \frac{m_{t+2}m_t - m_{t+1}^2 + m_{t+1}m_t}{m_t^2}.$$

Since $t = (\delta + \gamma)\mu$, we have $(\delta - \varepsilon)\mu \leq t \leq t + 2 \leq (\delta + \varepsilon)\mu$, provided that γ is sufficiently small as a function of δ and ε . Since the assumptions of the proposition imply that m_s is well approximated by μ^s for each $s \in \{t, t + 1, t + 2\}$, a straightforward computation yields

$$(1 + \delta + \gamma/2)\mu \leq \mathbb{E}[\tilde{X}] \leq (1 + \delta + 3\gamma/2)\mu \quad \text{and} \quad \text{Var}[\tilde{X}] \leq 10\eta\mu^2 + 2\mu,$$

provided that η is sufficiently small. Therefore, Chebyshev’s inequality yields

$$\mathbb{P}(|\tilde{X} - (1 + \delta + \gamma)\mu| \geq \gamma\mu) \leq \mathbb{P}(|\tilde{X} - \mathbb{E}[\tilde{X}]| \geq \gamma\mu/2) \leq \frac{4 \cdot (10\eta\mu^2 + 2\mu)}{\gamma^2\mu^2} \leq \varepsilon. \quad (62)$$

Next, using (61) with the function $g(x) = \mathbb{1}[|x - (1 + \delta + \gamma)\mu| < \gamma\mu] \cdot (m_t/x^t)$, we see that

$$\mathbb{P}(X \geq (1 + \delta)\mu) \geq \mathbb{P}(|X - (1 + \delta + \gamma)\mu| \leq \gamma\mu) = \mathbb{E}[g(\tilde{X})].$$

When $g(\tilde{X})$ is nonzero, then \tilde{X} is bounded from above by $(1 + \delta + 2\gamma)\mu$, and thus

$$\begin{aligned} \mathbb{E}[g(\tilde{X})] &\geq \mathbb{P}(g(\tilde{X}) \neq 0) \cdot \frac{m_t}{(1 + \delta + 2\gamma)^t} \\ &= (1 - \mathbb{P}(|\tilde{X} - (1 + \delta + \gamma)\mu| \geq \gamma\mu)) \cdot \frac{m_t}{(1 + \delta + 2\gamma)^t} \stackrel{(62)}{\geq} \frac{(1 - \varepsilon) \cdot m_t}{(1 + \delta + 2\gamma)^t}. \end{aligned}$$

Using Lemma 8.4 with $x = (1 + \gamma)\mu$, we obtain

$$\log(1 + \delta + 2\gamma)^t \leq I(t/(1 + \gamma)\mu) \cdot (1 + \gamma)\mu + t \log((1 + \gamma)\mu) + \frac{t + 1}{(1 + \gamma)\mu}.$$

On the other hand, our assumptions imply that, when η is small,

$$\log m_t \geq t \log \mu + \log(1 - \eta) \geq t \log \mu - 2\eta.$$

Combining the above bounds, we obtain

$$-\log \mathbb{P}(X \geq (1 + \delta)\mu) \leq I(t/(1 + \gamma)\mu) \cdot (1 + \gamma)\mu + t \log(1 + \gamma) + \frac{t + 1}{(1 + \gamma)\mu} + 2\eta - \log(1 - \varepsilon).$$

Recalling that $t = (\delta + \gamma)\mu$, the continuity of I and the fact that $I(\delta) > 0$ imply that

$$-\log \mathbb{P}(X \geq (1 + \delta)\mu) \leq I(\delta) \cdot \mu + \varepsilon \Psi_X(\delta) = (1 + \varepsilon)\Psi_X(\delta),$$

provided that γ is sufficiently small and η is sufficiently small. \square

8.2. Cluster analysis. We will deduce both Propositions 8.1 and 8.2 from Proposition 8.3 and Lemma 8.5, stated below, by analysing the component structure of certain random hypergraphs. Since the proofs turn out to be quite similar, we adopt a general point of view from the start. Suppose that \mathcal{H} is a hypergraph and, given some $p \in (0, 1)$, denote by \mathcal{H}_p the random induced subhypergraph of \mathcal{H} obtained by keeping every vertex with probability p , independently. The *dependency graph* $G_{\mathcal{H}}$ is the graph on the vertex set $E(\mathcal{H})$ whose edges are all pairs $\{\sigma_1, \sigma_2\}$ such that $\sigma_1 \cap \sigma_2 \neq \emptyset$. A *cluster* is a set $E' \subseteq E(\mathcal{H})$ that induces a connected subgraph in $G_{\mathcal{H}}$. We write $D_s(\mathcal{H}_p)$ for the number of clusters of size s whose elements are edges in \mathcal{H}_p .

Lemma 8.5. *For all positive real numbers c and η , there exists a positive constant K such that the following holds. Let \mathcal{H} be a uniform hypergraph, let $p \in (0, 1)$, and define $X = e_{\mathcal{H}_p}$ and $\mu = \mathbb{E}[X]$. Assume that $K \leq \mu \leq \sqrt{e_{\mathcal{H}}}/K$ and that $\mathbb{E}[D_s(\mathcal{H}_p)] \leq \exp(-Ks)$ for every integer s such that $2 \leq s \leq c\mu$. Then $|M_t(X) - \mu^t| \leq \eta\mu^t$ for every integer t such that $1 \leq t \leq c\mu$.*

Proof. Let $t \leq c\mu$ be a positive integer and let $\mathcal{H}(t)$ denote the set of all sequences of t distinct edges of \mathcal{H} . For each sequence $\bar{\sigma} = (\sigma_1, \dots, \sigma_t) \in \mathcal{H}(t)$, let $X_{\bar{\sigma}}$ be the indicator random variable for the event $\sigma_1 \cup \dots \cup \sigma_t \subseteq V(\mathcal{H}_p)$. Denote the uniformity of \mathcal{H} by k . Our definitions readily imply that $X^t = \sum_{\bar{\sigma} \in \mathcal{H}(t)} X_{\bar{\sigma}}$, that $|\mathcal{H}(t)| = e_{\mathcal{H}}^t$, and that $\mathbb{E}[X_{\bar{\sigma}}] \geq p^{kt}$ for all $\bar{\sigma} \in \mathcal{H}(t)$. Thus

$$M_t(X) = \mathbb{E}[X^t] = \sum_{\bar{\sigma} \in \mathcal{H}(t)} \mathbb{E}[X_{\bar{\sigma}}] \geq e_{\mathcal{H}}^t \cdot p^{kt}.$$

Since, for every $x \geq t$,

$$x^t = x^t \prod_{s=0}^{t-1} \left(1 - \frac{s}{x}\right) \geq x^t \left(1 - \sum_{s=0}^{t-1} \frac{s}{x}\right) \geq x^t \left(1 - \frac{t^2}{x}\right),$$

and $t \leq c\mu \leq c\sqrt{e_{\mathcal{H}}}/K \leq e_{\mathcal{H}}$ for sufficiently large K , we have

$$M_t(X) \geq \left(1 - \frac{t^2}{e_{\mathcal{H}}}\right) \cdot e_{\mathcal{H}}^t p^{kt} = \left(1 - \frac{t^2}{e_{\mathcal{H}}}\right) \cdot \mu^t \geq \left(1 - \frac{c^2}{K^2}\right) \cdot \mu^t \geq (1 - \eta) \cdot \mu^t,$$

provided that K is sufficiently large.

It remains to prove the upper bound. It will be convenient to partition the set $\mathcal{H}(t)$ of sequences according to the component structure of the subgraph of $G_{\mathcal{H}}$ induced by the elements of the sequence. More precisely, given a nonnegative integer ℓ and integers s_1, \dots, s_{ℓ} such that $2 \leq s_1 \leq \dots \leq s_{\ell}$, let $\mathcal{H}(t; s_1, \dots, s_{\ell})$ be the family of all $\bar{\sigma} = (\sigma_1, \dots, \sigma_t) \in \mathcal{H}(t)$ such that the set $\{\sigma_1, \dots, \sigma_t\}$ induces a subgraph in $G_{\mathcal{H}}$ whose ℓ nontrivial connected components (maximal clusters) have sizes s_1, \dots, s_{ℓ} , so that this graph has $t - (s_1 + \dots + s_{\ell})$ isolated vertices.⁷ Observe that, for every collection W_1, \dots, W_{ℓ} of connected subsets of vertices $G_{\mathcal{H}}$ with sizes s_1, \dots, s_{ℓ} , respectively, there are at most $t^{s_1 + \dots + s_{\ell}} \cdot e_{\mathcal{H}}^{t - (s_1 + \dots + s_{\ell})}$ sequences $\bar{\sigma} = (\sigma_1, \dots, \sigma_t) \in \mathcal{H}(t)$ such that the nontrivial connected components of $\{\sigma_1, \dots, \sigma_t\}$ are exactly W_1, \dots, W_{ℓ} ; indeed, there

⁷This includes the case $\ell = 0$ in which $\mathcal{H}(t; \emptyset)$ corresponds to induced subgraphs of $G_{\mathcal{H}}$ all of whose connected components are isolated vertices.

are at most $t^{s_1+\dots+s_\ell}$ ways to choose the locations of the vertices in $W_1 \cup \dots \cup W_\ell$ in a sequence of length t and, for each such choice, at most $e_{\mathcal{H}}^{t-(s_1+\dots+s_\ell)}$ choices for the remaining elements of the sequence. We conclude that

$$\sum_{\bar{\sigma} \in \mathcal{H}(t; s_1, \dots, s_\ell)} \mathbb{E}[X_{\bar{\sigma}}] \leq \mu^{t-(s_1+\dots+s_\ell)} \cdot \prod_{i=1}^{\ell} \mathbb{E}[D_{s_i}(\mathcal{H}_p)] \cdot t^{s_i}$$

and, consequently, summing over all ℓ and all sequences s_1, \dots, s_ℓ and using the assumed upper bound on the expectation of $D_s(\mathcal{H}_p)$, valid for each $s \leq t$,

$$\begin{aligned} M_t(X) &\leq \sum_{s=0}^t \mu^{t-s} \cdot \sum_{\ell \geq 0} \sum_{\substack{s_1+\dots+s_\ell=s \\ 2 \leq s_1 \leq \dots \leq s_\ell}} \prod_{i=1}^{\ell} \mathbb{E}[D_{s_i}(\mathcal{H}_p)] \cdot t^{s_i} \\ &\leq \sum_{s=0}^t \mu^{t-s} \cdot \sum_{\ell \geq 0} \sum_{\substack{s_1+\dots+s_\ell=s \\ 2 \leq s_1 \leq \dots \leq s_\ell}} \prod_{i=1}^{\ell} \exp(-Ks_i + s_i \log t) \\ &= \sum_{s=0}^t \mu^t \cdot \sum_{\ell \geq 0} \sum_{\substack{s_1+\dots+s_\ell=s \\ 2 \leq s_1 \leq \dots \leq s_\ell}} \exp(-Ks + s \log(t/\mu)). \end{aligned}$$

Since, for every $s \geq 0$, there are at most 2^s sequences s_1, \dots, s_ℓ of positive integers whose sum is s (this includes the case when $s = 0$, when the only such sequence is the empty sequence), we have

$$M_t(X) \leq \mu^t \cdot \sum_{s=0}^t \exp(-Ks + s \log(t/\mu) + s).$$

Finally, since $t/\mu \leq c$, we may choose $K = K(c, \eta)$ so that

$$M_t(X) \leq \mu^t \cdot \sum_{s=0}^t \left(\frac{\eta}{1+\eta} \right)^s \leq (1+\eta)\mu^t,$$

completing the proof. \square

8.3. Proof of Proposition 8.1. Let \mathcal{H} be the hypergraph on the vertex set $\llbracket N \rrbracket$ whose edges are k -term arithmetic progressions in $\llbracket N \rrbracket$, so that $X = X_{N,p}^{k\text{-AP}} = e_{\mathcal{H}_p}$. Let $\mu = \mathbb{E}[X]$, let $\eta = \eta(\varepsilon, \delta)$ be the constant from the statement of Proposition 8.3, and let $K = K(\varepsilon, \delta, \eta)$ be the constant from the statement of Lemma 8.5.

For any two integers $a, b \in \llbracket N \rrbracket$ with $a < b$, there is at most one k -term arithmetic progression that starts with a and ends with b (and exactly one such progression if $b - a$ is divisible by $k - 1$). Therefore, $N^2/(2k) \leq e_{\mathcal{H}} \leq N^2$ for all large enough N . In particular, since we assume that $CN^{-1} \leq p^{k/2} \leq C^{-1}N^{-1} \log N$, we find that

$$\frac{C^2}{2k} \leq \mu \leq \left(\frac{\log N}{C} \right)^2 \tag{63}$$

and thus $\max\{1/\eta, K\} \leq \mu \leq \sqrt{e_{\mathcal{H}}}/K$ whenever C is large. The claimed estimate on $-\log \mathbb{P}(X \geq (1+\delta)\mathbb{E}[X])$ will follow from Proposition 8.3 and Lemma 8.5 once we verify that $D_s(\mathcal{H}_p)$, the number of clusters of s arithmetic progressions of length k in the set $\llbracket N \rrbracket_p$, satisfies

$$\mathbb{E}[D_s(\mathcal{H}_p)] \leq \exp(-Ks)$$

for every s satisfying $2 \leq s \leq (\delta + \varepsilon)\mu$.

In order to do so, let $\mathcal{D}(s, m)$ be the set of all clusters $\{\sigma_1, \dots, \sigma_s\}$ of s arithmetic progressions of length k in $\llbracket N \rrbracket$ such that $|\sigma_1 \cup \dots \cup \sigma_s| = m$; we also let $D_{s,m}$ be the number of such clusters whose union is contained in the random set $\llbracket N \rrbracket_p$. When $s \geq 2$, the union of any s

distinct k -term arithmetic progressions contains between $k + 1$ and ks numbers, and therefore $D_{s,m} = 0$ unless $k + 1 \leq m \leq ks$. Thus, we can write

$$D_s(\mathcal{H}_p) = \sum_{m=k+1}^{ks} D_{s,m}.$$

For each integer m , let a_m denote the number of m -element subsets of $\llbracket N \rrbracket$ that are the union of a single, nonempty (but possibly trivial) cluster of k -term arithmetic progressions. Since a progression is uniquely determined by its first and second element, it follows that, for each s ,

$$\mathbb{E}[D_{s,m}] \leq a_m p^m \binom{m^2}{s}. \quad (64)$$

Claim 8.6. *For every integer $m \geq 1$,*

$$a_m \leq N^2 \cdot (2kmN)^{\frac{m-k}{k-1}}.$$

Proof. We prove the claimed upper bound on a_m by induction on m . It is vacuously true when $m < k$, since then $a_m = 0$, or when $m = k$, as $a_k = e_{\mathcal{H}} \leq N^2$. Assume now that $m \geq k + 1$ and let A be an arbitrary set counted by a_m . Since $m > k$, the set A must be a union of at least two different progressions. Moreover, there are a proper subset $A' \subsetneq A$ that is a union of a (nonempty) cluster of k -term arithmetic progressions and a k -term progression σ that intersects A' such that $A = A' \cup \sigma$; note that the number of σ_i 's whose union is A' may be significantly smaller than the number that was used to generate A . By construction, we have that $|A'| = |A| - k + |A' \cap \sigma| = m - k + |A' \cap \sigma|$. Since there are at most $k|A'|N \leq kmN$ arithmetic progressions of length k that intersect A' in exactly one element and at most $k^2|A'|^2 \leq k^2m^2$ progressions that intersect A' in two or more elements,

$$a_m \leq kmN \cdot a_{m-k+1} + k^2m^2 \cdot (a_{m-k+2} + \cdots + a_{m-1}).$$

It follows from our inductive assumption that

$$\begin{aligned} a_m &\leq kmN \cdot N^2 \cdot (2kmN)^{\frac{m-2k+1}{k-1}} + k^2m^2 \cdot k \cdot N^2 \cdot (2kmN)^{\frac{m-k-1}{k-1}} \\ &= N^2 \cdot \left((2kmN)^{\frac{m-k}{k-1}} / 2 + k^3m^2 \cdot (2kmN)^{\frac{m-k}{k-1} - \frac{1}{k-1}} \right). \end{aligned}$$

Finally, as (63) implies that $m \leq ks \leq k(\delta + \varepsilon)\mu \leq k(\delta + \varepsilon)(C^{-1} \log N)^2$ and $N \geq C$, then

$$k^3m^2 \cdot (2kmN)^{-1/(k-1)} \leq 1/2,$$

provided that C is sufficiently large. This implies the claimed upper bound on a_m . \square

Assume now that $k + 1 \leq m \leq ks$. Since $s \leq (\delta + \varepsilon)\mu$, inequality (63) implies that m is only polylogarithmic in N ; on the other hand, $p \leq (C^{-1}N^{-1} \log N)^{2/k}$. Since $k \geq 3$ and $N \geq C$, there is a positive constant γ that depends only on k such that

$$(2kmN)^{1/(k-1)} p \leq (2kmN)^{1/(k-1)} \cdot (C^{-1}N^{-1} \log N)^{2/k} \leq N^{-2(k+1)\gamma} \leq N^{-2m\gamma/(m-k)}.$$

In particular, Claim 8.6 implies that

$$a_m p^m \leq N^2 p^k \cdot \left((2kmN)^{1/(k-1)} p \right)^{m-k} \leq N^2 p^k \cdot N^{-2m\gamma} \leq N^{-m\gamma},$$

where for the last inequality we use that $N^2 p^k$ is at most polylogarithmic in N and $N \geq C$. Combining this bound with (64) we conclude that

$$\mathbb{E}[D_{s,m}] \leq N^{-m\gamma} \cdot \binom{m^2}{s} \leq \exp \left(-m\gamma \log N + s \log \left(\frac{em^2}{s} \right) \right).$$

Let $f: (0, \infty) \rightarrow (0, \infty)$ be the function defined by

$$f(x) = \exp \left(-x\gamma \log N + s \log \left(\frac{ex^2}{s} \right) \right) = \exp \left(-x\gamma \log N + 2s \log x + s \log(e/s) \right),$$

so that $\mathbb{E}[D_{s,m}] \leq f(m)$. Elementary calculus shows that f is maximised at $x = 2s/(\gamma \log N)$. Therefore,

$$\mathbb{E}[D_{s,m}] \leq f\left(\frac{2s}{\gamma \log N}\right) = \exp\left(-2s + s \log\left(\frac{4es}{\gamma^2(\log N)^2}\right)\right).$$

Since our assumptions imply that, see (63),

$$\frac{s}{(\log N)^2} \leq \frac{(\delta + \varepsilon)\mu}{(\log N)^2} \leq \frac{\delta + \varepsilon}{C^2},$$

we may conclude that $\mathbb{E}[D_{s,m}] \leq \exp(-(K+k)s)$, provided that C is sufficiently large. Therefore, if C is sufficiently large,

$$\mathbb{E}[D_s(\mathcal{H}_p)] = \sum_{m=k+1}^{ks} \mathbb{E}[D_{s,m}] \leq ks \cdot \exp(-Ks - ks) \leq \exp(-Ks).$$

This completes the proof.

8.4. Proof of Proposition 8.2. Let H be a connected, Δ -regular graph and let \mathcal{H} be the hypergraph on the vertex set $\binom{[n]}{2}$ whose edges are copies of H in K_n , so that $X = X_{n,p}^H = e_{\mathcal{H}_p}$. Let $\mu = \mathbb{E}[X]$, let $\eta = \eta(\varepsilon, \delta)$ be the constant from the statement of Proposition 8.3, and let $K = K(\varepsilon, \delta, \eta)$ be the constant from the statement of Lemma 8.5.

Since $(n/v_H)^{v_H} \leq \binom{n}{v_H} \leq e_{\mathcal{H}} \leq n^{v_H}$ for all large enough n , our assumption $Cn^{-1} \leq p^{\Delta/2} \leq C^{-1}n^{-1}(\log n)^{\frac{1}{v_H-2}}$ and the fact that $2e_H = \Delta v_H$ imply that

$$\left(\frac{C}{v_H}\right)^{v_H} \leq \mu \leq \frac{(\log n)^{1+\frac{2}{v_H-2}}}{C^{v_H}}, \quad (65)$$

and thus $\max\{1/\eta, K\} \leq \mu \leq \sqrt{e_{\mathcal{H}}}/K$ whenever C is sufficiently large. The claimed estimate on $-\log \mathbb{P}(X \geq (1+\delta)\mathbb{E}[X])$ will follow from Proposition 8.3 and Lemma 8.5 once we verify that $D_s(\mathcal{H}_p)$, the number of clusters of s copies of H in the random graph $G_{n,p}$, satisfies

$$\mathbb{E}[D_s(\mathcal{H}_p)] \leq \exp(-Ks)$$

for every s satisfying $2 \leq s \leq (\delta + \varepsilon)\mu$.

To this end, for every $s \geq 1$, every $k \geq 1$, and every $m \geq 1$, let $\mathcal{D}(s, k, m)$ denote the set of all clusters $\{\sigma_1, \dots, \sigma_s\}$ of s distinct copies of H in K_n such that the graph $\sigma_1 \cup \dots \cup \sigma_s$ has k vertices (of nonzero degree) and m edges. We further let $D_{s,k,m}$ denote the number of such clusters whose union is contained in $G_{n,p}$. When $s \geq 2$, the union of any s distinct copies of H contains between v_H and $v_H s$ vertices and between $e_H + 1$ and $e_H s$ edges, and thus $D_{s,k,m} = 0$ unless $v_H \leq k \leq v_H s$ and $e_H + 1 \leq m \leq e_H s$. We can therefore write

$$D_s(\mathcal{H}_p) = \sum_{k=v_H}^{v_H s} \sum_{m=e_H+1}^{e_H s} D_{s,k,m}.$$

Claim 8.7. *There exists a positive constant γ such that, for every $s \geq 2$, every $k \geq v_H$, and every $m \geq e_H + 1$,*

$$\mathbb{E}[D_{s,k,m}] \leq n^{-2\gamma m} \binom{k^2}{m} \binom{(2m)^{v_H/2}}{s}.$$

Proof. We first show that, for every $s \geq 1$, the set $\mathcal{D}(s, k, m)$ is empty unless

$$m - e_H \geq \left(\frac{\Delta}{2} + \frac{1}{2v_H}\right) \cdot (k - v_H). \quad (66)$$

We prove this fact by induction on s . The case $s = 1$ holds vacuously, as the set $\mathcal{D}(1, k, m)$ is nonempty only when $k = v_H$ and $m = e_H$. Assume now that $s \geq 2$ and let G be the union of copies of H that form some cluster in $\mathcal{D}(s, k, m)$. By definition, G has k vertices and m edges. Furthermore, for some $s' < s$, $k' \leq k$ and $m' < m$, there exist a subgraph $G' \subseteq G$ and a copy σ

of H in K_n that intersects (the edge set of) G' such that G' is the union of copies of H that form some cluster in $\mathcal{D}(s', k', m')$, and $G = G' \cup \sigma$. We note that s' may be strictly smaller than $s - 1$. Let $J \subseteq H$ be the subgraph of H that is isomorphic to $\sigma \cap G'$, so that $m = m' + e_H - e_J$ and $k = k' + v_H - v_J$. It follows from the inductive assumption that

$$\begin{aligned} m - e_H &= m' - e_J \geq \left(\frac{\Delta}{2} + \frac{1}{2v_H} \right) \cdot (k' - v_H) + e_H - e_J \\ &= \left(\frac{\Delta}{2} + \frac{1}{2v_H} \right) \cdot (k - v_H) - \left(\frac{\Delta}{2} + \frac{1}{2v_H} \right) \cdot (v_H - v_J) + e_H - e_J. \end{aligned}$$

We claim that the above inequality implies (66). This is obviously true when $J = H$. If J is a proper subgraph of H , then $2e_J \leq \Delta v_J - 1$, since H is connected and Δ -regular, and therefore

$$e_H - e_J \geq \frac{\Delta v_H}{2} - \frac{\Delta v_J - 1}{2} = \left(\frac{\Delta}{2} + \frac{1}{2(v_H - v_J)} \right) \cdot (v_H - v_J) \geq \left(\frac{\Delta}{2} + \frac{1}{2v_H} \right) \cdot (v_H - v_J),$$

which gives (66).

To complete the proof of the claim, note that

$$\mathbb{E}[D_{s,k,m}] \leq n^k p^m \cdot \binom{k^2}{m} \cdot \binom{N(H, k, m)}{s},$$

where $N(H, k, m)$ denotes the largest number of copies of H in a graph with k vertices and m edges. Inequality (66) implies that

$$n^k p^m = n^{v_H} p^{e_H} \cdot n^{k-v_H} p^{m-e_H} \leq n^{v_H} p^{e_H} \cdot \left(np^{\Delta/2+1/(2v_H)} \right)^{2v_H(m-e_H)/(\Delta v_H+1)}.$$

Since $p^{\Delta/2} \leq C^{-1}n^{-1} \log n$, the quantity $n^{v_H} p^{e_H}$ is only polylogarithmic in n and $np^{\Delta/2+1/(2v_H)} \leq n^{-\gamma'}$ for some positive constant γ' . As $m \geq e_H + 1$ and n is large, then there is a positive constant γ such that $n^k p^m \leq n^{-2\gamma m}$. The claimed upper bound on $\mathbb{E}[D_{s,k,m}]$ now follows from Theorem 5.4, which implies that $N(H, k, m) \leq (2m)^{v_H/2}$. \square

Claim 8.7 and the inequality $\binom{a}{b} \leq (ea/b)^b$ imply that

$$\begin{aligned} \mathbb{E}[D_{s,k,m}] &\leq \exp \left(-2\gamma m \log n + m \log(ek^2/m) + s \log(e(2m)^{v_H/2}/s) \right) \\ &\leq \exp \left(-\gamma m \log n + s \log(e(2m)^{v_H/2}/s) \right), \end{aligned}$$

where the second inequality as $k \leq v_H s \leq v_H(\delta + \varepsilon)\mu$ and μ is at most polylogarithmic in n , see (65). Let $f: (0, \infty) \rightarrow (0, \infty)$ be the function defined by

$$f(x) = \exp \left(-\gamma x \log n + s \log(e(2x)^{v_H/2}/s) \right),$$

so that $\mathbb{E}[D_{s,k,m}] \leq f(m)$. Elementary calculus shows that f is maximised at $x = sv_H/(2\gamma \log n)$. Therefore,

$$\mathbb{E}[D_{s,k,m}] \leq f \left(\frac{sv_H}{2\gamma \log n} \right) = \exp \left(-\frac{sv_H}{2} + \frac{sv_H}{2} \log \left(\frac{e^{2/v_H} v_H s^{1-2/v_H}}{\gamma \log n} \right) \right).$$

Since our assumptions imply that, see (65),

$$\frac{s^{1-2/v_H}}{\log n} \leq \frac{((\delta + \eta)\mu)^{1-2/v_H}}{\log n} \leq \frac{(\delta + \eta)^{1-2/v_H}}{C^{v_H-2}},$$

we may conclude that $\mathbb{E}[D_{s,k,m}] \leq \exp(-(K + v_H e_H)s)$, provided that C is sufficiently large. Therefore, if C is sufficiently large,

$$\mathbb{E}[D_s(\mathcal{H}_p)] = \sum_{k=v_H}^{v_H s} \sum_{m=e_H+1}^{e_H s} \mathbb{E}[D_{s,k,m}] \leq s^2 v_H e_H \cdot \exp(-Ks - v_H e_H s) \leq \exp(-Ks).$$

This completes the proof.

9. BEYOND POLYNOMIALS WITH NONNEGATIVE COEFFICIENTS

Although Theorem 3.1 applies only to the case where $X = X(Y)$ is a polynomial with nonnegative coefficients, the proof can be adapted to yield a similar result for all nonnegative functions $X: \{0, 1\}^N \rightarrow \mathbb{R}_{\geq 0}$. In this case, the degree assumption in Theorem 3.1 has to be replaced by a more general assumption on the ‘complexity’ of X .

Given an $I \subseteq \llbracket N \rrbracket$ and a $z \in \{0, 1\}^N$, we let

$$F(I, z) = \{y \in \{0, 1\}^N : y_i = z_i \text{ for all } i \in I\};$$

we call sets of this form *subcubes*. If F is a subcube, then there is a unique set I such that $F = F(I, z)$, for some z . We can thus define the *codimension* of F by $\text{codim } F = |I|$. Given a nonnegative function X on the hypercube, we define the *complexity* of X to be the smallest integer d for which it is possible to represent X as a linear combination with nonnegative coefficients of indicator functions of subcubes with codimension at most d . The complexity of X is well defined, and at most N , since $X = \sum_{z \in \{0, 1\}^N} X(z) \mathbb{1}_{F(\llbracket N \rrbracket, z)}$ is such a linear combination. Note also that the complexity of every polynomial with nonnegative coefficients and degree d is at most d .

Assume now that Y is a random variable taking values in $\{0, 1\}^N$ and that $X = X(Y)$. Given a subcube $F \subseteq \{0, 1\}^N$, we write $\mathbb{E}_F[X] = \mathbb{E}[X \mid Y \in F]$ for the expectation of X conditioned on $Y \in F$. We further define $\Phi_X: \mathbb{R}_{\geq 0} \rightarrow \mathbb{R}_{\geq 0} \cup \{\infty\}$ by

$$\Phi_X(\delta) = \min \left\{ -\log \mathbb{P}(Y \in F) : F \subseteq \{0, 1\}^N \text{ is a subcube with } \mathbb{E}_F[X] \geq (1 + \delta) \mathbb{E}[X] \right\}. \quad (67)$$

If X is an increasing function of Y , this definition coincides with our earlier definition of $\Phi_X(\delta)$, because then the minimum is achieved at a subcube of the form $F(I, \mathbf{1})$, where $\mathbf{1}$ is the N -dimensional all-ones vector. One may adapt the proof of Theorem 3.1 to show the following. (A precise proof of this theorem can be found in [23], which was written after this work was completed).

Theorem 9.1. *For every positive integer d and all positive real numbers ε and δ with $\varepsilon < 1/2$, there is a positive $K = K(d, \varepsilon, \delta)$ such that the following holds. Let Y be a sequence of N independent $\text{Ber}(p)$ random variables for some $p \in (0, 1 - \varepsilon]$ and assume that $X = X(Y)$ is non-negative and has complexity at most d and satisfies $\Phi_X(\delta - \varepsilon) \geq K \log(1/p)$. Denote by \mathcal{F}^* the collection of all subcubes $F \subseteq \{0, 1\}^N$ satisfying*

$$(F1) \quad \mathbb{E}_F[X] \geq (1 + \delta - \varepsilon) \mathbb{E}[X],$$

$$(F2) \quad \text{codim } F \leq K \cdot \Phi_X(\delta + \varepsilon).$$

Then,

$$\mathbb{P}(X \geq (1 + \delta) \mathbb{E}[X]) \leq (1 + \varepsilon) \cdot \mathbb{P}(Y \in F \text{ for some } F \in \mathcal{F}^*).$$

We note that this theorem does not exactly match Theorem 3.1; indeed, we only assert that the upper tail event is dominated by the appearance of a subcube in \mathcal{F}^* . It is possible to further restrict the family \mathcal{F}^* , analogous to the extraction of cores from seeds. We do not pursue this direction here.

Theorem 9.1 can be used to study the upper tail problem for induced subgraph counts. Suppose that H is a fixed graph and $X = X_{n,p}^{H\text{-ind}}$ is the number of induced copies of H in the random graph $G_{n,p}$. Let $N = \binom{n}{2}$ and, for an arbitrary bijection $\sigma_n: \binom{\llbracket n \rrbracket}{2} \rightarrow \llbracket N \rrbracket$, let Y_i be the indicator random variable of the event that $\sigma_n^{-1}(i)$ is an edge in $G_{n,p}$. Then we can write

$$X = \sum_{\substack{H' \subseteq K_n \\ H' \cong H}} \prod_{e \in E(H)} Y_{\sigma_n(e)} \prod_{e \in \binom{V(H)}{2} \setminus E(H)} (1 - Y_{\sigma_n(e)}).$$

In particular, the complexity of X is bounded by e_H and Theorem 9.1 applies. On the other hand, it is clear that X is generally not monotone, so one cannot use Theorem 3.1. This direction was successfully pursued by Cohen [23].

10. CONCLUDING REMARKS

Nonregular graphs. It is an open problem to extend Theorems 1.5 and 1.6 to nonregular graphs. It is straightforward to extend Theorem 1.6 to the more general case of strictly balanced graphs; however, note that Šileikis and Warnke [66] constructed balanced graphs for which the conclusion of Theorem 1.6 does not hold. In the localised regime, the results in [18, 24, 25, 30] apply to arbitrary (as opposed to only Δ -regular) graphs H ; however, these works require polynomially-suboptimal assumptions on the density p . Recently, Šileikis and Warnke [63] determined the order of the logarithmic upper tail probability for the number of copies of the star graph $K_{1,s}$ in $G_{n,p}$.

The phase transition between the Poisson and the localised regimes. We believe that the logarithmic upper tail probabilities of the random variables considered in this paper are always determined by either the Poisson behaviour, the localised behaviour, or the coexistence of the two (in the regime where they are commensurate). More precisely, we believe that for both $X = X_{n,p}^H$ (for a connected, Δ -regular H) and $X = X_{N,p}^{k\text{-AP}}$,

$$-\log \mathbb{P}(X \geq (1 + \delta) \mathbb{E}[X]) = (1 \pm o(1)) \cdot \min_{0 \leq \delta' \leq \delta} (\Phi_X(\delta') + \Psi_X(\delta - \delta')), \quad (68)$$

as long as $\mathbb{E}[X] \rightarrow \infty$ and $p \rightarrow 0$. Let $p^* = p^*(\delta, n)$ be such that $\Phi_X(\delta) = \Psi_X(\delta)$. Note that if $p \ll p^*$, then $\Psi_X(\delta) \ll \Phi_X(\delta)$ and we recover Theorems 1.4 and 1.6, whereas if $p \gg p^*$, we have $\Psi_X(\delta) \gg \Phi_X(\delta)$ and (68) implies (in some cases a stronger version of) Theorems 1.3, 1.5, and 1.7. If $p = \Theta(p^*)$, then both terms are of the same order and the conjecture allows for the upper tail to be dominated by configurations exhibiting features of both the Poisson and localised regimes.

Structural theorems for non-complete graphs. In the case where $X = X_{n,p}^H$ for a connected, Δ -regular graph H , we have neither determined the asymptotics of $\Phi_X(\delta)$ in the range $np^\Delta \rightarrow c \in (0, \infty)$ nor given a structural description of the upper tail event $\{X \geq (1 + \delta) \mathbb{E}[X]\}$ for any density p . Doing the former would yield the logarithmic upper tail probability of X , via Proposition 7.1; it is likely that the value of $\Phi_X(\delta)$ is given by a mixture of the clique construction and a ‘hub-like’ construction in which a constant number of vertices have degrees linear in n . As for the latter, the method used to prove Theorem 5.11 can be generalised to yield an analogous statement in which K_r is replaced with an arbitrary Δ -regular graph H . Armed with such a ‘stability’ statement, it is relatively straightforward to show that, when $np^\Delta \rightarrow 0$, the random graph $G_{n,p}$ conditioned on the upper tail event $\{X \geq (1 + \delta) \mathbb{E}[X]\}$ contains an ‘almost-clique’ of the ‘right’ size, as was the case when $H = K_r$. We were not able to prove such a structural statement in the complementary range $np^\Delta = \Omega(1)$.

Stability results for arithmetic progressions. An interesting problem is to characterise the near-minimisers of the optimisation problem for $\Phi_X(\delta)$ when $X = X_{N,p}^{k\text{-AP}}$. More precisely, we ask for a description of all subsets $I \subseteq \llbracket N \rrbracket$ that satisfy $\mathbb{E}_I[X] \geq (1 + \delta) \mathbb{E}[X]$ and $|I| \leq (1 + \varepsilon) \Phi_X(\delta + \varepsilon)$. As a consequence of Theorem 3.1 and the entropic stability of $X_{N,p}^{k\text{-AP}}$, which we established in the proof of Proposition 4.2, such a result would imply a structural characterisation of the upper tail event. Since the dominant contribution to the difference $\mathbb{E}_I[X] - \mathbb{E}[X]$ comes from k -term arithmetic progressions contained in I , this problem is equivalent to understanding the structure of sets $I \subseteq \mathbb{Z}$ that are near-maximisers of the number of k -term arithmetic progressions (among subsets of a given size). The structure of true maximisers was described by Green and Sisask [37] in the case for $k = 3$.

Decomposing the upper tail measure. Let \bar{Y} be the random variable obtained by conditioning Y on the upper tail event $\{X(Y) \geq (1 + \delta) \mathbb{E}[X]\}$ and let \tilde{Y} be the random variable obtained by first choosing a uniformly random solution I of the optimisation problem for $\Phi_X(\delta)$ and then conditioning Y on $\prod_{i \in I} Y_i = 1$. It would be very interesting to determine necessary and sufficient conditions so that \bar{Y} and \tilde{Y} are close in some metric. In particular, are the assumptions of

Theorem 3.1 sufficiently strong to imply this? This question is closely related to the more general problem of decomposing a Gibbs measure into a mixture of product measures. The work of Eldan and Gross [31], and, more recently, of Austin [4], gives general conditions for the existence of such a decomposition.

Moderate deviations. Throughout this paper, we have assumed that δ is a fixed, positive constant. It is interesting and natural to study the probability of $\{X \geq (1 + \delta) \mathbb{E}[X]\}$ when δ is allowed to depend on N and p . In the case where $\delta \mathbb{E}[X]$ is of the same order as $\sqrt{\text{Var}(X)}$, one can often prove a Central Limit Theorem, see [6, 7, 61]. The regime in which $\delta \mathbb{E}[X] \gg \sqrt{\text{Var}(X)}$ but $\delta \rightarrow 0$ is often referred to as the moderate deviation regime. One expects that, under reasonable assumptions, the logarithmic upper tail probability $-\log \mathbb{P}(X \geq (1 + \delta) \mathbb{E}[X])$ is of order $\min\{(\delta \mathbb{E}[X])^2 / \text{Var}(X), \Phi_X(\delta)\}$; this has been verified in certain cases, see [10, 33, 36, 38, 63, 69]. Our methods can be adapted to the moderate deviation regime. In an upcoming work [40], we calculate the logarithmic upper tail probability for $X = X_{N,p}^{k\text{-AP}}$ for nearly all pairs (p, δ) for which localisation is believed to occur—that is, when $\Phi_X(\delta) \ll (\delta \mathbb{E}[X])^2 / \text{Var}(X)$.

Other random graph models. Upper tails for subgraph counts have been considered in random graph models other than $G_{n,p}$, such as exponential random graphs [19], random geometric graphs [20], random regular graphs [39, 65], and (dense) uniform random graphs [27]. The framework developed here can be generalised to other (non-product) measures on the hypercube, providing a possible approach to such questions. It is likely that this requires adapting the notions of cores and entropic stability to the model.

Acknowledgements. We thank Bhaswar Bhattacharya, Asaf Cohen, Nicholas Cook, Ronen Eldan, Michael Krivelevich, Eyal Lubetzky, Matas Šileikis, Lutz Warnke, and Yufei Zhao for helpful comments and discussions. We are indebted to the three anonymous referees whose comments greatly improved the presentation of this paper.

REFERENCES

- [1] N. Alon. On the number of subgraphs of prescribed type of graphs with a given number of edges. *Israel J. Math.*, 38(1-2):116–130, 1981.
- [2] F. Augeri. A transportation approach to the mean-field approximation. [arXiv:1903.08021](https://arxiv.org/abs/1903.08021).
- [3] F. Augeri. Nonlinear large deviation bounds with applications to Wigner matrices and sparse Erdős–Rényi graphs. *Ann. Probab.*, 48(5):2404–2448, 2020.
- [4] T. Austin. The structure of low-complexity Gibbs measures on product spaces. *Ann. Probab.*, 47(6):4002–4023, 2019.
- [5] J. Balogh, R. Morris, W. Samotij, and L. Warnke. The typical structure of sparse K_{r+1} -free graphs. *Trans. Amer. Math. Soc.*, 368(9):6439–6485, 2016.
- [6] A. D. Barbour, M. Karoński, and A. Ruciński. A central limit theorem for decomposable random variables with applications to random graphs. *J. Combin. Theory Ser. B*, 47(2):125–145, 1989.
- [7] Y. Barhoumi-Andréani, C. Koch, and H. Liu. Bivariate fluctuations for the number of arithmetic progressions in random sets. *Electron. J. Probab.*, 24:Paper No. 145, 32, 2019.
- [8] A. Basak and R. Basu. Upper tail large deviations of regular subgraph counts in Erdős–Rényi graphs in the full localized regime. [arXiv:1912.11410](https://arxiv.org/abs/1912.11410).
- [9] B. B. Bhattacharya, S. Ganguly, E. Lubetzky, and Y. Zhao. Upper tails and independence polynomials in random graphs. *Adv. Math.*, 319:313–347, 2017.
- [10] B. B. Bhattacharya, S. Ganguly, X. Shao, and Y. Zhao. Upper tail large deviations for arithmetic progressions in a random set. *Int. Math. Res. Not. IMRN*, (1):167–213, 2020.
- [11] B. Bollobás. Threshold functions for small subgraphs. *Math. Proc. Camb. Philos. Soc.*, 90(2):197–206, 1981.
- [12] B. Bollobás. *Random graphs*. Cambridge University Press, Cambridge, 2nd edition, 2001.
- [13] J. A. Bondy and U. S. R. Murty. *Graph theory*, volume 244 of *Graduate Texts in Mathematics*. Springer, New York, 2008.
- [14] S. Boucheron, G. Lugosi, and P. Massart. *Concentration inequalities*. Oxford University Press, Oxford, 2013.
- [15] J. Briët and S. Gopi. Gaussian width bounds with applications to arithmetic progressions in random settings. *Int. Math. Res. Not. IMRN*, (22):8673–8696, 2020.

- [16] S. Chatterjee. The missing log in large deviations for triangle counts. *Random Structures Algorithms*, 40(4):437–451, 2012.
- [17] S. Chatterjee. *Large deviations for random graphs*, volume 2197 of *Lecture Notes in Mathematics*. Springer, Cham, 2017. Lecture notes from the 45th Probability Summer School held in Saint-Flour (École d’Été de Probabilités de Saint-Flour), June 2015.
- [18] S. Chatterjee and A. Dembo. Nonlinear large deviations. *Adv. Math.*, 299:396–450, 2016.
- [19] S. Chatterjee and P. Diaconis. Estimating and understanding exponential random graph models. *Ann. Statist.*, 41(5):2428–2461, 2013.
- [20] S. Chatterjee and M. Harel. Localization in random geometric graphs with too many edges. *Ann. Probab.*, 48(2):574–621, 2020.
- [21] S. Chatterjee and S. R. S. Varadhan. The large deviation principle for the Erdős–Rényi random graph. *European J. Combin.*, 32(7):1000–1017, 2011.
- [22] F. R. K. Chung, R. L. Graham, P. Frankl, and J. B. Shearer. Some intersection theorems for ordered sets and graphs. *J. Combin. Theory Ser. A*, 43:23–37, 1986.
- [23] A. Cohen. The upper tail problem for induced 4-cycles in sparse random graphs. M.Sc. thesis, 2021.
- [24] N. Cook and A. Dembo. Large deviations of subgraph counts for sparse Erdős–Rényi graphs. *Adv. Math.*, 373:107289, 53, 2020.
- [25] N. A. Cook, A. Dembo, and H. T. Pham. Regularity method and large deviation principles for the Erdős–Rényi hypergraph. [arXiv:2102.09100](https://arxiv.org/abs/2102.09100).
- [26] R. DeMarco and J. Kahn. Tight upper tail bounds for cliques. *Random Structures Algorithms*, 41(4):469–487, 2012.
- [27] A. Dembo and E. Lubetzky. A large deviation principle for the Erdős–Rényi uniform random graph. *Electron. Commun. Probab.*, 23, 2018.
- [28] A. Dembo and O. Zeitouni. *Large deviations techniques and applications*, volume 38 of *Applications of Mathematics*. Springer, New York, 2nd edition, 1998.
- [29] R. Diestel. *Graph theory*, volume 173 of *Graduate Texts in Mathematics*. Springer, Berlin, 5th edition, 2017.
- [30] R. Eldan. Gaussian-width gradient complexity, reverse log-Sobolev inequalities and nonlinear large deviations. *Geom. Funct. Anal.*, 28(6):1548–1596, 2018.
- [31] R. Eldan and R. Gross. Decomposition of mean-field Gibbs distributions into product measures. *Electron. J. Probab.*, 23, 2018.
- [32] P. Erdős. On the number of complete subgraphs contained in certain graphs. *Magyar Tud. Akad. Mat. Kutató Int. Közl.*, 7:459–464, 1962.
- [33] G. Fiz Pontiveros, S. Griffiths, M. Secco, and O. Serra. Deviation probabilities for arithmetic progressions and other regular discrete structures. [arXiv:1910.12835](https://arxiv.org/abs/1910.12835).
- [34] E. Friedgut and J. Kahn. On the number of copies of one hypergraph in another. *Israel J. Math.*, 105:251–256, 1998.
- [35] D. Galvin. Three tutorial lectures on entropy and counting. [arXiv:1406.7872](https://arxiv.org/abs/1406.7872).
- [36] C. Goldschmidt, S. Griffiths, and A. Scott. Moderate deviations of subgraph counts in the Erdős–Rényi random graphs $G(n, m)$ and $G(n, p)$. *Trans. Amer. Math. Soc.*, 373(8):5517–5585, 2020.
- [37] B. Green and O. Sisask. On the maximal number of 3-term arithmetic progressions in subsets of $\mathbb{Z}/p\mathbb{Z}$. *Bulletin London Math. Soc.*, 40(6):945–955, 2008.
- [38] S. Griffiths, C. Koch, and M. Secco. Deviation probabilities for arithmetic progressions and irregular discrete structures. [arXiv:2012.09280](https://arxiv.org/abs/2012.09280).
- [39] B. Gunby. Upper tails of subgraph counts in sparse regular graphs. [arXiv:2010.00658](https://arxiv.org/abs/2010.00658).
- [40] M. Harel, F. Mousset, and W. Samotij. Upper tails for arithmetic progressions in the moderate deviation regime. Manuscript.
- [41] T. E. Harris. A lower bound for the critical probability in a certain percolation process. *Proc. Cambridge Philos. Soc.*, 56:13–20, 1960.
- [42] W. Hoeffding. Probability inequalities for sums of bounded random variables. *J. Amer. Statist. Assoc.*, 58(301):13–30, 1963.
- [43] S. Janson. Poisson approximation for large deviations. *Random Structures Algorithms*, 1(2):221–229, 1990.
- [44] S. Janson, K. Oleszkiewicz, and A. Ruciński. Upper tails for subgraph counts in random graphs. *Israel J. Math.*, 142:61–92, 2004.
- [45] S. Janson and A. Ruciński. The infamous upper tail. *Random Structures Algorithms*, 20(3):317–342, 2002.
- [46] S. Janson and A. Ruciński. The deletion method for upper tail estimates. *Combinatorica*, 24(4):615–640, 2004.
- [47] S. Janson and A. Ruciński. Upper tails for counting objects in randomly induced subhypergraphs and rooted random graphs. *Ark. Mat.*, 49(1):79–96, 2011.
- [48] S. Janson and L. Warnke. The lower tail: Poisson approximation revisited. *Random Structures Algorithms*, 48(2):219–246, 2016.

- [49] M. Karoński and A. Ruciński. On the number of strictly balanced subgraphs of a random graph. In *Graph Theory, Łągow 1981*, volume 1018 of *Lecture Notes in Mathematics*, pages 79–83. Springer, 1983.
- [50] G. Katona. A theorem of finite sets. In *Proc. Colloq. Theory of Graphs, Tihany 1966*, pages 187–207. Academic Press, New York, 1966.
- [51] P. Keevash. Shadows and intersections: Stability and new proofs. *Adv. Math.*, 218:1685–1703, 2008.
- [52] J. H. Kim and V. H. Vu. Concentration of multivariate polynomials and its applications. *Combinatorica*, 20(3):417–434, 2000.
- [53] J. H. Kim and V. H. Vu. Divide and conquer martingales and the number of triangles in a random graph. *Random Structures Algorithms*, 24(2):166–174, 2004.
- [54] G. Kozma and W. Samotij. Lower tails via relative entropy. [arXiv:2104.04850](https://arxiv.org/abs/2104.04850).
- [55] J. B. Kruskal. The number of simplices in a complex. In *Mathematical Optimization Techniques*, volume 10, pages 251–278. Univ. of California Press, Berkeley, CA, 1963.
- [56] E. Lubetzky and Y. Zhao. On replica symmetry of large deviations in random graphs. *Random Structures Algorithms*, 47(1):109–146, 2015.
- [57] E. Lubetzky and Y. Zhao. On the variational problem for upper tails in sparse random graphs. *Random Structures Algorithms*, 50(3):420–436, 2017.
- [58] F. Mousset, A. Noever, K. Panagiotou, and W. Samotij. On the probability of nonexistence in binomial subsets. *Ann. Probab.*, 48(1):493–525, 2020.
- [59] S. Mukherjee and B. B. Bhattacharya. Replica symmetry in upper tails of mean-field hypergraphs. *Adv. in Appl. Math.*, 119:102047, 25, 2020.
- [60] I. Rivin. Counting cycles and finite dimensional L^p norms. *Adv. in Appl. Math.*, 29(4):647–662, 2002.
- [61] A. Ruciński. When are small subgraphs of a random graph normally distributed? *Probab. Theory Related Fields*, 78(1):1–10, 1988.
- [62] M. Šileikis. On the upper tail of counts of strictly balanced subgraphs. *Electron. J. Combin.*, 19(1):4, 2012.
- [63] M. Šileikis and L. Warnke. Upper tail bounds for stars. *Electron. J. Combin.*, 27(1):Paper No. 1.67, 23, 2020.
- [64] M. Talagrand. Concentration of measure and isoperimetric inequalities in product spaces. *Publ. Math. Inst. Hautes Études Sci.*, 81(1):73–205, 1995.
- [65] P. van der Hoorn, G. Lippner, and E. Mossel. Regular graphs with linearly many triangles. [arXiv:1904.02212](https://arxiv.org/abs/1904.02212).
- [66] M. Šileikis and L. Warnke. A counterexample to the DeMarco-Kahn upper tail conjecture. *Random Structures Algorithms*, 55(4):775–794, 2019.
- [67] V. H. Vu. On the concentration of multivariate polynomials with small expectation. *Random Structures Algorithms*, 16(4):344–363, 2000.
- [68] V. H. Vu. Concentration of non-Lipschitz functions and applications. *Random Structures Algorithms*, 20(3):262–316, 2002.
- [69] L. Warnke. Upper tails for arithmetic progressions in random subsets. *Israel J. Math.*, 221(1):317–365, 2017.
- [70] L. Warnke. On the missing log in upper tail estimates. *J. Combin. Theory Ser. B*, 140:98–146, 2020.
- [71] Y. Zhao. On the lower tail variational problem for random graphs. *Combin. Probab. Comput.*, 26(2):301–320, 2017.

MATAN HAREL, DEPARTMENT OF MATHEMATICS, NORTHEASTERN UNIVERSITY, BOSTON, MA 02115, USA
Email address: m.harel@northeastern.edu

FRANK MOUSSET, SCHOOL OF MATHEMATICAL SCIENCES, TEL AVIV UNIVERSITY, TEL AVIV 6997801, ISRAEL
Email address: moussetfrank@gmail.com

WOJCIECH SAMOTIJ, SCHOOL OF MATHEMATICAL SCIENCES, TEL AVIV UNIVERSITY, TEL AVIV 6997801, ISRAEL
Email address: samotij@tauex.tau.ac.il