Counting extensions revisited

Matas Šileikis* and Lutz Warnke†

7 November 2019; revised 28 July 2021

Abstract

We consider rooted subgraphs in random graphs, i.e., extension counts such as (i) the number of triangles containing a given vertex or (ii) the number of paths of length three connecting two given vertices. In 1989, Spencer gave sufficient conditions for the event that, with high probability, these extension counts are asymptotically equal for all choices of the root vertices. For the important strictly balanced case, Spencer also raised the fundamental question as to whether these conditions are necessary. We answer this question by a careful second moment argument, and discuss some intriguing problems that remain open.

1 Introduction

Subgraph counts and their many natural generalizations are central topics in random graph theory: since the 1960's they are a constant source of beautiful problems and conjectures, which have repeatedly inspired the development of important new probabilistic techniques and insights (see [7, 1, 15, 12]).

In this paper we consider rooted subgraph counts in the binomial random graph $\mathbb{G}_{n,p}$, i.e., so-called extension counts [27, 31, 20, 35] such as (i) the number of triangles containing a given vertex or (ii) the number of paths of length three connecting two given vertices. In combinatorics and related areas, the need for studying such extension counts arises frequently in probabilistic proofs and applications, including zero-one laws in random graphs [27, 20, 32], games on random graphs [19, 22], random graph processes [4, 3, 5, 11, 6], sparse random analogues of classical extremal and Ramsey results [23, 26, 2], and many more, such as [30, 24, 35, 33, 39, 17, 34, 21, 36]. Consequently the investigation of extension counts is not only a natural problem in probabilistic combinatorics, but also an important issue from the applications point of view.

After initial groundwork of Shelah and Spencer [27] as well as Spencer [30] on (rooted subgraph) extension counts, in 1989 Spencer [31] proved sufficient conditions for the event that, with high probability[1], these extension counts are asymptotically equal in $\mathbb{G}_{n,p}$ for all choices of the root vertices. For the important strictly balanced case, he also raised the fundamental question whether these sufficient conditions (see (3) below) are qualitatively necessary. In this paper we answer Spencer's 30-year old question by a careful second moment argument (see Theorem 1 below), rectifying a surprising gap in the random graph literature. We also discuss some further partial results and intriguing open problems (see Sections 1.2–1.3 below).

1.1 Main result

To fix notation, by a *rooted graph* (G, H) we mean a graph $H = (V(H), E(H))$ and an induced subgraph $G \subseteq H$ with labeled 'root' vertices $V(G) = \{1, \dots, v_G\}$. Given a tuple $\mathbf{x} = (x_1, \dots, x_{v_G})$ of distinct vertices from some 'host' graph, a (G, H)-*extension* of $\mathbf{x}$ is a copy of the graph $H_G := (V(H), E(H) \setminus E(G))$ in which each vertex $j \in V(G)$ is mapped onto x_j. Note that if $\mathbf{x}$ spans a copy of G in the host graph (i.e., if the function $j \mapsto x_j$ maps edges of G to edges in the host graph), then every (G, H)-extension of $\mathbf{x}$ corresponds to a copy of H. Since the edges between root vertices do not affect the definition of a (G, H)-extension, the

*Institute of Computer Science of the Czech Academy of Sciences, Pod Vodárenskou věží 2, 182 07 Prague, Czech Republic. E-mail: matas.sileikis@gmail.com. With institutional support RVO:67985807. Research supported by the Czech Science Foundation, grant number GA19-08740S.

†School of Mathematics, Georgia Institute of Technology, Atlanta GA 30332, USA. E-mail: warnke@math.gatech.edu. Research partially supported by NSF Grant DMS-1703516 and a Sloan Research Fellowship.

[1] As usual, we say that an event holds *whp* (with high probability) if it holds with probability tending to 1 as $n \to \infty$.

reader may without loss of generality assume that $V(G)$ is an independent set of H in the results below, cf. [15, 17] (allowing for G that are not independent will be convenient in some proofs, though). For brevity, we write $[n]_{v_G}$ for the set of all *roots*, i.e., tuples $\mathbf{x} = (x_1, \ldots, x_{v_G})$ of distinct vertices from $[n] := \{1, \ldots, n\}$. Let $X_{\mathbf{x}} = X_{G,H}(\mathbf{x})$ denote the number of (G, H)-extensions of $\mathbf{x}$ in the binomial random graph $\mathbb{G}_{n,p}$. Note that the expected value

$$\mu = \mu_{G,H} := \mathbb{E}X_{\mathbf{x}} \asymp n^{v_H - v_G} p^{e_H - e_G} \tag{1}$$

does not depend[2] on the particular choice of $\mathbf{x}$. To avoid trivialities, we henceforth assume that H has more edges than G, i.e., that $e_H > e_G$. Extending the standard density notation for unrooted subgraphs, we define

$$m(G, H) := \max_{G \subsetneq J \subseteq H} d(G, J) \quad \text{with} \quad d(G, J) := \frac{e_J - e_G}{v_J - v_G}, \tag{2}$$

and say that (G, H) is *strictly balanced* if $d(G, J) < d(G, H)$ for all $G \subsetneq J \subsetneq H$. We also call (G, H) *grounded* if at least one root vertex $j \in V(G)$ is connected to a non-root vertex $w \in V(H) \setminus V(G)$.

Spencer derived in 1989 sufficient conditions for the event that, with high probability, all extension counts satisfy $X_{\mathbf{x}} \sim \mu$, i.e., are asymptotically equal. In the important case when (G, H) is strictly balanced, [31, Theorem 2] states that for every fixed $\varepsilon \in (0, 1]$ there is a constant $K(\varepsilon) > 0$ such that

$$\lim_{n\to\infty} \mathbb{P}\Big(\max_{\mathbf{x} \in [n]_{v_G}} |X_{\mathbf{x}} - \mu| < \varepsilon\mu \Big) = 1 \quad \text{if } \mu \geqslant K(\varepsilon) \log n. \tag{3}$$

Spencer remarked that his constant satisfies $K(\varepsilon) \to \infty$ as $\varepsilon \to 0$, and speculated that this is probably also necessary, see [31, Remark on p.249]. In other words, he raised the question whether his sufficient condition is qualitatively best possible.

Our main result answers this fundamental question: (4) shows that the 'correct' dependence is $K(\varepsilon) = \Theta(\varepsilon^{-2})$ in the grounded case, even when $\varepsilon = \varepsilon(n) \to 0$ at some polynomial rate. For completeness, (5) also shows that the logarithm in the sufficient condition (3) is unnecessary in the less interesting ungrounded case (where extension counts are essentially unrooted subgraph counts, cf. example (b) in Figure 1).

Theorem 1 (Main result: strictly balanced case)**.** *Let (G, H) be a rooted graph that is strictly balanced. There are constants $c, C, \alpha > 0$ such that, for all $p = p(n) \in [0, 1]$ and $\varepsilon = \varepsilon(n) \in [n^{-\alpha}, 1]$, the following holds:*

(i) *If the rooted graph (G, H) is grounded, then*

$$\lim_{n\to\infty} \mathbb{P}\Big(\max_{\mathbf{x} \in [n]_{v_G}} |X_{\mathbf{x}} - \mu| < \varepsilon\mu \Big) = \begin{cases} 0 & \text{if } \varepsilon^2\mu \leqslant c \log n, \\ 1 & \text{if } \varepsilon^2\mu \geqslant C \log n. \end{cases} \tag{4}$$

(ii) *If the rooted graph (G, H) is not grounded, then*

$$\lim_{n\to\infty} \mathbb{P}\Big(\max_{\mathbf{x} \in [n]_{v_G}} |X_{\mathbf{x}} - \mu| < \varepsilon\mu \Big) = \begin{cases} 0 & \text{if } \varepsilon^2\mu \to 0, \\ 1 & \text{if } \varepsilon^2\mu \to \infty. \end{cases} \tag{5}$$

In concrete words, (4)–(5) of Theorem 1 give thresholds for the concentration of extension counts in terms of $\varepsilon^2\mu$, similar to the thresholds in terms of the edge probability p that are well-known for many properties of $\mathbb{G}_{n,p}$. The role of the expression $\varepsilon^2\mu$ in (4)–(5) can be made plausible by pretending that $X_{\mathbf{x}}$ behaves like a binomial random variable with expectation μ (the actual behaviour is of course more involved), in which case Chernoff-type tail bounds of the form $\mathbb{P}(|X_{\mathbf{x}} - \mu| \geqslant \varepsilon\mu) \leqslant e^{-\Omega(\varepsilon^2\mu)}$ hold. Indeed, considering the union bound over the $\Theta(n^{v_G})$ roots $\mathbf{x}$, it then seems plausible that the 1-statement follows when $\varepsilon^2\mu$ is at least a large enough multiple of $\log n$. An intuitive reason why the $\log n$ factor is absent in the ungrounded threshold (5) is that here the $X_{\mathbf{x}}$ are strongly correlated and in fact almost equal (e.g., in example (b) from Figure 1 each $X_{\mathbf{x}}$ is well-approximated by the total number of triangles), so there should be no need to use a union bound.

The main contribution of Theorem 1 is the 0-statement in the grounded threshold (4), which was missing in previous work: our proof uses a careful second moment argument (combining correlation inequalities and

[2] Here $a_n \asymp b_n$ is a convenient shorthand for $a_n = \Theta(b_n)$, following standard asymptotic notation as in [15, p. 9].

Figure 1: Examples of rooted graphs, with the root vertex circled and primal subgraphs marked in bold: (a) strictly balanced and grounded, (b) strictly balanced and not grounded, (c) with a unique primal that is grounded, and (d) with a unique primal that is not grounded. Our main result Theorem 1 applies to (a),(b), Theorem 2 applies to (a),(c), Theorem 3 applies to (b),(d), and Theorem 4 applies to all of them.

counting arguments with Janson's inequality) in order to establish that, with high probability, there exists a root $\mathbf{x}$ with $X_{\mathbf{x}} \geqslant (1+\varepsilon)\mu$, i.e., with too many (G,H)-extensions. This is closely related to the task of obtaining good lower bounds on $\mathbb{P}(X_{\mathbf{x}} \geqslant (1+\varepsilon)\mu)$, which are not so well understood as upper bounds; see [16, 18, 9, 29]. To sidestep this conceptual obstacle, in Section 3 we therefore work with (easier to estimate) auxiliary events that enforce $X_{\mathbf{x}} \geqslant (1+\varepsilon)\mu$ via 'disjoint' extensions, and we believe that our approach might also be useful for establishing 'lower bounds' in other problems.

1.2 Partial results: beyond the strictly balanced case

We also establish some threshold results for extension counts of rooted graphs (G,H) that are not necessarily strictly balanced. Here things are more complicated, since we now need to take into account all subgraphs $J \subseteq H$ containing the root G, in particular those that satisfy $d(G,J) = m(G,H)$; cf. [30, 31, 24, 15]. We call such subgraphs J *primal*, and for brevity also say that J is *grounded* if (G,J) is grounded. The partial results Theorems 2–3 below cover all strictly balanced (G,H), and they in particular imply that Theorem 1 also holds with $\varepsilon^2\Phi$ instead of $\varepsilon^2\mu$ (possibly after modifying the constants c, C, α), where

$$\Phi = \Phi_{G,H} := \min_{G \subseteq J \subseteq H : e_J > e_G} \mu_{G,J}. \tag{6}$$

There is no contradiction here: the extra assumption $\varepsilon \geqslant n^{-\alpha}$ ensures that the conclusions of the 0- and 1-statements of Theorem 1 coincide regardless of whether we use $\varepsilon^2\Phi$ or $\varepsilon^2\mu$ (cf. Section 5.2). It thus comes as no surprise that in our main result Theorem 1 the technical assumption $\varepsilon \geqslant n^{-\alpha}$ is indeed[3] necessary.

The following result covers the case where (G,H) has only one primal subgraph which also happens to be grounded, such as in examples (a) and (c) from Figure 1; this case includes the rooted graphs in Theorem 1 (i) since in that case H is a unique primal subgraph.

Theorem 2 (Unique and grounded primal case). *Let (G,H) be a rooted graph with a unique primal subgraph J. If (G,J) is grounded, then there are constants $c, C, \alpha > 0$ such that, for all $p = p(n) \in [0,1]$ and $\varepsilon = \varepsilon(n) \in [n^{-\alpha}, 1]$,*

$$\lim_{n\to\infty} \mathbb{P}\Big(\max_{\mathbf{x}\in[n]_{v_G}} |X_{\mathbf{x}} - \mu| < \varepsilon\mu\Big) = \begin{cases} 0 & \text{if } \varepsilon^2\Phi \leqslant c\log n, \\ 1 & \text{if } \varepsilon^2\Phi \geqslant C\log n. \end{cases} \tag{7}$$

The heuristic idea is that the main contribution to deviations of $X_{\mathbf{x}} = X_{G,H}(\mathbf{x})$ comes from those of $X_{G,J}(\mathbf{x})$, and, since (G,J) is strictly balanced and grounded, the problem thus intuitively reduces to Theorem 1 (i).

The following result covers the case where no primal subgraph of (G,H) is grounded, such as in examples (b) and (d) from Figure 1; this case includes the rooted graphs in Theorem 1 (ii).

Theorem 3 (No grounded primals case). *Let (G,H) be a rooted graph with no grounded primal subgraphs. There is a constant $\alpha > 0$ such that, for all $p = p(n) \in [0,1]$ and $\varepsilon = \varepsilon(n) \in [n^{-\alpha}, 1]$,*

$$\lim_{n\to\infty} \mathbb{P}\Big(\max_{\mathbf{x}\in[n]_{v_G}} |X_{\mathbf{x}} - \mu| < \varepsilon\mu\Big) = \begin{cases} 0 & \text{if } \varepsilon^2\Phi \to 0, \\ 1 & \text{if } \varepsilon^2\Phi \to \infty. \end{cases} \tag{8}$$

[3]For examples (a) and (b) from Figure 1 with $\varepsilon \asymp n^{-1/2}$ and $\varepsilon \asymp n^{-1}$, when $p \asymp n^{-1/4}$ it is routine to check that $\Phi \to \infty$, $\varepsilon^2\Phi \to 0$ and $\varepsilon^2\mu \gg \log n$ in both cases. Hence the 0-statement holds by (9) of Theorem 4, showing that (4)–(5) of Theorem 1 fail.

Figure 2: The rooted graphs used in Propositions 5–6, with the root vertex circled: for (e) Spencer's general 1-statement is not optimal, and for (f) the natural condition $\varepsilon^2\Phi \gg \log n$ does not imply the 1-statement.

Similar to Theorem 1 (ii), the intuition is that all $X_{\mathbf{x}}$ are approximately equal once we know the number of unrooted copies of a certain subgraph of H (e.g., in example (d) from Figure 1 this special subgraph is K_4).

Theorems 2–3 give thresholds for the concentration of extension counts in terms of $\varepsilon^2\Phi$. For general (G, H) we do not have such a threshold, but the following result intuitively states that the transition from the 0-statement to the 1-statement always happens at some point as $\varepsilon^2\Phi$ changes from $o(1)$ to $n^{\Omega(1)}$.

Theorem 4 (General case: approximate conditions)**.** *Let (G, H) be a rooted graph. For all $p = p(n) \in [0, 1]$ and $\varepsilon = \varepsilon(n) \in (0, 1]$ with $1 - p = \Omega(1)$ and $\Phi \to \infty$,*

$$\lim_{n\to\infty} \mathbb{P}\Big(\max_{\mathbf{x}\in[n]_{v_G}} |X_{\mathbf{x}} - \mu| < \varepsilon\mu\Big) = \begin{cases} 0 & \text{if } \varepsilon^2\Phi \to 0, \\ 1 & \text{if } \varepsilon^2\Phi = n^{\Omega(1)}. \end{cases} \tag{9}$$

The 1-statement in (9) implies [31, Corollary 4], which in turn strengthens a result that played a key role in the study of zero-one laws [27] due to Shelah and Spencer (since the 'safe' assumptions from [31, 27] imply $\Phi = n^{\Omega(1)}$ via Remark 1 (iv) from Section 2).

1.3 Discussion: open problems and cautionary examples

For rooted subgraph extension counts, the main open problem is to fully determine the thresholds for concentration, i.e., to close the gap in (9) of Theorem 4 (and to weaken the conditions of Theorems 1–3).

Problem 1. *Determine the 'correct' conditions for the* 0- *and* 1-*statements of any rooted graph (G, H).*

Our understanding of Problem 1 is still far from satisfactory. Indeed, even for fixed $\varepsilon \in (0, 1]$ the correct 1-statement condition remains open, which we now illustrate for the rooted graph (e) from Figure 2. In this case, any (G, H)-extension can be viewed as a combination of a (G, K_4)-extension and a (K_4, H)-extension. The proof of Spencer's general 1-statement result [31, Theorem 3] combines this decomposition with his strictly balanced result (3) for (G, K_4) and (K_4, H), leading to a sufficient condition of form $\min\{\mu_{G,K_4}, \mu_{K_4,H}\} \geqslant K'(\varepsilon)\log n$ (cf. [31, Section 2]). The following result shows that this sufficient condition can be weakened in some range, demonstrating that Spencer's general 1-statement condition is not always optimal.

Proposition 5. *Let (G, H) be the rooted graph (e) depicted in Figure 2. Set $\omega := np^2$. For all $p = p(n) \in [0, 1]$ and $\varepsilon = \varepsilon(n) \in (0, 1]$ such that $\omega \ll \log n$ and $\varepsilon^2\omega^3 \gg \log n$, we have $\varepsilon^2\mu_{G,K_4} \gg \log n \gg \varepsilon^2\mu_{K_4,H}$ but $\mathbb{P}(\max_{\mathbf{x}\in[n]_{v_G}} |X_{\mathbf{x}} - \mu| < \varepsilon\mu) \to 1$ as $n \to \infty$.*

It is not hard to see that in the setting of Proposition 5 we have $\varepsilon^2\Phi \asymp \varepsilon^2\mu_{G,K_4} \gg \log n$, which together with Theorems 2–3 suggests that maybe $\varepsilon^2\Phi \gg \log n$ is always a sufficient condition[4] for the 1-statement (which would sharpen Theorem 4). However, the following result shows that this speculation is false for the rooted graph (f) depicted in Figure 2, indicating that Problem 1 is more tricky than one might think.

Proposition 6. *Let (G, H) be the rooted graph (f) depicted in Figure 2. Set $\omega := np^2$. For all $p = p(n) \in [0, 1]$ and $\varepsilon = \varepsilon(n) \in (0, 1]$ such that $\omega \ll (\log n)^{0.39}$ and $\varepsilon^2\omega^3 \gg \log n$, we have $\varepsilon^2\Phi \asymp \varepsilon^2\mu_{G,K_4} \gg \log n$ but $\mathbb{P}(\max_{\mathbf{x}\in[n]_{v_G}} |X_{\mathbf{x}} - \mu| < \varepsilon\mu) \to 0$ as $n \to \infty$.*

[4]Further support comes from the fact that $X_{\mathbf{x}}$ is asymptotically normal, see Claim 17 (ii) in Appendix A and the variance estimate (10) from Section 2, which makes it plausible that $\mathbb{P}(|X_{\mathbf{x}} - \mu| \geqslant \varepsilon\mu) \leqslant e^{-\Omega((\varepsilon\mu)^2/\operatorname{Var} X_{\mathbf{x}})} \leqslant e^{-\Omega(\varepsilon^2\Phi)} \ll n^{-v_G}$ holds, which in turn would then establish the 1-statement by taking the union bound over all $\Theta(n^{v_G})$ roots $\mathbf{x}$.

Overall, we hope that the above intriguing examples and open problems will stimulate more research into rooted subgraph counts. When (G, H) is strictly balanced and grounded, then we conjecture that (7) holds for suitable $c, C > 0$ under the natural assumptions $\mu \to \infty$ and $1-p = \Omega(1)$, i.e., without assuming $\varepsilon \geqslant n^{-\alpha}$. We leave it as an open problem to formulate a conjecture for the general solution to Problem 1, which in many cases is closely related to determining the regime where $\mathbb{P}(|X_{\mathbf{x}} - \mu| \geqslant \varepsilon\mu)$ changes from $n^{-o(1)}$ to $n^{-\omega(1)}$, say. In the concluding remarks we also discuss a potential connection to extreme value theory (see Section 7).

1.4 Organization of the paper

In Section 2 we introduce some auxiliary results, which also imply Theorem 4. In Section 3 we prove our main result Theorem 1 (i) for strictly balanced (G, H) that are grounded. In Sections 4 and 5.1 we prove Theorems 2 and 3, i.e., cover the case where no grounded primal of (G, H) exists, and the case where the primal of (G, H) is unique and grounded, respectively. In Section 5.2, we prove Theorem 1 (ii) for strictly balanced (G, H) that are not grounded. In Section 6 we prove the cautionary examples from Propositions 5–6. Finally, Section 7 contains some concluding remarks and problems.

2 Preliminaries

In this section we collect some useful basic observations, and a partial result which implies Theorem 4. First, by adapting the textbook argument [15, Lemma 3.5] for (unrooted) subgraph counts, for any rooted graph (G, H) it is standard to see that the variance of $X_{G,H}(\mathbf{x})$ satisfies

$$\sigma^2 = \sigma^2_{G,H} := \operatorname{Var} X_{G,H}(\mathbf{x}) \asymp (1-p)\mu^2_{G,H}/\Phi_{G,H} \tag{10}$$

for any edge probability $p = p(n) \in (0, 1]$, where $\mu = \mu_{G,H}$ and $\Phi = \Phi_{G,H}$ are as defined in (1) and (6); cf. [28]. Next, inspired by similar statements for subgraph counts [15, Lemma 3.6], using the relation $\mu_{G,J} \asymp (n^{1/d(G,J)}p)^{e_J - e_G}$ for all $G \subseteq J \subseteq H$ with $e_J > e_G$, it is straightforward to establish the following useful properties. Recall that $m(G, H)$ and $\Phi = \Phi_{G,H}$ are defined in (2) and (6), respectively.

Remark 1. *For any rooted graph (G, H), the following hold for all $p = p(n) \in [0, 1]$:*
(i) *$\Phi \to \infty$ is equivalent to $p \gg n^{-1/m(G,H)}$.*
(ii) *$\Phi = \Omega(1)$ is equivalent to $p = \Omega(n^{-1/m(G,H)})$.*
(iii) *If $\Phi \asymp 1$, then $\mu_{G,J} \asymp 1$ for any $G \subseteq J \subseteq H$ that is primal for (G, H).*
(iv) *If $p = \Omega(n^{-1/m(G,H)+\eta})$ for some constant $\eta \geqslant 0$, then $\Phi = \Omega(n^{\eta})$.*

Finally, the approximate result Theorem 4 immediately follows from the following slightly more general theorem, whose technical statement will be convenient in several later proofs. In particular, in some ranges of the parameters, we will be able to deduce the desired 1- or 0-statements directly from (11)–(12) below.

Theorem 7. *For any rooted graph (G, H), the following hold for all $p = p(n) \in [0, 1]$:*
(i) *If $\Phi = \Omega(1)$ and $(t/\mu)^2\Phi \geqslant n^{\Omega(1)}$, then*

$$\lim_{n\to\infty} \mathbb{P}\Big(\max_{\mathbf{x}\in[n]_{v_G}} |X_{\mathbf{x}} - \mu| < t\Big) = 1. \tag{11}$$

(ii) *If $\varepsilon = \varepsilon(n) \in (0, 1]$ and either (a) $\Phi(1-p) \to \infty$ and $\varepsilon^2\Phi/(1-p) \to 0$, or (b) $\Phi \to 0$, then*

$$\lim_{n\to\infty} \mathbb{P}\Big(\max_{\mathbf{x}\in[n]_{v_G}} |X_{\mathbf{x}} - \mu| \geqslant \varepsilon\mu\Big) = 1. \tag{12}$$

Remark 2. *In* (i)*, the conclusion* (11) *holds with probability $1 - o(n^{-\tau})$ for any constant $\tau > 0$.*

We defer the simple proof of Theorem 7 to Appendix A, and only mention the main ideas here. Claim (ii) exploits that $X_{\mathbf{x}}$ is asymptotically normal when $\Phi(1-p) \to \infty$. Claim (i) is based on Markov's inequality and a central moment estimate $\mathbb{E}(X_{\mathbf{x}} - \mu)^{2m} \leqslant C_m\sigma^{2m} \leqslant D_m(\mu^2/\Phi)^m$ that is a by-product of the usual asymptotic normality proof via the method of moments (see Claim 17 in Appendix A). This approach for obtaining tail estimates 'without much effort' does not seem to be as widely known in probabilistic combinatorics, and we believe that it will be useful in other applications (e.g., it yields a simple direct proof of [31, Corollary 4]).

3 Strictly balanced and grounded case (Theorem 1)

In this section we prove the threshold (4) of Theorem 1 (i) for strictly balanced rooted graphs (G, H) that are grounded (see Section 5.2 for the less interesting ungrounded case).

The 0-statement in (4) is the main difficulty, and here the plan is to use a second moment argument to show the existence of a root $\mathbf{x} \in [n]_{v_G}$ with too many (G, H)-extensions, i.e., with $X_{\mathbf{x}} \geqslant (1+\varepsilon)\mu$. Unfortunately, even an asymptotic estimate of the relevant first moment is challenging, since the upper tail probability $\mathbb{P}(X_{\mathbf{x}} \geqslant (1+\varepsilon)\mu)$ is hard to estimate up to a $1+o(1)$ factor (this is an instance of the 'infamous' upper tail problem [16, 29]). To sidestep this technical difficulty, we instead show the existence of a root $\mathbf{x} \in [n]_{v_G}$ which attains $X_{\mathbf{x}} = \lceil (1+\varepsilon)\mu \rceil$ due to exactly $\lceil (1+\varepsilon)\mu \rceil$ extensions that are vertex-disjoint outside of $\mathbf{x}$. The crux is that these auxiliary events are more tractable: we can estimate the relevant first and second moments up to the required $1+o(1)$ factors via a careful mix of Harris' Lemma [13], Janson's inequality [14, 8, 25], and counting arguments. It turns out that here the extra assumption $\varepsilon \geqslant n^{-\alpha}$ is helpful: it will allow us to focus on fairly small edge probabilities $p = p(n)$ that are close to $n^{-1/d(G,H)}$, which intuitively makes it easier to show that various events are approximately independent (as tacitly required by the second moment method); see Section 3.2 for the details.

The 1-statement in (4) is simpler (and nowadays fairly routine). For edge probabilities $p = p(n)$ that are close to $n^{-1/d(G,H)}$, we use a standard union bound argument, estimating the lower tail $\mathbb{P}(X_{\mathbf{x}} \leqslant (1-\varepsilon)\mu)$ via Janson's inequality [14, 15, 25] and the upper tail $\mathbb{P}(X_{\mathbf{x}} \geqslant (1+\varepsilon)\mu)$ via an inequality of Warnke [38]. For edge probabilities $p = p(n)$ much larger than $n^{-1/d(G,H)}$, it turns out that we can simply use the partial result Theorem 7 (i) due to the extra assumption $\varepsilon \geqslant n^{-\alpha}$; see Section 3.3 for the details.

3.1 Technical preliminaries

Our upcoming arguments exploit two standard properties of strictly balanced rooted graphs: (i) for fairly small edge probabilities $p = p(n)$, the expectation $\mu = \mu_{G,H}$ is significantly smaller than any other expectation $\mu_{G,J}$ with $G \subsetneq J \subsetneq H$ (note that $\mu_{G,H}/\mu_{G,J} \asymp n^{v_H - v_J} p^{e_H - e_J} \ll 1$ via (13) below), and (ii) after removing the root vertices from H, the remaining graph $H - V(G)$ is connected. Both mimic well-known properties from the unrooted case, so we defer the routine proof of Lemma 8 to Section 3.4.

Lemma 8. *For any strictly balanced rooted graph (G, H), the following hold:*

(i) *There is a constant $\beta = \beta(G, H) > 0$ such that, for all $p = p(n) \in [0, 1]$ with $p = O(n^{-1/d(G,H)+\beta})$,*

$$\max_{G \subsetneq J \subsetneq H} n^{v_H - v_J} p^{e_H - e_J} \ll n^{-\beta}. \tag{13}$$

(ii) *The graph $H - V(G)$, obtained from H by deleting the vertices of G, is connected.*

3.2 The 0-statement

Our second-moment-based proof of the 0-statement in (4) of Theorem 1 hinges on the following key lemma. Given a root $\mathbf{x} \in [n]_{v_G}$, let $\mathcal{E}_{\mathbf{x}}$ denote the event that, in $\mathbb{G}_{n,p}$, the root $\mathbf{x}$ has exactly $z := \lceil (1+\varepsilon)\mu \rceil$ many (G, H)-extensions, and all of them are pairwise vertex-disjoint (i.e., sharing no vertices outside $\mathbf{x}$). We also say that two roots $\mathbf{x}_1, \mathbf{x}_2 \in [n]_{v_G}$ are *disjoint* if they share no elements as (unordered) sets.

Lemma 9. *Let (G, H) be a rooted graph that is strictly balanced and grounded. There are constants $c, \gamma > 0$ such that, for all $\varepsilon = \varepsilon(n) \in (0, 1]$ and $p = p(n) \in [0, 1]$ with $p \leqslant n^{-1/d(G,H)+\gamma}$, $\mu \geqslant 1/2$ and $\varepsilon^2 \mu \leqslant c \log n$, the following holds: for all roots $\mathbf{x} \in [n]_{v_G}$ we have*

$$\mathbb{P}(\mathcal{E}_{\mathbf{x}}) \gg n^{-1/2}, \tag{14}$$

and for all disjoint roots $\mathbf{x}_1, \mathbf{x}_2 \in [n]_{v_G}$ we have

$$\mathbb{P}(\mathcal{E}_{\mathbf{x}_1}, \mathcal{E}_{\mathbf{x}_2}) \leqslant (1+o(1))\mathbb{P}(\mathcal{E}_{\mathbf{x}_1})\mathbb{P}(\mathcal{E}_{\mathbf{x}_2}). \tag{15}$$

Proof of the 0-statement in (4) of Theorem 1. Let $c, \gamma > 0$ be the constants given by Lemma 9. Fix arbitrary $0 < \alpha < \gamma/2$. First, when $p > n^{-1/d(G,H)+\gamma}$, then $\varepsilon \geqslant n^{-\alpha}$ and Remark 1 (iv) imply $\varepsilon^2 \mu \geqslant n^{-2\alpha} \cdot \Phi_{G,H} =$

$\Omega(n^{\gamma-2\alpha}) \gg \log n$, so the condition of the 0-statement cannot be satisfied and hence there is nothing to prove. Next, when $\mu < 1/2$, then $(1+\varepsilon)\mu \leqslant 2\mu < 1$ and $\varepsilon \leqslant 1$ imply that the interval $((1-\varepsilon)\mu, (1+\varepsilon)\mu)$ contains no integers, and so the 0-statement again holds trivially.

Thus we can henceforth assume $\mu \geqslant 1/2$ and $p \leqslant n^{-1/d(G,H)+\gamma}$, as required by Lemma 9. For convenience, we set $s := \lfloor n/v_G \rfloor \asymp n$, and choose disjoint roots $\mathbf{x}_1, \ldots, \mathbf{x}_s \in [n]_{v_G}$. Writing $Y := |\{i \in [s] : \mathcal{E}_{\mathbf{x}_i} \text{ holds}\}|$, to prove the 0-statement of Theorem 1 we shall now show that $Y > 0$ whp, i.e., that $\mathbb{P}(Y > 0) \to 1$ as $n \to \infty$. Using (14) we obtain $\mathbb{E}Y = \sum_{1 \leqslant i \leqslant s} \mathbb{P}(\mathcal{E}_{\mathbf{x}_i}) \gg s \cdot n^{-1/2} \asymp n^{1/2} \to \infty$. Together with (15) it follows that

$$\mathbb{E}Y^2 \leqslant \sum_{1 \leqslant i,j \leqslant s:\, i \neq j} \mathbb{P}(\mathcal{E}_{\mathbf{x}_i},\, \mathcal{E}_{\mathbf{x}_j}) + \sum_{1 \leqslant i \leqslant s} \mathbb{P}(\mathcal{E}_{\mathbf{x}_i}) \;\leqslant\; (1+o(1)) \cdot (\mathbb{E}Y)^2 + \mathbb{E}Y \;\sim\; (\mathbb{E}Y)^2.$$

Now Chebyshev's inequality readily yields $\mathbb{P}(Y = 0) \leqslant \operatorname{Var} Y/(\mathbb{E}Y)^2 \to 0$ as $n \to \infty$, completing the proof. $\square$

The remainder of Section 3.2 is dedicated to the proof of Lemma 9. For concreteness, for $\beta > 0$ as given by Lemma 8 (i), we choose the constants $\gamma, c \in (0, 1/2)$ such that

$$\gamma e_H \;<\; \min\bigl\{\beta/v_H,\, \beta/2,\, 1/2,\, 1-2c\bigr\}. \tag{16}$$

Recalling $\mu \asymp n^{v_H - v_G} p^{e_H - e_G}$ and $\varepsilon \leqslant 1$, using the assumptions $\mu \geqslant 1/2$ and $p \leqslant n^{-1/d(G,H)+\gamma}$, we infer

$$1/2 \;\leqslant\; \mu \;\leqslant\; z = \lceil (1+\varepsilon)\mu \rceil \;\leqslant\; O(n^{\gamma e_H}) \ll \min\bigl\{n^{1/2}, n^{\beta/2}\bigr\}, \tag{17}$$

and

$$p \leqslant \Bigl(n^{-(v_H - v_G) + \gamma(e_H - e_G)}\Bigr)^{\frac{1}{e_H - e_G}} \leqslant \Bigl(n^{-1+1/2}\Bigr)^{\frac{1}{e_H - e_G}} \ll 1/2, \tag{18}$$

with room to spare. With foresight, given $\mathbf{x} \in [n]_{v_G}$, we denote by $N = N_{G,H}(\mathbf{x})$ the number of (G,H)-extensions of $\mathbf{x}$ in K_n. Note that $N \asymp n^{v_H - v_G}$ does not depend on the particular choice of $\mathbf{x}$.

3.2.1 The first moment: inequality (14)

We start with (14), i.e., a lower bound for $\mathbb{P}(\mathcal{E}_{\mathbf{x}})$. Recall that every $\mathbf{x} \in [n]_{v_G}$ has N extensions in K_n. The plan is to show that $\mathbb{P}(\mathcal{E}_{\mathbf{x}})$ is comparable with $\mathbb{P}(\mathrm{Bin}(N, p^{e_H - e_G}) = z)$. More precisely, we will show that

$$\mathbb{P}(\mathcal{E}_{\mathbf{x}}) \;\geqslant\; (1+o(1)) \cdot \binom{N}{z} p^{(e_H - e_G)z} (1 - p^{e_H - e_G})^{N-z}. \tag{19}$$

In view of $z \approx (1+\varepsilon)\mu = (1+\varepsilon)Np^{e_H - e_G}$, using Stirling's formula it then will be routine to deduce that the lower bound in (19) is $\Theta(z^{-1/2}) \cdot e^{-\Theta(\varepsilon^2 \mu)}$, which together with (16)–(17) and the assumption $\varepsilon^2 \mu \leqslant c \log n$ will eventually imply the desired inequality (14); see (28)–(29) below.

Turning to the technical details, given $\mathbf{x} \in [n]_{v_G}$, let $\mathfrak{H}(\mathbf{x})$ denote the set of all (unordered) collections of $z = \lceil (1+\varepsilon)\mu \rceil$ vertex-disjoint (G,H)-extensions of $\mathbf{x}$ in K_n. Given $\mathcal{C} \in \mathfrak{H}(\mathbf{x})$, let $\mathcal{C}^c$ denote the remaining $N - z$ extensions of $\mathbf{x}$ in K_n. Given a collection $\mathcal{S}$ of extensions of $\mathbf{x}$, we write $\mathcal{I}_{\mathcal{S}}$ for the event that all extensions in $\mathcal{S}$ are present in $\mathbb{G}_{n,p}$, and $\mathcal{D}_{\mathcal{S}}$ for the event that all extensions in $\mathcal{S}$ are not present in $\mathbb{G}_{n,p}$. Note that

$$\mathbb{P}(\mathcal{E}_{\mathbf{x}}) = \sum_{\mathcal{C} \in \mathfrak{H}(\mathbf{x})} \mathbb{P}(\mathcal{I}_{\mathcal{C}},\, \mathcal{D}_{\mathcal{C}^c}) = \sum_{\mathcal{C} \in \mathfrak{H}(\mathbf{x})} \mathbb{P}(\mathcal{I}_{\mathcal{C}}) \mathbb{P}(\mathcal{D}_{\mathcal{C}^c} \mid \mathcal{I}_{\mathcal{C}}) \geqslant |\mathfrak{H}(\mathbf{x})| \min_{\mathcal{C} \in \mathfrak{H}(\mathbf{x})} \mathbb{P}(\mathcal{I}_{\mathcal{C}}) \mathbb{P}(\mathcal{D}_{\mathcal{C}^c} \mid \mathcal{I}_{\mathcal{C}}), \tag{20}$$

where the minimum is of course only formal: by symmetry the probabilities are the same for every $\mathcal{C} \in \mathfrak{H}(\mathbf{x})$. To estimate $|\mathfrak{H}(\mathbf{x})|$, note that given $i \leqslant z$ vertex-disjoint extensions, the number of choices for another vertex-disjoint extension is $N - O(zn^{v_H - v_G - 1})$. Since $\mathfrak{H}(\mathbf{x})$ consists of unordered collections of extensions, using $N \asymp n^{v_H - v_G}$ and $z \ll n^{1/2}$ (see (17)) together with $1 - x = e^{-x(1+o(1))}$ as $x \to 0$ it follows that

$$|\mathfrak{H}(\mathbf{x})| = \frac{\bigl(N - O(zn^{v_H - v_G - 1})\bigr)^z}{z!} = \frac{N^z}{z!} \cdot \Bigl(1 - O\bigl(z/n\bigr)\Bigr)^z \sim \frac{N^z}{z!} \sim \binom{N}{z}. \tag{21}$$

Since the extensions in $\mathcal{C} \in \mathfrak{H}(\mathbf{x})$ are disjoint, we have

$$\mathbb{P}(\mathcal{I}_{\mathcal{C}}) = p^{(e_H - e_G)z}. \tag{22}$$

For the remaining lower bound on $\mathbb{P}(\mathcal{D}_{\mathcal{C}^c}|\mathcal{I}_{\mathcal{C}})$, the idea is to apply Harris' Lemma [13] and then use Lemma 8 (i) to show that the effect of 'overlapping' pairs of extensions is negligible.

Claim 10. *Let $\mathbf{x} \in [n]_{v_G}$. Then, for all $\mathcal{C} \in \mathfrak{H}(\mathbf{x})$, we have*

$$\mathbb{P}(\mathcal{D}_{\mathcal{C}^c}|\mathcal{I}_{\mathcal{C}}) \;\geqslant\; (1+o(1)) \cdot (1 - p^{e_H - e_G})^{N-z}. \tag{23}$$

Proof. We fix $\mathcal{C} \in \mathfrak{H}(\mathbf{x})$, and define the auxiliary graph $F := \big([n],\ \bigcup_{H_1 \in \mathcal{C}} E(H_1)\big)$. Note that after conditioning on the event $\mathcal{I}_{\mathcal{C}}$, in $\mathbb{G}_{n,p}$ each possible edge from $E(K_n) \setminus E(F)$ is still included independently with probability p. Therefore Harris' Lemma (see, e.g., [1, Theorem 6.3.2]) implies that

$$\mathbb{P}(\mathcal{D}_{\mathcal{C}^c}|\mathcal{I}_{\mathcal{C}}) \geqslant \prod_{H_2 \in \mathcal{C}^c} \big(1 - p^{e_H - e_G - e(H_2 \cap F)}\big). \tag{24}$$

Note that there are at most $N - z$ extensions $H_2 \in \mathcal{C}^c$ with $e(H_2 \cap F) = 0$, each contributing a factor of $1 - p^{e_H - e_G}$ to the right-hand side of (24). Every other extension $H_2 \in \mathcal{C}^c$ contains at least one edge not in F (since by Lemma 8 (ii), after deleting the root vertices $\mathbf{x}$, all graphs in $\{H_1 - \mathbf{x} : H_1 \in \mathcal{C}\}$ are vertex-disjoint and connected), so that $p^{e_H - e_G - e(H_2 \cap F)} \leqslant p \leqslant 1/2$ by (18). Since $1 - x \geqslant e^{-2x}$ for $x \leqslant 1/2$, from (24) it follows that

$$\mathbb{P}(\mathcal{D}_{\mathcal{C}^c}|\mathcal{I}_{\mathcal{C}}) \geqslant (1 - p^{e_H - e_G})^{N-z} \cdot \exp\Big(-2 \sum_{\substack{H_2 \in \mathcal{C}^c:\\ e(H_2 \cap F) \geqslant 1}} p^{e_H - e_G - e(H_2 \cap F)}\Big). \tag{25}$$

To estimate the sum in (25), note that if $H_2 \in \mathcal{C}^c$ shares an edge with F, then $E(H_2 \cap F)$ corresponds to a (G,J)-extension of $\mathbf{x}$ for some $G \subsetneq J \subsetneq H$. The number of such extensions is at most $(v_H z)^{v_J - v_G} = O(z^{v_H})$, with room to spare. Given a (G,J)-extension, it can be further extended to some $H_2 \in \mathcal{C}^c$ in at most $n^{v_H - v_J}$ ways. Using $e_H - e_G - (e_J - e_G) = e_H - e_J$ together with (17) and (13), it follows that

$$\sum_{\substack{H_2 \in \mathcal{C}^c:\\ e(H_2 \cap F) \geqslant 1}} p^{e_H - e_G - e(H_2 \cap F)} \leqslant \sum_{G \subsetneq J \subsetneq H} O\Big(z^{v_H} n^{v_H - v_J} \cdot p^{e_H - e_J}\Big) \ll n^{\gamma e_H v_H - \beta} = o(1), \tag{26}$$

which together with (25) establishes inequality (23). ☐

Combining estimates (20)–(23), we readily obtain inequality (19). To establish (14), it remains to estimate the right-hand side of (19) via the following well-known form of Stirling's formula (see, e.g., [7, equation (1.4)]):

$$n! = \sqrt{2\pi n}\left(\frac{n}{e}\right)^n e^{\alpha_n} \qquad \text{with} \qquad \alpha_n = O(n^{-1}). \tag{27}$$

With foresight, let $t := z - \mu = \varepsilon\mu + O(1)$, and define $\varphi(x) := (1+x)\log(1+x) - x$ for $x > -1$. Recalling (17) we have $1 \leqslant z \ll n^{1/2} \ll N$. Using Stirling's formula (27) together with $\mu = Np^{e_H - e_G}$ and $z = \mu + t$, then a simple (but slightly tedious) calculation along the lines of the Appendix of [37] gives

$$\begin{aligned}\binom{N}{z} p^{(e_H - e_G)z}(1 - p^{e_H - e_G})^{N-z} &\geqslant \frac{\exp\Big(-O\big(N^{-1} + z^{-1} + (N-z)^{-1}\big)\Big)}{\sqrt{2\pi z(1 - z/N)}} \cdot \left(\frac{\mu}{z}\right)^z \left(\frac{N-\mu}{N-z}\right)^{N-z} \\ &\geqslant \Omega(z^{-1/2}) \cdot \exp\Big(-\mu\varphi\big(t/\mu\big) - (N-\mu)\varphi\big(-t/(N-\mu)\big)\Big).\end{aligned} \tag{28}$$

Note that $\log(1+x) \leqslant x$ implies $\varphi(x) \leqslant x^2$. Using (17) we readily infer $t^2/\mu = \varepsilon^2\mu + O(1)$. Furthermore, (16) implies $p^{e_H - e_G} \leqslant n^{-1/2}$, so that $N \geqslant n^{1/2}\mu \gg \mu$. Using the estimates (17) and $\varepsilon^2\mu \leqslant c\log n$ together with $\gamma e_H/2 + c < 1/2$ (see (16)), it now follows that (28) is at least

$$\Omega\big(z^{-1/2}\big) \cdot \exp\Big(-\big(1 + O\big(n^{-1/2}\big)\big)\varepsilon^2\mu\Big) \;\geqslant\; \Omega(1) \cdot \exp\Big(-\big(\gamma e_H/2 + c\big)\log n\Big) \gg n^{-1/2}, \tag{29}$$

which together with (19) completes the proof of inequality (14) from Lemma 9.

3.2.2 The second moment: inequality (15)

Now we turn to (15), i.e., an upper bound for $\mathbb{P}(\mathcal{E}_{\mathbf{x}_1}, \mathcal{E}_{\mathbf{x}_2})$ when $\mathbf{x}_1$, $\mathbf{x}_2$ are disjoint. Recalling (20), note that

$$\mathbb{P}(\mathcal{E}_{\mathbf{x}_1}, \mathcal{E}_{\mathbf{x}_2}) = \sum_{\mathcal{C}_1 \in \mathfrak{H}(\mathbf{x}_1)} \sum_{\mathcal{C}_2 \in \mathfrak{H}(\mathbf{x}_2, \mathcal{C}_1)} \mathbb{P}(\mathcal{I}_{\mathcal{C}_1 \cup \mathcal{C}_2}, \mathcal{D}_{\mathcal{C}_1^c \cup \mathcal{C}_2^c}), \tag{30}$$

where we (with foresight) define

$$\mathfrak{H}(\mathbf{x}_2, \mathcal{C}_1) := \big\{\mathcal{C}_2 \in \mathfrak{H}(\mathbf{x}_2) : \mathbb{P}(\mathcal{I}_{\mathcal{C}_1 \cup \mathcal{C}_2}, \mathcal{D}_{\mathcal{C}_1^c \cup \mathcal{C}_2^c}) > 0\big\}. \tag{31}$$

Guided by the heuristics that the various events are approximately independent, the plan is to show that

$$\mathbb{P}(\mathcal{I}_{\mathcal{C}_1 \cup \mathcal{C}_2}, \mathcal{D}_{\mathcal{C}_1^c \cup \mathcal{C}_2^c}) \leqslant (1+o(1))\mathbb{P}(\mathcal{I}_{\mathcal{C}_1}, \mathcal{D}_{\mathcal{C}_1^c}) \cdot \mathbb{P}(\mathcal{I}_{\mathcal{C}_2}, \mathcal{D}_{\mathcal{C}_2^c}), \tag{32}$$

though the actual details will be slightly more involved. Ignoring these complications for now, note that (32) would together with (30), (20) and $\mathfrak{H}(\mathbf{x}_2, \mathcal{C}_1) \subseteq \mathfrak{H}(\mathbf{x}_2)$ indeed imply the desired inequality (15).

Turning to the technical details, by applying Harris' Lemma (noting that $\mathcal{I}_{\mathcal{C}_1 \cup \mathcal{C}_2}$ is an increasing event and that $\mathcal{D}_{\mathcal{C}_1^c \cup \mathcal{C}_2^c}$ is a decreasing event; see the definitions above [1, Theorem 6.3.2]) to the right-hand side of (30) we obtain that

$$\mathbb{P}(\mathcal{E}_{\mathbf{x}_1}, \mathcal{E}_{\mathbf{x}_2}) \leqslant \sum_{\mathcal{C}_1 \in \mathfrak{H}(\mathbf{x}_1)} \sum_{\mathcal{C}_2 \in \mathfrak{H}(\mathbf{x}_2, \mathcal{C}_1)} \mathbb{P}(\mathcal{I}_{\mathcal{C}_1 \cup \mathcal{C}_2})\mathbb{P}(\mathcal{D}_{\mathcal{C}_1^c \cup \mathcal{C}_2^c}). \tag{33}$$

Recalling that every $\mathbf{x} \in [n]_{v_G}$ has N extensions in K_n, Harris' Lemma also gives the lower bound $\mathbb{P}(\mathcal{D}_{\mathcal{C}_1^c \cup \mathcal{C}_2^c}) \geqslant (1-p^{e_H - e_G})^{2(N-z)}$. We will now prove an asymptotically matching upper bound that does *not* depend on the choice of $\mathcal{C}_1$ and $\mathcal{C}_2$ (similarly as in Claim 10). Here the idea is to apply a form of Janson's inequality [8, 15, 1], and then again use Lemma 8 (i) to argue that 'overlaps' have negligible contribution.

Claim 11. *Let $\mathbf{x}_1, \mathbf{x}_2 \in [n]_{v_G}$ be disjoint. Then, for all $\mathcal{C}_1 \in \mathfrak{H}(\mathbf{x}_1)$ and $\mathcal{C}_2 \in \mathfrak{H}(\mathbf{x}_2)$, we have*

$$\mathbb{P}(\mathcal{D}_{\mathcal{C}_1^c \cup \mathcal{C}_2^c}) \leqslant (1+o(1)) \cdot (1-p^{e_H - e_G})^{2(N-z)}. \tag{34}$$

Proof. Let $\mathcal{S}$ be the family of edge-sets, each of size $e_H - e_G$, corresponding to extensions in $\mathcal{C}_1^c \cup \mathcal{C}_2^c$ (each extension of $\mathbf{x}_1$ or $\mathbf{x}_2$ is uniquely determined by its edge-set, since H has no isolated vertices outside of $V(G)$ by Lemma 8 (ii)). Note that if an extension in $\mathcal{C}_1^c$ is also an extension in $\mathcal{C}_2^c$, then it must contain some vertex from $\mathbf{x}_2$ (because (G, H) is grounded). Since $\mathbf{x}_1, \mathbf{x}_2$ are disjoint, the number of such duplicate extensions is $O(n^{v_H - v_G - 1})$, which implies that $|\mathcal{S}| \geqslant 2(N-z) - O(n^{v_H - v_G - 1})$. Setting $X := \sum_{E \in \mathcal{S}} \mathbb{1}_{\{E \subseteq \mathbb{G}_{n,p}\}}$, note that the event $\mathcal{D}_{\mathcal{C}_1^c \cup \mathcal{C}_2^c}$ is precisely the event that $X = 0$. Since $p \leqslant 1/2$ (see (18)) implies $1/(1-p^{e_H - e_G}) \leqslant 2$ and $(1-p^{e_H - e_G})^{-1} \leqslant e^{2p^{e_H - e_G}}$, by invoking the Boppana–Spencer [8] variant of Janson's inequality (see, e.g., [15, Remark 2.20] or [1, Theorem 8.1.1]) it then follows that

$$\mathbb{P}(\mathcal{D}_{\mathcal{C}_1^c \cup \mathcal{C}_2^c}) = \mathbb{P}(X=0) \leqslant (1-p^{e_H - e_G})^{|\mathcal{S}|} \cdot e^{\Delta/(1-p^{e_H - e_G})} \leqslant (1-p^{e_H - e_G})^{2(N-z)} \cdot e^{O(n^{v_H - v_G - 1} p^{e_H - e_G} + \Delta)}, \tag{35}$$

where

$$\Delta := \sum_{\substack{(E_1, E_2) \in \mathcal{S} \times \mathcal{S}: \\ 1 \leqslant |E_1 \cap E_2| < e_H - e_G}} p^{|E_1 \cup E_2|}. \tag{36}$$

Using $\mu = Np^{e_H - e_G} \asymp n^{v_H - v_G} p^{e_H - e_G}$ together with (17), it follows that

$$n^{v_H - v_G - 1} p^{e_H - e_G} \asymp \mu \cdot n^{-1} \ll n^{1/2 - 1} = o(1). \tag{37}$$

Turning to the Δ-term, note that $|\mathcal{S}| p^{e_H - e_G} \leqslant 2(N-z)p^{e_H - e_G} \leqslant 2\mu$. By proceeding analogously to the estimates in (25)–(26), using (17) and (13) it routinely follows that

$$\Delta \leqslant \sum_{E_1 \in \mathcal{S}} p^{e_H - e_G} \sum_{\substack{E_2 \in \mathcal{S}: \\ 1 \leqslant |E_1 \cap E_2| < e_H - e_G}} p^{e_H - e_G - |E_1 \cap E_2|} \leqslant O\Big(\mu \cdot \sum_{G \subsetneq J \subsetneq H} n^{v_H - v_J} p^{e_H - e_J}\Big) = o(1), \tag{38}$$

which together with (35)–(37) establishes inequality (34). $\square$

To sum up, by inserting the estimates (22) and (34) into (33), we readily arrive at

$$\mathbb{P}(\mathcal{E}_{\mathbf{x}_1}, \mathcal{E}_{\mathbf{x}_2}) \;\leqslant\; (1+o(1)) \cdot p^{(e_H - e_G)z}(1-p^{e_H-e_G})^{2(N-z)} \sum_{\mathcal{C}_1 \in \mathfrak{H}(\mathbf{x}_1)} \sum_{\mathcal{C}_2 \in \mathfrak{H}(\mathbf{x}_2, \mathcal{C}_1)} \mathbb{P}(\mathcal{I}_{\mathcal{C}_2} \mid \mathcal{I}_{\mathcal{C}_1}). \tag{39}$$

Anticipating that the main contribution comes from pairs $\mathcal{C}_1, \mathcal{C}_2$ of 'disjoint' collections, we partition

$$\mathfrak{H}(\mathbf{x}_1) := \mathfrak{H}_0(\mathbf{x}_1, \mathbf{x}_2) \cup \mathfrak{H}_{\geqslant 1}(\mathbf{x}_1, \mathbf{x}_2), \tag{40}$$

where $\mathfrak{H}_0(\mathbf{x}_1, \mathbf{x}_2)$ contains the collections $\mathcal{C}_1 \in \mathfrak{H}(\mathbf{x}_1)$ for which the auxiliary graph

$$F = F(\mathcal{C}_1) := \Big([n],\, \textstyle\bigcup_{H' \in \mathcal{C}_1} E(H')\Big) \tag{41}$$

contains no extensions of $\mathbf{x}_2$, and $\mathfrak{H}_{\geqslant 1}(\mathbf{x}_1, \mathbf{x}_2)$ contains the remaining ones. Since $\mathbf{x}_1, \mathbf{x}_2$ are disjoint and (G, H) is grounded, every $\mathcal{C}_1 \in \mathfrak{H}_{\geqslant 1}(\mathbf{x}_1, \mathbf{x}_2)$ must contain at least one extension overlapping with $\mathbf{x}_2$ (in at least one vertex). From (21), $N \asymp n^{v_H - v_G}$ and $z \ll n$ (see (17)) it follows that, for some constant $A = A(G, H) > 0$,

$$|\mathfrak{H}_{\geqslant 1}(\mathbf{x}_1, \mathbf{x}_2)| \leqslant A n^{v_H - v_G - 1} \cdot \binom{N}{z-1} \asymp n^{v_H - v_G - 1} \cdot \frac{z}{N} \cdot |\mathfrak{H}(\mathbf{x}_1)| \ll |\mathfrak{H}(\mathbf{x}_1)|. \tag{42}$$

Exploiting the groundedness assumption, we next show that pairs $\mathcal{C}_1, \mathcal{C}_2$ can only overlap in at most $v_G = O(1)$ extensions (see Claim 12), and that overlapping pairs effectively have negligible contribution (see Claim 13).

Claim 12. *Let $\mathbf{x}_1, \mathbf{x}_2 \in [n]_{v_G}$ be disjoint. Then, for all $\mathcal{C}_1 \in \mathfrak{H}(\mathbf{x}_1)$, the graph $F = F(\mathcal{C}_1)$ defined in* (41) *contains at most v_G vertex-disjoint extensions of $\mathbf{x}_2$.*

Proof. The graph $F - \mathbf{x}_1$, obtained by removing the vertices $\mathbf{x}_1$ from F, consists of isolated vertices and vertex-disjoint copies of the graph $H - V(G)$, which, by Lemma 8 (ii), is connected. Let H' be obtained from $H - E(G)$ by removing isolated root vertices (if any). Since (G, H) is grounded, we have $e_{H-V(G)} < e_{H'}$. Note that H' is connected (since it equals $H - V(G)$ with some root vertices connected to it) and therefore $F - \mathbf{x}_1$ is H'-free. It follows that any extension of $\mathbf{x}_2$ that is present in F must intersect $\mathbf{x}_1$, so there are at most $|\mathbf{x}_1| = v_G$ such vertex-disjoint extensions of $\mathbf{x}_2$. □

Claim 13. *Let $\mathbf{x}_1, \mathbf{x}_2 \in [n]_{v_G}$ be disjoint. Then*

$$\sum_{\mathcal{C}_1 \in \mathfrak{H}(\mathbf{x}_1)} \sum_{\mathcal{C}_2 \in \mathfrak{H}(\mathbf{x}_2, \mathcal{C}_1)} \mathbb{P}(\mathcal{I}_{\mathcal{C}_2} \mid \mathcal{I}_{\mathcal{C}_1}) \;\leqslant\; (1+o(1)) \sum_{\mathcal{C}_1 \in \mathfrak{H}(\mathbf{x}_1)} \sum_{\mathcal{C}_2 \in \mathfrak{H}(\mathbf{x}_2)} \mathbb{P}(\mathcal{I}_{\mathcal{C}_2}). \tag{43}$$

Proof of Claim 13. In the first step we estimate $\sum_{\mathcal{C}_2 \in \mathfrak{H}(\mathbf{x}_2, \mathcal{C}_1)} \mathbb{P}(\mathcal{I}_{\mathcal{C}_2} \mid \mathcal{I}_{\mathcal{C}_1})$ using a counting argument that accounts for the different kinds of overlaps of $\mathcal{C}_2$ with the graph $F = F(\mathcal{C}_1)$ defined in (41). Turning to the details, as in the proof of Claim 11 we will think of (G, H)-extensions as edge-sets of size $e_H - e_G$. Recall that $|\mathcal{C}_1| = |\mathcal{C}_2| = z = \lceil (1+\varepsilon)\mu \rceil$. Suppose that the graph F contains k extensions of $\mathbf{x}_2$. If $\mathcal{C}_2 \in \mathfrak{H}(\mathbf{x}_2, \mathcal{C}_1)$ then all these k extensions must be present in $\mathcal{C}_2$, since otherwise $\mathbb{P}(\mathcal{I}_{\mathcal{C}_1 \cup \mathcal{C}_2}, \mathcal{D}_{\mathcal{C}_1^c \cup \mathcal{C}_2^c}) \leqslant \mathbb{P}(\mathcal{I}_{\mathcal{C}_1}, \mathcal{D}_{\mathcal{C}_2^c}) = 0$ contradicting $\mathcal{C}_2 \in \mathfrak{H}(\mathbf{x}_2, \mathcal{C}_1)$. List the remaining extensions in $\mathcal{C}_2$ as $E_1, \ldots, E_{z-k}$ in an arbitrary order. Note that each E_i is not fully contained in $E(F)$, and thus the intersection $E_i \cap E(F)$ is the edge-set of some (G, J_i)-extension of $\mathbf{x}_2$ for some graph J_i satisfying $G \subseteq J_i \subsetneq H$ (the case $J_i = G$ occurs when the extension E_i is edge-disjoint from F). When these intersections are given by $J_1, \ldots, J_{z-k}$, then we clearly have

$$\mathbb{P}\left(\mathcal{I}_{\mathcal{C}_2} \mid \mathcal{I}_{\mathcal{C}_1}\right) = \prod_{i=1}^{z-k} p^{e_H - e_G - (e_{J_i} - e_G)} = \prod_{i=1}^{z-k} p^{e_H - e_{J_i}}.$$

Furthermore, the number of sequences $E_1, \ldots, E_{z-k}$ corresponding to intersections $J_1, \ldots, J_{z-k}$ is bounded from above by

$$\prod_{i=1}^{z-k} \big(v_G + (v_H - v_G)z\big)^{v_{J_i} - v_G} \hat{N}_{J_i, H},$$

where $\hat{N}_{J,H} := N_{G,H} = N$ if $G = J$ and $\hat{N}_{J,H} := n^{v_H - v_J}$ otherwise. Hence, summing over all possible choices of $J_1, \ldots, J_{z-k}$ and dividing by $(z-k)!$ (since we sum over unordered collections $\mathcal{C}_2$), it follows that

$$\begin{aligned}\sum_{\mathcal{C}_2 \in \mathfrak{H}(\mathbf{x}_2, \mathcal{C}_1)} \mathbb{P}(\mathcal{I}_{\mathcal{C}_2} \mid \mathcal{I}_{\mathcal{C}_1}) &\leqslant \frac{1}{(z-k)!} \sum_{\substack{J_1, \ldots, J_{z-k}: \\ G \subseteq J_i \subsetneq H}} \prod_{i=1}^{z-k} \bigl(v_G + (v_H - v_G)z\bigr)^{v_{J_i} - v_G} \hat{N}_{J_i,H} p^{e_H - e_{J_i}} \\ &\leqslant \frac{z^k}{z!} \cdot \Biggl(\sum_{G \subseteq J \subsetneq H} \bigl(v_G + (v_H - v_G)z\bigr)^{v_J - v_G} \hat{N}_{J,H} p^{e_H - e_J} \Biggr)^{z-k}. \end{aligned} \tag{44}$$

Noting that $\hat{N}_{G,H} p^{e_H - e_G} = \mu$, using (26) and $\mu \asymp z$ we bound the sum in (44) from above by, say,

$$\mu + O\Bigl(\sum_{G \subsetneq J \subsetneq H} z^{v_H} n^{v_H - v_J} p^{e_H - e_J} \Bigr) \leqslant \mu + o(1) = \mu \cdot \Bigl(1 + o\bigl(z^{-1}\bigr)\Bigr). \tag{45}$$

From the assumptions $\varepsilon \leqslant 1$ and $\mu \geqslant 1/2$ it follows that $z \leqslant (1+\varepsilon)\mu + 1 \leqslant 4\mu$, say. Therefore, in view of (44)–(45), using $\mu = N p^{e_H - e_G}$ and (21) it follows that

$$\sum_{\mathcal{C}_2 \in \mathfrak{H}(\mathbf{x}_2, \mathcal{C}_1)} \mathbb{P}(\mathcal{I}_{\mathcal{C}_2} \mid \mathcal{I}_{\mathcal{C}_1}) \leqslant \left(\frac{z}{\mu}\right)^k \frac{(N p^{e_H - e_G})^z}{z!} \Bigl(1 + o\bigl(z^{-1}\bigr)\Bigr)^{z-k} \;\leqslant\; (1 + o(1)) \cdot 4^k |\mathfrak{H}(\mathbf{x}_2)| p^{(e_H - e_G)z}, \tag{46}$$

whenever the graph F defined in (41) contains exactly k extensions of $\mathbf{x}_2$.

In the second step we sum the above estimate (46) over all $\mathcal{C}_1 \in \mathfrak{H}(\mathbf{x}_1)$. Recalling the partition (40), note that $k = 0$ when $\mathcal{C}_1 \in \mathfrak{H}_0(\mathbf{x}_1, \mathbf{x}_2)$, and that $k \leqslant v_G$ otherwise (see Claim 12). From (46) it follows that

$$\sum_{\mathcal{C}_1 \in \mathfrak{H}(\mathbf{x}_1)} \sum_{\mathcal{C}_2 \in \mathfrak{H}(\mathbf{x}_2, \mathcal{C}_1)} \mathbb{P}(\mathcal{I}_{\mathcal{C}_2} \mid \mathcal{I}_{\mathcal{C}_1}) \leqslant (1 + o(1)) \cdot \Bigl(|\mathfrak{H}_0(\mathbf{x}_1, \mathbf{x}_2)| + 4^{v_G} |\mathfrak{H}_{\geqslant 1}(\mathbf{x}_1, \mathbf{x}_2)| \Bigr) \cdot |\mathfrak{H}(\mathbf{x}_2)| p^{(e_H - e_G)z}.$$

In view of (42), the factor in the above parentheses is at most $(1 + o(1)) \cdot |\mathfrak{H}(\mathbf{x}_1)|$, say, which together with $p^{(e_H - e_G)z} = \mathbb{P}(\mathcal{I}_{\mathcal{C}_2})$ from (22) then completes the proof of inequality (43). $\square$

Finally, inserting the estimates (43), $p^{(e_H - e_G)z} = \mathbb{P}(\mathcal{I}_{\mathcal{C}_1})$, and (23) into (39), it follows that

$$\mathbb{P}(\mathcal{E}_{\mathbf{x}_1}, \mathcal{E}_{\mathbf{x}_2}) \;\leqslant\; (1 + o(1)) \sum_{\mathcal{C}_1 \in \mathfrak{H}(\mathbf{x}_1)} \mathbb{P}(\mathcal{I}_{\mathcal{C}_1}) \mathbb{P}(\mathcal{D}_{\mathcal{C}_1^c} | \mathcal{I}_{\mathcal{C}_1}) \sum_{\mathcal{C}_2 \in \mathfrak{H}(\mathbf{x}_2)} \mathbb{P}(\mathcal{I}_{\mathcal{C}_2}) \mathbb{P}(\mathcal{D}_{\mathcal{C}_2^c} | \mathcal{I}_{\mathcal{C}_2}),$$

which together with (20) completes the proof of inequality (15) and thus Lemma 9 (which in turn implies the 0-statement in (4) of Theorem 1, as discussed). $\square$

3.3 The 1-statement

Our proof of the 1-statement in (4) of Theorem 1 is based on a fairly standard union bound argument.

Proof of the 1*-statement in* (4) *of Theorem 1.* Fix an arbitrary constant $\tau > 0$. For $\beta > 0$ as given by Lemma 8 (i), fix constants $0 < \gamma \leqslant \beta$ and $0 < \alpha < \gamma/2$ as in the proof of the 0-statement (see Section 3.2). If $p > n^{-1/d(G,H)+\gamma}$, then Remark 1 (iv) implies $\Phi_{G,H} = \Omega(n^\gamma)$, and using $\varepsilon^2 \Phi_{G,H} = \Omega(n^{\gamma - 2\alpha}) = n^{\Omega(1)}$ we see that the 1-statement of Theorem 1 follows from Theorem 7 (i) with $t = \varepsilon\mu$.

In the remaining (main) case $p \leqslant n^{-1/d(G,H)+\gamma}$, we fix a root $\mathbf{x} \in [n]_{v_G}$. Since there are $O(n^{v_G})$ many such roots, for the 1-statement of Theorem 1 it suffices to show that, for $C > 0$ large enough,

$$\mathbb{P}\left(|X_{\mathbf{x}} - \mu| \geqslant \varepsilon\mu\right) = o\bigl(n^{-(v_G + \tau)}\bigr) \qquad \text{if } \varepsilon^2 \mu \geqslant C \log n. \tag{47}$$

To avoid clutter, we shall henceforth use the convention that all implicit constants c_i may depend on (G, H). For the lower tail we shall apply Janson's inequality [25, Theorem 1] analogously to the textbook argument [15, 18] for unrooted subgraph counts, which in view of (13) from Lemma 8 (i) routinely gives

$$\mathbb{P}\left(X_{\mathbf{x}} \leqslant (1 - \varepsilon)\mu\right) \;\leqslant\; \exp\Bigl(-c_1 \varepsilon^2 \mu\Bigr) \leqslant n^{-c_1 C} = o\bigl(n^{-(v_G + \tau)}\bigr) \tag{48}$$

for $C > (v_G + \tau)/c_1$ (similar to (36) and (38), the relevant Δ-term of Janson's inequality, which here is defined in terms of the family $\mathcal{S}$ of edge-sets corresponding to extensions of $\mathbf{x}$ in K_n, satisfies $\Delta = o(1)$ by (17) and (13)). For the upper tail we shall apply [38, Theorem 32] in the setting described in [38, Example 20] (the conditions (Hℓ), (P), (Pq) are defined in [38, Section 4.1]). The underlying hypergraph $\mathcal{H} = \mathfrak{H}(\mathbf{x})$ consists of the edge-sets of extensions of $\mathbf{x}$, thus having vertex-set $V(\mathcal{H}) = E(K_n)$. We set the parameters to $N = n^2$, $\ell = 1$, $q = k = e_H - e_G$, and $K = v_G + 2\tau$. The quantity μ_j from [38, Example 20] satisfies $\max_{1 \leqslant j < q} \mu_j \leqslant \max_{G \subsetneq J \subsetneq H} n^{v_H - v_J} p^{e_H - e_J} \ll n^{-\beta}$ by Lemma 8 (i). Invoking [38, Theorem 32], it then follows that

$$\mathbb{P}\left(X_{\mathbf{x}} \geqslant (1+\varepsilon)\mu\right) \leqslant \big(1 + o(1)\big) \cdot \exp\Big(-\min\big\{c_2 \varepsilon^2 \mu,\ (v_G + 2\tau)\log n\big\}\Big) = o\big(n^{-(r+\tau)}\big) \tag{49}$$

for $C > (v_G + \tau)/c_2$, completing the proof of (47) and thus the 1-statement in (4) of Theorem 1. □

Remark 3 (Theorem 1: stronger 1-statement)**.** *The above proof yields, in view of Remark 2, the following stronger conclusion: for any fixed $\tau > 0$ there is a constant $C = C(\tau, G, H) > 0$ such that the 1-statement in (4) of Theorem 1 holds with probability $1 - o(n^{-\tau})$.*

3.4 Deferred proof of Lemma 8

For completeness, we now give the routine proof of Lemma 8 deferred from Section 3.1.

Proof of Lemma 8. (i): Set $\Psi_{J,H} := n^{v_H - v_J} p^{e_H - e_J}$. In the case $v_J = v_H$, for any $\beta > 0$ satisfying $1/d(G,H) > 2\beta$ we have $\Psi_{J,H} = p^{e_H - e_J} \ll n^{-(e_H - e_J)\beta} \leqslant n^{-\beta}$. Thus we can henceforth assume $v_J < v_H$. Since G is an induced subgraph of H and thus of J, we also have $v_G < v_J$. Since (G,H) is strictly balanced we have $d(G,J) < d(G,H)$, which implies

$$d(J,H) = \frac{(e_H - e_G) - (e_J - e_G)}{(v_H - v_G) - (v_J - v_G)} = \frac{(v_H - v_G)d(G,H) - (v_J - v_G)d(G,J)}{(v_H - v_G) - (v_J - v_G)} > d(G,H). \tag{50}$$

Hence $1/d(G,H) > 1/d(J,H) + 2\beta$ for $\beta > 0$ sufficiently small, so that $p = O(n^{-1/d(G,H)+\beta}) \ll n^{-1/d(J,H)-\beta}$. Observe that $e_H > e_J$, since otherwise $e_H = e_J$ and $v_H > v_J$ imply $d(G,J) > d(G,H)$, contradicting that (G,H) is strictly balanced. Hence $\Psi_{J,H} = (n^{1/d(J,H)}p)^{e_H - e_J} \ll n^{-\beta}$, completing the proof of (13).

(ii): Assume the contrary. Then we can split $V(H) \setminus V(G)$ into two nonempty sets V_1 and V_2 such that there are no edges between V_1 and V_2. Writing $H_i := H[V(G) \cup V_i]$, we readily obtain

$$d(G,H) = \frac{e_H - e_G}{v_H - v_G} = \frac{\sum_{i \in [2]}(e_{H_i} - e_G)}{\sum_{i \in [2]}(v_{H_i} - v_G)} = \frac{\sum_{i \in [2]}(v_{H_i} - v_G)d(G,H_i)}{\sum_{i \in [2]}(v_{H_i} - v_G)} \leqslant \max_{i \in [2]} d(G,H_i).$$

Since (G,H) is strictly balanced we have $d(G,H_i) < d(G,H)$, yielding the desired contradiction. □

4 No grounded primals case (Theorem 3)

In this section we prove Theorem 3 by focusing on a maximal primal subgraph $J_{\max}$ of (G,H); we remark that $J_{\max}$ is in fact unique (the union of all primal subgraphs), but we do not need this. Our arguments hinge on the basic observation that, since $J_{\max}$ is by assumption not grounded (i.e., there are no edges between $V(G)$ and $V(J_{\max}) \setminus V(G)$), extension counts $X_{G,J_{\max}}(\mathbf{x})$ are essentially the same as the number of *unrooted* copies of the graph $K := J_{\max} - V(G)$, where the vertices of G are deleted from $J_{\max}$.

For the 1-statement this heuristically means that if $X_{G,J_{\max}}(\mathbf{x})$ is concentrated for *some* $\mathbf{x}$, then $X_{G,J_{\max}}(\mathbf{x})$ is concentrated for *all* $\mathbf{x}$ (the reason being that not too many copies of K can overlap with any root $\mathbf{x}'$, see Lemma 16 below). Furthermore, using Theorem 7 (i) it turns out that whp each copy of $J_{\max}$ extends to the 'right' number of H-copies (here the crux will be that $\Phi_{J_{\max},H} = n^{\Omega(1)}$ follows from Remark 1 (iv) and Lemma 14 below). Combining these two estimates then allows us to deduce that whp $X_{G,H}(\mathbf{x})$ is concentrated for all $\mathbf{x}$; see Section 4.3 for the details.

For the 0-statement we shall proceed similarly, the main difference is that, for a *fixed* $\mathbf{x}$, we start by arguing that $X_{G,J_{\max}}(\mathbf{x})$ is not concentrated, i.e., whp far away from its expected value. This allows us to deduce that $\mathbf{x}$ has whp the wrong number of (G,H)-extensions (since by Theorem 7 (i) whp each copy of $J_{\max}$ again extends to the right number of copies of H); see Section 4.2 for the details.

4.1 Setup and technical preliminaries

In the upcoming arguments it will, as in [31], often be convenient to treat extensions as sequences of vertices. Given a rooted graph (G, H) with labeled vertices $V(G) = \{1, \ldots, v_G\}$ and $V(H) \setminus V(G) = \{v_G + 1, \ldots, v_H\}$, *an ordered* (G, H)*-extension of* $\mathbf{x} = (x_1, \ldots, x_{v_G}) \in [n]_{v_G}$ is a sequence $\mathbf{y} = (y_{v_G+1}, \ldots, y_{v_H})$ of distinct vertices from $[n] \setminus \{x_1, \ldots, x_{v_G}\}$ such that the injection which maps each vertex $j \in V(G)$ onto x_j and each vertex $i \in V(H) \setminus V(G)$ onto y_i, also maps every edge $f \in E(H) \setminus E(G)$ onto an edge. Given a root $\mathbf{x} \in [n]_{v_G}$, let $Y_{G,H}(\mathbf{x})$ denote the number of ordered (G, H)-extensions of $\mathbf{x}$ in $\mathbb{G}_{n,p}$. Note that

$$\nu_{G,H} := \mathbb{E}Y_{G,H}(\mathbf{x}) = (n - v_G)(n - v_G - 1) \cdots (n - v_H + 1) \cdot p^{e_H - e_G} \tag{51}$$

does not depend on the particular choice of $\mathbf{x}$. Let $\mathrm{aut}(G, H)$ denote the number of automorphisms of H that fix the set $V(G)$. Since each extension corresponds to $\mathrm{aut}(G, H)$ many ordered extensions, we obtain

$$Y_{G,H}(\mathbf{x}) = \mathrm{aut}(G, H) \cdot X_{G,H}(\mathbf{x}), \tag{52}$$

$$\nu_{G,H} = \mathrm{aut}(G, H) \cdot \mu_{G,H}, \tag{53}$$

where $\mu_{G,H} = \mathbb{E}X_{G,H}(\mathbf{x})$ is defined as in (1). One further useful elementary observation is that, for any induced $G \subseteq J \subseteq H$, we have

$$\nu_{G,J} \cdot \nu_{J,H} = \nu_{G,H}. \tag{54}$$

Our arguments will also exploit the following technical property of maximal primal subgraphs.

Lemma 14. *If* $J_{\max} \subsetneq H$ *is a maximal primal of the rooted graph* (G, H)*, then* $m(J_{\max}, H) < m(G, H)$.

Proof. Fix $J_{\max} \subsetneq J \subseteq H$. Using maximality of $J_{\max} \supsetneq G$, we infer $d(G, J) < m(G, H)$ and $d(G, J_{\max}) = m(G, H)$. Proceeding analogously to inequality (50), it routinely follows that

$$d(J_{\max}, J) = \frac{(v_J - v_G)d(G, J) - (v_{J_{\max}} - v_G)d(G, J_{\max})}{(v_J - v_G) - (v_{J_{\max}} - v_G)} < m(G, H),$$

which completes the proof by maximizing over all feasible J. □

4.2 The 0-statement

As discussed, for the 0-statement of Theorem 3 the core idea is to show that $X_{G,J_{\max}}(\mathbf{x})$ is not concentrated for some $\mathbf{x} \in [n]_{v_G}$, and that $X_{J_{\max},H}(\mathbf{y})$ is concentrated for all $\mathbf{y} \in [n]_{v_{J_{\max}}}$, see (58)–(59) below.

Proof of the 0*-statement of Theorem 3.* Assuming $\varepsilon \geqslant n^{-\alpha}$ with $\alpha < 1/2$ (as we may), we have $\varepsilon^2\Phi_{G,H} = \Omega(n^{1-2\alpha}p^{e_H - e_G}) \gg p^{e_H - e_G}$, so the assumption $\varepsilon^2\Phi_{G,H} \to 0$ implies $p \to 0$ and thus $1 - p = \Theta(1)$. Since (G, H) has no grounded primals, the desired 0-statement now follows by combining the conclusions of Theorem 7 (ii) for the cases $\Phi_{G,H} \to 0$ and $\Phi_{G,H} \to \infty$ with the conclusion of Lemma 15 below for $\Phi_{G,H} \asymp 1$ (formally using, as usual, the subsubsequence principle [15, Section 1.2]). □

Lemma 15. *Let* (G, H) *be a rooted graph with no grounded primal subgraphs. Then, for all* $p = p(n) \in [0, 1]$ *and* $\varepsilon = \varepsilon(n) \in (0, 1]$ *with* $\Phi_{G,H} \asymp 1$ *and* $\varepsilon \to 0$,

$$\lim_{n\to\infty} \mathbb{P}\Big(\max_{\mathbf{x}\in[n]_{v_G}} |X_{\mathbf{x}} - \mu| \geqslant \varepsilon\mu\Big) = 1. \tag{55}$$

Proof. Note that by increasing ε if necessary, we may henceforth assume $\varepsilon \geqslant n^{-\alpha}$ for any constant $\alpha > 0$ (since increasing ε can only decrease the probability on the left-hand side of (55) above). Let $J_{\max}$ be a maximal primal subgraph of (G, H). By Remark 1 (ii)–(iii), the assumption $\Phi_{G,H} \asymp 1$ implies

$$\mu_{G,J_{\max}} \asymp 1, \tag{56}$$

$$p = \Omega\big(n^{-1/m(G,H)}\big). \tag{57}$$

Turning to the details, we start with the claim that, whp,

$$\max_{\mathbf{x}\in[n]_{v_G}} |X_{G,J_{\max}}(\mathbf{x}) - \mu_{G,J_{\max}}| > 3\varepsilon\mu_{G,J_{\max}}, \tag{58}$$

$$\max_{\mathbf{y}\in[n]_{v_{J_{\max}}}} |X_{J_{\max},H}(\mathbf{y}) - \mu_{J_{\max},H}| < \tfrac{1}{2}\varepsilon\mu_{J_{\max},H}. \tag{59}$$

To show that this claim implies the desired 0-statement, we consider ordered extensions and note that multiplying (58) and (59) by $\mathrm{aut}(G, J_{\max})$ and $\mathrm{aut}(J_{\max}, H)$, respectively, we can replace X by Y and μ by ν, cf. (52) and (53). Observe that each ordered (G, H)-extension corresponds to a unique pair of extensions: one of $\mathbf{x}$ with respect to $(G, J_{\max})$ and one of $\mathbf{y}$ (which consists of $\mathbf{x}$ plus the vertices of the first extension) with respect to $(J_{\max}, H)$. Consequently, recalling the identity (54), inequalities (58)–(59) imply that there is $\mathbf{x} \in [n]_{v_G}$ such that either

$$Y_{G,H}(\mathbf{x}) > (1 + 3\varepsilon)\nu_{G,J_{\max}} \cdot (1 - \varepsilon/2)\nu_{J_{\max},H} > (1 + \varepsilon)\nu_{G,H} \tag{60}$$

or

$$Y_{G,H}(\mathbf{x}) < (1 - 3\varepsilon)\nu_{G,J_{\max}} \cdot (1 + \varepsilon/2)\nu_{J_{\max},H} < (1 - \varepsilon)\nu_{G,H}, \tag{61}$$

which in view of (52) and (53) establishes the desired 0-statement (after rescaling by $\mathrm{aut}(G, H)$).

It remains to show that (58) and (59) hold whp, and we start with (58). Consider the unrooted graph $K := J_{\max} - V(G)$, where the vertices of G are deleted from $J_{\max}$. By construction, we have $v_K = v_{J_{\max}} - v_G$. Since $J_{\max}$ is not grounded, we also have $e_K = e_{J_{\max}} - e_G$. Using (56) we infer

$$\mu_K \asymp n^{v_K}p^{e_K} = n^{v_{J_{\max}}-v_G}p^{e_{J_{\max}}-e_G} \asymp \mu_{G,J_{\max}} \asymp 1, \tag{62}$$

which by Markov's inequality implies that the number of K-copies is whp at most $n/(2v_K)$, say (with room to spare). This means that either (i) there are no K-copies, in which case $X_{G,J_{\max}}(\mathbf{x}) = 0$ for all $\mathbf{x} \in [n]_{v_G}$, or (ii) the K-copies span at most $n/2$ vertices, in which case there is one $\mathbf{x}_1 \in [n]_{v_G}$ that is disjoint from all K-copies and another set $\mathbf{x}_2 \in [n]_{v_G}$ that intersects at least one K-copy, so that $X_{G,J_{\max}}(\mathbf{x}_1) = X_K > X_{G,J_{\max}}(\mathbf{x}_2)$. In both cases it follows that (58) holds whp, since (56) and $\varepsilon \to 0$ imply that the interval $(1 \pm 3\varepsilon)\mu_{G,J_{\max}}$ does not contain zero, and moreover contains at most one integer.

Turning to (59), note that (59) holds trivially when $J_{\max} = H$. Otherwise $m(J_{\max}, H) < m(G, H)$ by Lemma 14, so that (57) implies $p = \Omega(n^{\gamma-1/m(J_{\max},H)})$ for some constant $\gamma > 0$. Using Remark 1 (iv), it follows that $\Phi_{J_{\max},H} = \Omega(n^{\gamma})$. Assuming $\varepsilon \geqslant n^{-\alpha}$ with $\alpha < \gamma/2$ (as we may), we infer $\varepsilon^2\Phi_{J_{\max},H} = \Omega(n^{\gamma/2-\alpha}) = n^{\Omega(1)}$. Applying Theorem 7 (i) with $t = \frac{1}{2}\varepsilon\mu_{J_{\max},H}$, now (59) holds whp. $\square$

4.3 The 1-statement

As discussed, for the 1-statement of Theorem 3 we rely on the fact that no vertex is contained in too many copies of the (unrooted) graph $J_{\max} - V(G)$, which is formalized by Lemma 16 below. As usual, given a graph K with $v_K \geqslant 1$, subgraphs $J \subseteq K$ with $v_J \geqslant 1$ that maximize the density $d_J := d(\emptyset, J) = e_J/v_J$ are called *primal* (consistently with rooted graphs terminology), and K is called *balanced* when K itself is primal.

Lemma 16. *Let K be a balanced graph with $e_K \geqslant 1$. There are constants $\beta, C > 0$ such that, for all $p = p(n) \in [0, 1]$ with $n^{\beta-1/d_K} \ll p = O(n^{\beta-1/d_K})$, in $\mathbb{G}_{n,p}$ whp every vertex $x \in [n]$ is contained in at most $C\lambda^{v_K-v_{G_{\min}}}$ copies of K, where $\lambda := np^{d_K}$ and $G_{\min} \subseteq K$ is a primal subgraph with the smallest number of vertices.*

We defer the density based proof to Appendix B (which is rather tangential to the main argument here), and now use Lemma 16 to prove the desired 1-statement of Theorem 3.

Proof of the 1-statement of Theorem 3. The assumptions $\varepsilon \leqslant 1$ and $\varepsilon^2\Phi_{G,H} \to \infty$ imply $\Phi_{G,H} \to \infty$, so Remark 1 (i) implies $p \gg n^{-1/m(G,H)}$. If $\varepsilon^2\Phi_{G,H} = n^{\Omega(1)}$, then the desired 1-statement follows from Theorem 7 (i), so we may further assume $\varepsilon^2\Phi_{G,H} \leqslant n^c$ for any constant $c > 0$ of our choice, which together with the assumption $\varepsilon \geqslant n^{-\alpha}$ implies $\Phi_{G,H} \leqslant n^{c+2\alpha}$. Using the contrapositive of Remark 1 (iv), by choosing $\alpha, c > 0$ sufficiently small (as we may) we thus henceforth can assume

$$n^{-1/m(G,H)} \ll p \ll n^{\beta-1/m(G,H)}, \tag{63}$$

where the constant $\beta > 0$ is as given by Lemma 16.

Turning to the details, let $J_{\max}$ be a maximal primal subgraph of (G, H). For convenience we use ordered extensions, as before. Note that $Y_{G,J_{\max}}(\mathbf{x})$ is the number of (unrooted) copies of graph $K := J_{\max} - V(G)$ that are disjoint from $\mathbf{x}$. For any vertex $x \in [n]$, let $Z_K(x)$ denote the number of copies of K containing x. We fix some $\mathbf{x}' \in [n]_{v_G}$, and start with the claim that there exists a constant $D > 0$ such that, whp,

$$|Y_{G,J_{\max}}(\mathbf{x}') - \nu_{G,J_{\max}}| < \tfrac{1}{8}\varepsilon\nu_{G,J_{\max}}. \tag{64}$$

$$\max_{\mathbf{y}\in[n]_{v_{J_{\max}}}} |Y_{J_{\max},H}(\mathbf{y}) - \nu_{J_{\max},H}| < \tfrac{1}{2}\varepsilon\nu_{J_{\max},H}, \tag{65}$$

$$\max_{x\in[n]} Z_K(x) \leqslant D\frac{\varepsilon\nu_{G,J_{\max}}}{\varepsilon^2\Phi_{G,H}}. \tag{66}$$

We now show that this claim implies the desired 1-statement. In view of (64), the first step is to use (66) to show that $Y_{G,J_{\max}}(\mathbf{x})$ is also concentrated for the remaining roots $\mathbf{x} \neq \mathbf{x}'$. Specifically, using (66) to bound the number of $(G, J_{\max})$-extensions of $\mathbf{x}$ that overlap with $\mathbf{x}'$ (and those of $\mathbf{x}'$ overlapping with $\mathbf{x}$), in view of the assumption $\varepsilon^2\Phi_{G,H} \to \infty$ it follows that, for every $\mathbf{x} \in [n]_{v_G}$,

$$|Y_{G,J_{\max}}(\mathbf{x}) - Y_{G,J_{\max}}(\mathbf{x}')| \leqslant O\Big(\sum_{x\in\mathbf{x}\cup\mathbf{x}'} Z_K(x)\Big) \ll \tfrac{1}{8}\varepsilon\nu_{G,J_{\max}}.$$

Together with (64) this implies that, say,

$$\max_{\mathbf{x}\in[n]_{v_G}} |Y_{G,J_{\max}}(\mathbf{x}) - \nu_{G,J_{\max}}| < \tfrac{1}{4}\varepsilon\nu_{G,J_{\max}}. \tag{67}$$

The second step exploits that by (65) each copy of $J_{\max}$ extends to the 'right' number of copies of H. Indeed, with analogous reasoning as for (60)–(61) from Section 4.2, by combining (65) with (67) it now follows (in view of (54)) that

$$\max_{\mathbf{x}\in[n]_{v_G}} Y_{G,H}(\mathbf{x}) < (1+\varepsilon/4)\nu_{G,J_{\max}} \cdot (1+\varepsilon/2)\nu_{J_{\max},H} < (1+\varepsilon)\nu_{G,H},$$

and similarly,

$$\min_{\mathbf{x}\in[n]_{v_G}} Y_{G,H}(\mathbf{x}) > (1-\varepsilon)\nu_{G,H},$$

which in view of (52)–(53) establishes the 1-statement of Theorem 3 (by rescaling by $\mathrm{aut}(G, H)$).

It remains to show that (64)–(66) hold whp, and we start with (64). Since $\Phi_{G,J_{\max}} \geqslant \Phi_{G,H}$ by definition, using Chebyshev's inequality together with the variance estimate (10) and $\varepsilon^2\Phi_{G,H} \to \infty$, it follows that

$$\mathbb{P}\left(|X_{G,J_{\max}}(\mathbf{x}') - \mu_{G,J_{\max}}| \geqslant \tfrac{1}{8}\varepsilon\mu_{G,J_{\max}}\right) \leqslant \frac{\mathrm{Var}\, X_{G,J_{\max}}(\mathbf{x})}{(\varepsilon/8)^2\mu_{G,J_{\max}}^2} \asymp \frac{1-p}{\varepsilon^2\Phi_{G,J_{\max}}} \leqslant \frac{1}{\varepsilon^2\Phi_{G,H}} \to 0,$$

which in view of (52)–(53) then implies that (64) holds whp (by rescaling by $\mathrm{aut}(G, J_{\max})$).

Next we establish (65). Note that the proof of (59) only relies on (57) (which here holds by (63)), and that we may assume $\varepsilon \geqslant n^{-\alpha}$ for sufficiently small $\alpha > 0$. Hence by the same argument as for (59), in view of (52)–(53) it follows that (65) holds whp (after rescaling by $\mathrm{aut}(J_{\max}, H)$).

Finally, we turn to the auxiliary estimate (66). Note that every subgraph $J \subseteq K$ with $v_J \geqslant 1$ satisfies $d_J = d(G, G \cup J)$. Hence J is primal for K if and only if $G \cup J$ is primal for $(G, J_{\max})$. Since $J_{\max} = G \cup K$ is primal for (G, H), it follows that K is balanced, with $d_K = d(G, J_{\max}) = m(G, H)$. Using assumption (63), we thus have $n^{-1/d_K} \ll p \ll n^{\beta-1/d_K}$. Invoking Lemma 16, there is a constant $C > 0$ such that, whp,

$$\max_{x\in[n]} Z_K(x) \leqslant C\lambda^{v_K - v_{G_{\min}}},$$

where $G_{\min}$ is a primal subgraph of K with the smallest number of vertices. In particular, we have $d_K = d_{G_{\min}}$. Since $J_{\max}$ is a vertex-disjoint union of the graphs K and G, using $G_{\min} \subseteq K$ we infer that $e_{G_{\min}} = e_{G\cup G_{\min}} - e_G$ and $v_{G_{\min}} = v_{G\cup G_{\min}} - v_G$. Recalling $\lambda = np^{d_K} = np^{d_{G_{\min}}}$, it now follows analogously to (62) that

$$\lambda^{v_K - v_{G_{\min}}} = \frac{n^{v_K}p^{e_K}}{n^{v_{G_{\min}}}p^{e_{G_{\min}}}} \asymp \frac{\mu_K}{\mu_{G_{\min}}} \asymp \frac{\mu_{G,J_{\max}}}{\mu_{G,G\cup G_{\min}}} \leqslant \frac{\mu_{G,J_{\max}}}{\Phi_{G,H}},$$

which together with $1 \leqslant 1/\varepsilon = \varepsilon/\varepsilon^2$ and (53) completes the proof of (66) for suitable $D > 0$. □

5 Further cases

5.1 Unique and grounded primal case (Theorem 2)

In this section we prove Theorem 2 by adapting the arguments from Section 4 (focusing on the unique primal $J = J_{\max}$). The key difference is that here we can use the 0- and 1-statements of our main result Theorem 1 to deduce that $X_{G,J}(\mathbf{x})$ is not concentrated for some $\mathbf{x}$, or concentrated for all $\mathbf{x}$, respectively. This then allows us to prove the desired 0- and 1-statements, since each copy of J again extends to the 'right' number of copies of H (by Theorem 7 (i), as in Section 4); see (69)–(70) and (73) below.

Proof of Theorem 2. If $\Phi_{G,H} \to 0$, then the 0-statement holds by Theorem 7 (ii). Therefore we henceforth can assume $\Phi_{G,H} = \Omega(1)$, which by Remark 1 (ii) is equivalent to

$$p = \Omega\big(n^{-1/m(G,H)}\big). \tag{68}$$

Note that the proof of (59) relies only on (57) (which is the same as (68)), the fact that $J_{\max}$ is the maximal primal (which also holds trivially for J in the current setting), and that we may assume $\varepsilon \geqslant n^{-\alpha}$ for sufficiently small $\alpha > 0$ (which we may also assume here). Hence by the same argument as for (59), after rescaling by $\mathrm{aut}(J,H)$ (see (52)–(53)) we here obtain that, whp,

$$\max_{\mathbf{y} \in [n]_{v_J}} |Y_{J,H}(\mathbf{y}) - \nu_{J,H}| < \tfrac{1}{2}\varepsilon \nu_{J,H}. \tag{69}$$

We start with the 1-statement. Since $\mu_{G,J} \geqslant \Phi_{G,H}$ by definition, the assumption $\varepsilon^2\Phi_{G,H} \geqslant C \log n$ implies $\varepsilon^2\mu_{G,J} \geqslant C \log n$. By uniqueness of the primal J, the rooted graph (G,J) is strictly balanced. Therefore (4) of Theorem 1 implies (after rescaling by $\mathrm{aut}(G,J)$) that, for suitable $\alpha, C > 0$, whp,

$$\max_{\mathbf{x} \in [n]_{v_G}} |Y_{G,J}(\mathbf{x}) - \nu_{G,J}| < \tfrac{1}{4}\varepsilon \nu_{G,J}. \tag{70}$$

The 1-statement of Theorem 2 now follows from (69) and (70) by exactly the same reasoning with which (65) and (67) from Section 4.3 implied the 1-statement of Theorem 3.

We now turn to the 0-statement. We again plan to apply (4) of Theorem 1 to the strictly balanced rooted graph (G,J), for which we need to check that the assumption $\varepsilon^2\Phi_{G,H} \leqslant c \log n$ implies the required condition $\varepsilon^2\mu_{G,J} \leqslant c \log n$. We will do this by showing that $\Phi_{G,H} = \mu_{G,J}$ for n large enough. First, note that the assumptions $\varepsilon \geqslant n^{-\alpha}$ and $\varepsilon^2\Phi_{G,H} \leqslant c \log n$ imply $\Phi_{G,H} = O(n^{2\alpha} \log n)$. By (68) and the contrapositive of Remark 1 (iv) we can thus assume that, say,

$$p \asymp n^{\theta - 1/m(G,H)} \qquad \text{with} \qquad \theta = \theta(n,p) \in [0, 3\alpha]. \tag{71}$$

Since the primal J is unique, we have $d(G,J) = m(G,H)$, and $d(G,K) < m(G,H)$ when $G \subsetneq K \subseteq H$ satisfies $J \neq K$. Hence there exists a constant $\gamma = \gamma(G,J,H) > 0$ such that, for any $G \subsetneq K \subseteq H$,

$$\mu_{G,K} \asymp \Big(np^{d(G,K)}\Big)^{v_K - v_G} \asymp \Big(n^{1 - \frac{d(G,K)}{m(G,H)} + \theta d(G,K)}\Big)^{v_K - v_G} = \begin{cases} \Omega(n^{\gamma}) & \text{if } K \neq J, \\ O(n^{3\alpha(e_J - e_G)}) & \text{if } K = J. \end{cases} \tag{72}$$

By taking $\alpha > 0$ small enough, it follows that $\Phi_{G,H} = \mu_{G,J}$ for n large enough, which (as discussed) establishes $\varepsilon^2\mu_{G,J} \leqslant c \log n$. Therefore (4) of Theorem 1 implies (after rescaling by $\mathrm{aut}(G,J)$) that, whp,

$$\max_{\mathbf{x} \in [n]_{v_G}} |Y_{G,J}(\mathbf{x}) - \nu_{G,J}| \geqslant 3\varepsilon \nu_{G,J}. \tag{73}$$

The 0-statement of Theorem 2 now follows from (73) and (69) by the same (routine) reasoning with which (58)–(59) from Section 4.2 implied the 0-statement of Theorem 3. □

Remark 4 (Theorem 2: stronger 1-statement)**.** *The above proof yields, in view of Remarks 2–3, the following stronger conclusion: for any fixed $\tau > 0$ there is a constant $C = C(\tau, G, H) > 0$ such that the 1-statement in (7) of Theorem 2 holds with probability $1 - o(n^{-\tau})$.*

5.2 Strictly balanced and ungrounded case (Theorem 1)

In this section we prove the threshold (5) of Theorem 1 (ii) for strictly balanced rooted graphs (G, H) that are not grounded, which turns out to be a simple corollary of Theorem 3. The crux is that, by decreasing $\alpha > 0$ (if necessary), we can ensure that the 0- and 1-statement conditions in (5) and (8) coincide.

Proof of (5) *of Theorem 1.* Recall that $\mu = \mu_{G,H}$ and $\Phi = \Phi_{G,H}$ are defined in (1) and (6), respectively. By assumption the unique primal H is not grounded, so Theorem 3 applies. Decreasing the constant $\alpha > 0$ from Theorem 3, we can assume that $\beta \geqslant 3\alpha$, where $\beta > 0$ is the constant given by Lemma 8 (i). We now distinguish two ranges of $p = p(n)$. First, when $p \leqslant n^{-1/d(G,H)+\beta}$, then (13) from Lemma 8 implies that $\mu = \Phi$ for n large enough (since (13) implies $\mu_{G,H}/\mu_{G,J} \asymp n^{v_H - v_J} p^{e_H - e_J} \ll 1$ for all $G \subseteq J \subsetneq H$ with $e_J > e_G$). Second, when $p \geqslant n^{-1/d(G,H)+\beta} \geqslant n^{-1/m(G,H)+3\alpha}$, then $\varepsilon \geqslant n^{-\alpha}$ and Remark 1 (iv) imply that $\min\{\varepsilon^2\mu, \varepsilon^2\Phi\} \geqslant n^{-2\alpha} \cdot \Phi = \Omega(n^{\alpha}) \to \infty$. Since in both ranges the 0- and 1-statement conditions in (5) and (8) coincide, it follows that Theorem 3 implies (5). □

6 Cautionary examples (Proposition 5 and 6)

In this section we prove Propositions 5–6 for the rooted graphs (e) and (f) depicted in Figure 2. The proof idea for Proposition 5 is to proceed in two rounds for a fixed vertex x: using Theorem 1 we first find about μ_{G,K_4} many (G, K_4)-extensions of $\mathbf{x} = (x)$, which we then extend to about $\mu_{G,H}$ many (G, H)-extensions of $\mathbf{x}$. The crux is that most of the relevant (K_4, H)-extensions from the second round evolve nearly independently, which ultimately allows us to surpass the conditions of Spencer's result (3) and Theorem 1 for (K_4, H).

Proof of Proposition 5. Recalling $\omega = np^2$, by assumption we have $\varepsilon^2\mu_{G,K_4} \asymp \varepsilon^3\omega^3 \gg \log n$ and $\varepsilon^2\mu_{K_4,H} \leqslant \mu_{K_4,H} \asymp \omega \ll \log n$, which readily implies $\log\omega \asymp \log\log n$ and $p = n^{-1/2+o(1)}$. Now it is not difficult to verify that $\Phi_{G,H} \asymp \mu_{G,K_4} \asymp \omega^3$ (either directly, or similarly as for (72) from Section 5.1). Turning to the details of the 1-statement, we start with the auxiliary claim that, whp, for each vertex x the following event $\mathcal{P}_x$ holds:

(i) The vertex-neighbourhood Γ_x of x has size $|\Gamma_x| \leqslant 9np$.
(ii) The collection $\mathcal{T}_x$ of all triangles spanned by Γ_x has size $|\mathcal{T}_x| = (1 \pm \varepsilon/9)\binom{n-1}{3}p^6$.
(iii) Every vertex $y \in \Gamma_x$ is contained in at most $D := 15$ triangles from $\mathcal{T}_x$.

Indeed, invoking the 1-statement of Theorem 1 with H equal to K_4 and G being the root vertex v, from $\varepsilon^2\mu_{G,K_4} \asymp \varepsilon^2\omega^3 \gg \log n$ it follows that, whp, (ii) holds for all vertices x. Since $np = n^{1/2+o(1)} \gg \log n$, using standard Chernoff bounds it is routine to see that, whp, (i) holds for all vertices x. We claim that if (iii) fails for some $y \in \Gamma_x$, then there are 4 triangles in $\mathcal{T}_x$ containing y that form either a flower (share no vertices other than y) or a book (all contain an edge yz for some $z \in \Gamma_x$): to see this, note that if we assume the contrary, then for a maximal flower (with at most 3 triangles) each edge of it is contained in at most 2 other $\mathcal{T}_x$-triangles, whence there are at most $3 + 6 \cdot 2 = 15$ triangles in $\mathcal{T}_x$ containing y. The probability that there is either a 4-flower or 4-book with all vertices connected to some extra vertex x is at most $n^{10}p^{21} + n^7p^{16} = n^{-1/2+o(1)} \to 0$. It follows that, whp, properties (i)–(iii) hold for all vertices x, establishing the claim.

We now fix a root vertex x, and expose the edges of $\mathbb{G}_{n,p}$ in two rounds: in the first round we expose all edges incident to x and all edges inside Γ_x, and then in the second round we expose all remaining edges. We henceforth condition on the outcome of the first exposure round, and assume that $\mathcal{P}_x$ holds. As usual, to avoid clutter we shall omit this conditioning from our notation. Given distinct vertices $a, b \in \Gamma_x$, let $Y_{a,b}$ denote the number of common neighbours of a and b in $[n] \setminus (\{x\} \cup \Gamma_x)$. Note that $\varepsilon\omega \gg \sqrt{\log n/\omega} \gg 1$ by assumption. Since, by (ii), $|\mathcal{T}_x| \asymp n^3p^6 = \omega^3 \ll \varepsilon\omega^4 \asymp \varepsilon\mu$, using (iii) it is not difficult to see that

$$Z_x := \sum_{abc \in \mathcal{T}_x} (Y_{a,b} + Y_{b,c} + Y_{a,c}) \qquad \text{satisfies} \qquad \big|X_{(x)} - Z_x\big| \leqslant 3|\mathcal{T}_x| \cdot D \ll \varepsilon\mu/2. \tag{74}$$

Using (i) and $\varepsilon\omega \gg 1$ (see above) we infer $1 + |\Gamma_x| \leqslant 10np = n^{1/2+o(1)} \ll n/\omega \ll \varepsilon n$, and together with (ii) it then follows that, say,

$$\mathbb{E}Z_x = 3|\mathcal{T}_x| \cdot \big(n - 1 - |\Gamma_x|\big)p^2 = (1 \pm \varepsilon/8)\mu. \tag{75}$$

In (74) we now write each $Y_{a',b'}$ as a sum of indicators of length 2 paths, which enables us to estimate the lower tail of Z_x via Janson's inequality. By distinguishing between pairs of edge-overlapping paths that share one or two endpoints, using (iii) it is standard to see that the relevant Δ term is at most $\mathbb{E}Z_x \cdot (2Dp+2D) = O(\mathbb{E}Z_x)$, say. Using $\varepsilon^2 \mathbb{E}Z_x \asymp \varepsilon^2 \omega^4 \gg \log n$, by invoking [25, Theorem 1] it then follows that

$$\mathbb{P}(Z_x \leqslant (1-\varepsilon/8)\mathbb{E}Z_x) \leqslant \exp\bigl(-\Omega(\varepsilon^2 \mathbb{E}Z_x)\bigr) = o(n^{-1}). \tag{76}$$

Using (iii) we also see that any path shares an edge with a total of at most $2D = O(1)$ paths, which enables us to estimate the upper tail of Z_v via concentration inequalities for random variables with 'controlled dependencies'. In particular, by invoking [15, Proposition 2.44] (see also [37, Theorem 9]) it follows that

$$\mathbb{P}(Z_x \geqslant (1+\varepsilon/8)\mathbb{E}Z_x) \leqslant \exp\bigl(-\Omega(\varepsilon^2 \mathbb{E}Z_x)\bigr) = o(n^{-1}). \tag{77}$$

To sum up, (74)–(77) and $1-\varepsilon/2 < (1 \pm \varepsilon/8)^2 < 1+\varepsilon/2$ imply $\mathbb{P}(|X_{(x)} - \mu| \geqslant \varepsilon\mu \mid \mathcal{P}_x) = o(n^{-1})$, which readily completes the proof of the desired 1-statement (since, whp, $\mathcal{P}_x$ holds for all n vertices x). ☐

The proof idea for Proposition 6 is to find a copy of K_4 with an edge that is contained in extremely many triangles. To this end we proceed in two steps, inspired by [29, Lemma 3]: in the first step we find $\Theta(n)$ many vertex-disjoint copies of K_4, and in the second step we then find the desired edge contained in many triangles.

Proof of Proposition 6. Note that $\mu \asymp \omega^5$. As in the proof of Proposition 5, we again have $\log\omega \asymp \log\log n$ and $\Phi_{G,H} \asymp \mu_{G,K_4} \asymp \omega^3$, so $\varepsilon^2\Phi_{G,H} \asymp \varepsilon^2 \mu_{G,K_4} \gg \log n$ by assumption. Noting $0.39 < 2/5$, we define

$$z := \Bigl\lceil 2\bigl((1+\varepsilon)\mu\bigr)^{1/2}\Bigr\rceil \asymp \omega^{5/2} = o(\log n/\log\omega). \tag{78}$$

Turning to the details of the desired 0-statement, let Y_{K_4} denote the size of the largest collection of vertex-disjoint copies of K_4 spanned by the vertices in $W := \{1, \ldots, \lfloor n/2 \rfloor\}$. It is routine to check that the minimum of $|W|^{v_G} p^{e_G}$, taken over all $G \subseteq K_4$ with $v_G \geqslant 1$, equals $|W| \approx n/2$ for n large enough. Since the induced subgraph of $\mathbb{G}_{n,p}$ spanned by W has the same distribution as $\mathbb{G}_{|W|,p}$, by invoking [15, Theorem 3.29] there is a constant $c > 0$ such that

$$\mathbb{P}(Y_{K_4} \geqslant cn) = 1 - o(1). \tag{79}$$

We now condition on the edges spanned by W, and assume that $Y_{K_4} \geqslant cn$. To avoid clutter, we shall again omit this conditioning from our notation (as in the proof of Proposition 6). We henceforth fix $\lceil cn \rceil$ vertex-disjoint copies of K_4 spanned by W, and from the i-th such copy we pick an edge $\{v_i, w_i\}$ and a further vertex $x_i \notin \{v_i, w_i\}$. Defining Z_i as the number of vertices in $[n] \setminus W$ that are common neighbours of v_i and w_i, using $\binom{m}{z} \geqslant (m/z)^z$ for $m \geqslant z$ together with $np^2 = \omega = o(z)$ and (78), it routinely follows that

$$\mathbb{P}(Z_i \geqslant z) \geqslant \binom{\lceil n/2 \rceil}{z} p^{2z}\bigl(1-p^2\bigr)^{\lceil n/2 \rceil - z} \geqslant \Bigl(\frac{np^2}{2z}\Bigr)^z e^{-np^2} = e^{-\Theta(z \log\omega)} \geqslant n^{-o(1)}.$$

Note that $Z_i \geqslant z$ implies $X_{(x_i)} \geqslant \binom{z}{2} \geqslant \bigl(z/2\bigr)^2 \geqslant (1+\varepsilon)\mu$. Since the random variables Z_i depend on disjoint sets of independent edges, it then follows that

$$\mathbb{P}(\max_{x \in [n]} X_{(x)} < (1+\varepsilon)\mu) \leqslant \mathbb{P}(\max_{1 \leqslant i \leqslant \lceil cn \rceil} Z_i < z) = \prod_{1 \leqslant i \leqslant \lceil cn \rceil} \mathbb{P}(Z_i < z) \leqslant \Bigl(1 - n^{-o(1)}\Bigr)^{\lceil cn \rceil} = o(1).$$

Hence $\mathbb{P}(\max_{x \in [n]} X_{(x)} \geqslant (1+\varepsilon)\mu \mid Y_{K_4} \geqslant cn) = 1 - o(1)$, which together with (79) completes the proof. ☐

7 Concluding remarks

The results and problems of this paper can also be viewed through the lens of *extreme value theory*, where a standard goal is to show that a (suitably shifted and normalized) maximum converges to a non-degenerate distribution. To see the connection, note that the proof of Theorem 1 (i) describes an interval on which $\max_{\mathbf{x} \in [n]_{v_G}} X_{\mathbf{x}}$ is whp concentrated. Our setting concerns discrete random variables (which can have complicated behaviour, cf. [10, Section 8.5]), with a correlation structure that seems quite unusual for the field. Hence, as a first step, it would already be interesting to establish a 'law of large numbers' result (even for a restricted class of (G, H), such as strictly balanced ones), which is the content of the following problem.

Problem 2. *Determine for what rooted graphs* (G, H) *and edge probabilities* $p = p(n)$ *there is a sequence* (a_n) *of real positive numbers such that* $(\max_{\mathbf{x}} X_{\mathbf{x}} - \mu)/a_n$ *converges to* 1 *in probability (as* $n \to \infty$*).*

Acknowledgements. We are grateful to the referees for helpful suggestions concerning the presentation.

References

[1] N. Alon and J. Spencer. *The Probabilistic Method*, 4th ed., John Wiley & Sons, Inc., Hoboken, NJ (2016).

[2] J. Baron and J. Kahn. On the cycle space of a random graph. *Random Struct. Alg.* **54** (2019), 39–68.

[3] T. Bohman, A. Frieze, and E. Lubetzky. Random triangle removal. *Adv. Math.* **280** (2015), 379–438.

[4] T. Bohman and P. Keevash. The early evolution of the H-free process. *Invent. Math.* **181** (2010), 291–336.

[5] T. Bohman and P. Keevash. Dynamic concentration of the triangle-free process. *Random Struct. Alg.* **58** (2021), 221–293.

[6] T. Bohman and L. Warnke. Large girth approximate Steiner triple systems. *J. London Math. Soc.* **100** (2019), 895–913.

[7] B. Bollobás. *Random Graphs*, 2nd ed., Cambridge University Press (2001).

[8] R. Boppana and J. Spencer. A useful elementary correlation inequality. *J. Combin. Theory Ser. A* **50** (1989), 305–307.

[9] S. Chatterjee. A general method for lower bounds on fluctuations of random variables. *Ann. Probab.* **47** (2019), 2140–2171.

[10] M.Falk, J. Hüsler, and R. Reiss. *Laws of small numbers: Extremes and rare events.* 3rd ed., Birkhäuser, Basel (2011).

[11] G. Fiz Pontiveros, S. Griffiths, and R. Morris. The triangle-free process and $R(3, k)$. *Mem. Amer. Math. Soc.* **263** (2020), no. 1274.

[12] A. Frieze and M. Karoński. *Introduction to Random Graphs.* Cambridge University Press (2016).

[13] T. Harris. A lower bound for the critical probability in a certain percolation process. *Proc. Cambridge Philos. Soc.* **56** (1960), 13–20.

[14] S. Janson. Poisson approximation for large deviations. *Random Struct. Alg.* **1** (1990), 221–229.

[15] S. Janson, T. Łuczak, and A. Ruciński. *Random Graphs.* Wiley-Interscience, New York (2000).

[16] S. Janson and A. Ruciński. The infamous upper tail. *Random Struct. Alg.* **20** (2002), 317–342.

[17] S. Janson and A. Ruciński. Upper tails for counting objects in randomly induced subhypergraphs and rooted random graphs. *Ark. Mat.* **49** (2011), 79–96.

[18] S. Janson and L. Warnke. The lower tail: Poisson approximation revisited. *Random Struct. Alg.* **48** (2016), 219–246.

[19] T. Łuczak and P. Prałat. Chasing robbers on random graphs: zigzag theorem. *Random Struct. Alg.* **37** (2010), 516–524.

[20] T. Łuczak and J. Spencer. When does the zero-one law hold? *J. Amer. Math. Soc.* **4** (1991), 451–468.

[21] T. Makai. The reverse H-free process for strictly 2-balanced graphs. *J. Graph Theory* **79** (2015), 125–144.

[22] A. Noever. Online Ramsey games for more than two colors. *Random Struct. Alg.* **50** (2017), 464–492.

[23] A. Noever and A. Steger. Local resilience for squares of almost spanning cycles in sparse random graphs. *Electron. J. Combin.* **24** (2017), Paper 4.8, 15 pp.

[24] A. Ruciński. Matching and covering the vertices of a random graph by copies of a given graph. *Discrete Math.* **105** (1992), 185–197.

[25] O. Riordan and L. Warnke. The Janson inequalities for general up-sets. *Random Struct. Alg.* **46** (2015), 391–395.

[26] M. Schacht and F. Schulenburg. Sharp thresholds for Ramsey properties of strictly balanced nearly bipartite graphs. *Random Struct. Alg.* **52** (2018), 3–40.

[27] S. Shelah and J. Spencer. Zero-one laws for sparse random graphs. *J. Amer. Math. Soc.* **1** (1988), 97–115.

[28] M. Šileikis. *Inequalities for Sums of Random Variables: a combinatorial perspective.* PhD thesis, AMU Poznań (2012). Available from `http://hdl.handle.net/10593/2870`

[29] M. Šileikis and L. Warnke. A counterexample to the DeMarco-Kahn Upper Tail Conjecture. *Random Struct. Alg.*, **55** (2019), 775–794.

[30] J. Spencer. Threshold functions for extension statements. *J. Combin. Theory Ser. A* **53** (1990), 286–305.

[31] J. Spencer. Counting extensions. *J. Combin. Theory Ser. A* **55** (1990), 247–255.

[32] J. Spencer. *The Strange Logic of Random Graphs.* Springer-Verlag, Berlin (2001).

[33] J. Spencer and G. Tóth. Crossing numbers of random graphs. *Random Struct. Alg.* **21** (2002), 347–358.

[34] R. Spöhel, A. Steger, and L. Warnke. General deletion lemmas via the Harris inequality. *J. Comb.* **4** (2013), 251–271.

[35] V. Vu. A large deviation result on the number of small subgraphs of a random graph. *Combin. Probab. Comput.* **10** (2001), 79–94.
[36] L. Warnke. On the method of typical bounded differences. *Combin. Probab. Comput.* **25** (2016), 269–299.
[37] L. Warnke. Upper tails for arithmetic progressions in random subsets. *Israel J. Math.* **221** (2017), 317–365.
[38] L. Warnke. On the missing log in upper tail estimates. *J. Combin. Theory Ser. B* **140** (2020), 98–146.
[39] B. Ycart and J. Ratsaby. The VC dimension of k-uniform random hypergraphs. *Random Struct. Alg.* **30** (2007), 564–572.

A Appendix: Proof of Theorem 7

Our proof of Theorem 7 from Section 2 hinges on the following fairly routine claim (based on central moment estimates), where we write $X_{\mathbf{x}} = X_{G,H}(\mathbf{x})$, as usual. Recall that $\mu = \mathbb{E}X_{\mathbf{x}}$ and $\sigma^2 = \operatorname{Var} X_{\mathbf{x}}$ do not depend on the particular choice of $\mathbf{x}$, and that $\Phi = \Phi_{G,H}$ is defined in (6).

Claim 17. *For any rooted graph (G, H), the following holds for all $p = p(n) \in [0, 1]$ and $\mathbf{x} \in [n]_{v_G}$:*

(i) *If $\Phi = \Omega(1)$, then $\mathbb{E}(X_{\mathbf{x}} - \mu)^m = O\big((\mu^2/\Phi)^{m/2}\big)$ for any fixed integer $m \geqslant 2$, where the implicit constant may only depend on m, G and H.*

(ii) *If $\Phi(1-p) \to \infty$, then $\mathbb{P}(|X_{\mathbf{x}} - \mu| < \delta\sigma) \to \mathbb{P}(|\eta| < \delta)$ for any fixed $\delta \in (0, \infty)$, where η is a standard normal random variable.*

Proof of Claim 17. Recalling the variance estimate (10) and the definition (6) of Φ, a straightforward extension of the textbook proof of [15, Theorem 6.5 and Remark 6.6] for (unrooted) subgraph counts yields

$$\mathbb{E}(X_{\mathbf{x}} - \mu)^m = \mathbb{1}_{\{m \text{ even}\}}(m-1)!!\sigma^m(1 + o(1)) + O\Big(\sum_{1 \leqslant \ell < m/2} \sigma^m \big(\Phi(1-p)\big)^{\ell - m/2}\Big), \tag{80}$$

where $(m-1)!! = (m-1) \cdot (m-3) \cdots 1$ when m is even. As $\ell - m/2 < 0$ the sum in (80) is $o(\sigma^m)$ when $\Phi(1-p) \to \infty$. Hence $\mathbb{E}(X_{\mathbf{x}} - \mu)^m/\sigma^m \to \mathbb{1}_{\{m \text{ even}\}}(m-1)!!$, so the method of moments (see, e.g., [15, Corollary 6.3]) implies that $(X_{\mathbf{x}} - \mu)/\sigma$ converges to η in distribution, which implies (ii).

Turning to (i), from (10) and $\Phi = \Omega(1)$ we infer that in (80) we have $\sigma^m = O\big(\mu^m/\Phi^{m/2}\big)$ and

$$\sigma^m \big(\Phi(1-p)\big)^{\ell - m/2} \asymp \mu^m/\Phi^{m/2} \cdot \Phi^{\ell - m/2}(1-p)^\ell = O\big(\mu^m/\Phi^{m/2}\big),$$

which shows that (80) implies (i). $\square$

Proof of Theorem 7. For (i) we may assume that in our lower bound on t we have $(t/\mu)^2\Phi \geqslant n^c$ for sufficiently large n, where $c > 0$ is a constant. Fix an arbitrary constant $\tau > 0$, and set $m := \lceil (v_G + \tau + 1)/c \rceil$. Using first a union bound, next Markov's inequality, and finally Claim 17 (i), it then readily follows that

$$\mathbb{P}\left(\max_{\mathbf{x} \in [n]_{v_G}} |X_{\mathbf{x}} - \mu| \geqslant t\right) \leqslant \sum_{\mathbf{x} \in [n]_{v_G}} \frac{\mathbb{E}(X_{\mathbf{x}} - \mu)^{2m}}{t^{2m}} \leqslant O(n^{v_G}) \cdot \left(\frac{\mu^2}{t^2\Phi}\right)^m = o(n^{-\tau}).$$

Turning to (ii) we fix $\mathbf{x} = (1, \ldots, v_G)$, say, and claim that $|X_{\mathbf{x}} - \mu| \geqslant \varepsilon\mu$ whp. In case (a) we fix $\delta > 0$. Combining the variance estimate (10) with our assumption $\varepsilon^2\Phi/(1-p) \to 0$, we infer for n large enough that

$$\varepsilon\mu/\sigma = \sqrt{\varepsilon^2\mu^2/\sigma^2} \asymp \sqrt{\varepsilon^2\Phi/(1-p)} \ll \delta.$$

Together with Claim 17 (ii), it follows that

$$\limsup_{n \to \infty} \mathbb{P}(|X_{\mathbf{x}} - \mu| < \varepsilon\mu) \leqslant \mathbb{P}(|\eta| < \delta).$$

Since $\delta > 0$ was arbitrary, now the basic fact $\lim_{\delta \to 0} \mathbb{P}(|\eta| \leqslant \delta) \to 0$ completes the proof in case (a). In the remaining case (b), after recalling $X_{\mathbf{x}} = X_{G,H}(\mathbf{x})$, then Markov's inequality readily implies

$$\mathbb{P}(X_{\mathbf{x}} \geqslant 1) \leqslant \min_{G \subsetneq J \subseteq H : e_J > e_G} \mathbb{P}(X_{G,J}(\mathbf{x}) \geqslant 1) \leqslant \Phi_{G,H} \to 0,$$

so that whp $|X_{\mathbf{x}} - \mu| = \mu \geqslant \varepsilon\mu$, completing the proof. $\square$

B Appendix: Proof of Lemma 16

Recalling the notation and definitions for unrooted graphs introduced at the beginning of Section 4.3, Lemma 16 is implied by claim (v) of the following more general auxiliary result, whose technical statement is optimized for ease of the proofs (which are partially inspired by [30, Lemma 4 and 7]).

Lemma 18. *Let K be a balanced graph with $e_K \geqslant 1$. There are constants $\beta, B, C > 0$ such that, for all $p = p(n) \in [0,1]$ with $n^{-1/d_K} \ll p = O(n^{\beta-1/d_K})$, the following holds whp in $\mathbb{G}_{n,p}$, writing $\lambda := np^{d_K}$:*

- (i) *If $G \subseteq K$ is primal for K, then any two copies of G are either vertex-disjoint or their intersection is isomorphic to a primal subgraph of K.*
- (ii) *If $G_0 \subsetneq G_1$ are both primal for K and there is no third primal F such that $G_0 \subsetneq F \subsetneq G_1$, then, for every copy G_0' of G_0, all copies of G_1 that contain G_0' are vertex-disjoint outside of $V(G_0')$.*
- (iii) *If G_0, G_1 are as in (ii), then every copy of G_0 is contained in at most $B\lambda^{v_{G_1}-v_{G_0}}$ copies of G_1.*
- (iv) *If $v \in V(K)$ and $G^{(v)} \subseteq K$ is a minimal primal subgraph of K containing v, then for each vertex $x \in [n]$ there is at most one $(v, G^{(v)})$-extension of x.*
- (v) *If $G_{\min} \subseteq K$ is primal for K with the smallest number of vertices, then every vertex $x \in [n]$ is contained in at most $C\lambda^{v_K - v_{G_{\min}}}$ copies of K.*

Proof. (i): Fix a graph $U := G_1 \cup G_2$ that is formed by the union of some two distinct overlapping copies G_1, G_2 of G. Since there are only finitely many such graphs, it is enough to show that $\mathbb{G}_{n,p}$ whp does not contain a copy of U when the intersection $I := G_1 \cap G_2$ is not isomorphic to a primal subgraph of K. Noting $e_U = 2e_G - e_I$ and $v_U = 2v_G - v_I$, using that I is not primal, i.e., $d_I < d_K = d_G$, it follows that

$$d_U = \frac{e_U}{v_U} = \frac{2e_G - e_I}{2v_G - v_I} = \frac{2v_G d_G - v_I d_I}{2v_G - v_I} > d_G.$$

Since $p = O(n^{\beta-1/d_K}) \ll n^{-1/d_U}$ for $\beta > 0$ small enough, using $\mu_U \asymp n^{v_U}p^{e_U} = (np^{d_U})^{v_U} \ll 1$ and Markov's inequality it readily follows that $\mathbb{G}_{n,p}$ whp does not contain a copy of U.

(ii): By (i) any two distinct copies of G_1 that contain the same copy of G_0 must intersect in a subgraph isomorphic to some primal J with $G_0 \subseteq J \subsetneq G_1$. The assumed properties of G_0, G_1 imply that J cannot contain G_0 properly. Hence $J = G_0$, which implies that all copies of G_1 are vertex-disjoint outside $V(G_0')$.

(iii): In K_n, for each a copy G_0' of G_0 the number of copies of G_1 containing G_0' is at most $\lfloor An^{v_{G_1}-v_{G_0}} \rfloor$ for some constant $A = A(K) \geqslant 1$. Defining $B := e^2 A$, let $\mathcal{E}_{G_0'}$ denote the event that there are at least $z := \lceil B\lambda^{v_{G_1}-v_{G_0}} \rceil$ copies of G_1 that (a) contain G_0' and (b) are vertex-disjoint outside of $V(G_0')$. Since the copies share no edges other than $E(G_0')$, a standard union bound argument yields

$$\mathbb{P}\left(\mathcal{E}_{G_0'}\right) \leqslant \binom{\lfloor An^{v_{G_1}-v_{G_0}} \rfloor}{z} p^{e_{G_0}+(e_{G_1}-e_{G_0})z}.$$

Since K is balanced and G_0, G_1 are its primal subgraphs, we see that $d_{G_0} = d_{G_1} = d_K$, from which it routinely follows that

$$d(G_0, G_1) = \frac{e_{G_1} - e_{G_0}}{v_{G_1} - v_{G_0}} = \frac{v_{G_1}d_K - v_{G_0}d_K}{v_{G_1} - v_{G_0}} = d_K.$$

Hence $n^{v_{G_1}-v_{G_0}}p^{e_{G_1}-e_{G_0}} = (np^{d(G_0,G_1)})^{v_{G_1}-v_{G_0}} = \lambda^{v_{G_1}-v_{G_0}}$. Together with $\binom{x}{z} \leqslant (ex/z)^z$ and the definition of z, it then follows that, say,

$$\mathbb{P}\left(\mathcal{E}_{G_0'}\right) \leqslant p^{e_{G_0}} \cdot \left(eA\lambda^{v_{G_1}-v_{G_0}}/z\right)^z \leqslant p^{e_{G_0}} \cdot \exp\bigl(-\lambda^{v_{G_1}-v_{G_0}}\bigr).$$

Using $d_{G_0} = d_K$, we also obtain $n^{v_{G_0}}p^{e_{G_0}} = (np^{d_{G_0}})^{v_{G_0}} = \lambda^{v_{G_0}}$. Noting that $\lambda \to \infty$ as $n \to \infty$, by summing over all possible copies G_0' (there are at most $n^{v_{G_0}}$ of them) it then follows that

$$\sum_{G_0'} \mathbb{P}\left(\mathcal{E}_{G_0'}\right) \leqslant n^{v_{G_0}}p^{e_{G_0}} \cdot \exp\bigl(-\lambda^{v_{G_1}-v_{G_0}}\bigr) \leqslant \lambda^{v_{G_0}} \cdot \exp\bigl(-\lambda\bigr) \to 0.$$

Now (iii) follows readily by a union bound argument and (ii).

(iv): Given $x \in [n]$, let G', G'' be two $(v, G^{(v)})$-extensions of x. Since G', G'' are both copies of a primal subgraph $G^{(v)}$, by (i) we have that $G' \cap G''$ is a copy of a primal subgraph $J \subseteq G^{(v)}$. Since each of the two isomorphisms mapping $G^{(v)}$ to G' and G'' maps v to x, we infer that $v \in V(J)$. But since $G^{(v)}$ is minimal among primals containing v, it follows that $J = G^{(v)}$ and thus $G' = G''$.

(v): Given $v \in V(K)$, set $G_0 := G^{(v)}$ and choose a maximal chain

$$G^{(v)} = G_0 \subsetneq G_1 \subsetneq \cdots \subsetneq G_\ell = K$$

of primal subgraphs of K, with the property that for any i there is no primal F for K satisfying $G_i \subsetneq F \subsetneq G_{i+1}$ (to clarify: since K is balanced and thus primal, such maximal chains always exist). For each (v, K)-extension of x we can select a unique sequence of copies

$$x \in G'_0 \subsetneq G'_1 \subsetneq \cdots \subsetneq G'_\ell$$

such that G'_0 is an (v, G_0)-extension of x, and that, for each $i \in [\ell]$, G'_i is a copy of G_i. Hence it is enough to bound the number of such sequences, assuming (iii) and (iv). By (iv) there is at most one choice for G'_0, and, given G'_{i-1} with $i \in [\ell]$, by (iii) there are at most $B\lambda^{v_{G_i} - v_{G_{i-1}}}$ choices for suitable G'_i. Multiplying these bounds, we obtain

$$\max_{x \in [n]} X_{v,K}(x) \leqslant \prod_{i \in [\ell]} B\lambda^{v_{G_i} - v_{G_{i-1}}} \asymp \lambda^{\sum_{i \in [\ell]} [v_{G_i} - v_{G_{i-1}}]} = \lambda^{v_{G_\ell} - v_{G_0}} = \lambda^{v_K - v_{G^{(v)}}}.$$

Summing over all $v \in V(K)$, using $v_K - v_{G^{(v)}} \leqslant v_K - v_{G_{\min}}$ and $\lambda \gg 1$ it follows that, for some $C = C(K) > 0$, each vertex is contained in at most $C\lambda^{v_K - v_{G_{\min}}}$ copies of K, completing the proof of Lemma 18. $\square$