

RESEARCH ARTICLE

How does the chromatic number of a random graph vary?

Annika Heckel¹  | Oliver Riordan²

¹Matematiska institutionen, Uppsala universitet, Uppsala, Sweden

²Mathematical Institute, University of Oxford, Radcliffe Observatory Quarter, Oxford, UK

Correspondence

Annika Heckel, Matematiska institutionen, Uppsala universitet, Box 480, 751 06 Uppsala, Sweden.
Email: annika.heckel@math.uu.se

Funding information

ERC, Grant/Award Numbers: 676632-RanDM, 772606-PTRCSP

Abstract

The chromatic number $\chi(G)$ of a graph G is a fundamental parameter, whose study was originally motivated by applications ($\chi(G)$ is the minimum number of internally compatible groups the vertices can be divided into, if the edges represent incompatibility). As with other graph parameters, it is also studied from a purely theoretical point of view, and here a key question is: what is its typical value? More precisely, how does $\chi(G_{n,1/2})$, the chromatic number of a graph chosen uniformly at random from all graphs on n vertices, behave? This quantity is a random variable, so one can ask (i) for upper and lower bounds on its typical values, and (ii) for bounds on how much it varies: what is the width (for example, standard deviation) of its distribution? On (i) there has been considerable progress over the last 45 years; on (ii), which is our focus here, remarkably little. One would like both upper and lower bounds on the width of the distribution, and ideally a description of the (appropriately scaled) limiting distribution. There is a well-known upper bound of Shamir and Spencer of order \sqrt{n} , improved slightly by Alon to $\sqrt{n}/\log n$, but no non-trivial lower bound was known until 2019, when the first author proved that the width is at least $n^{1/4-o(1)}$ for infinitely many n , answering a long-standing question of Bollobás. In this paper we have two main aims: first, we shall prove a much stronger lower bound

on the width. We shall show unconditionally that, for some values of n , the width is at least $n^{1/2-o(1)}$, matching the upper bounds up to the error term. Moreover, conditional on a recently announced sharper explicit estimate for the chromatic number, we improve the lower bound to order $\sqrt{n} \log \log n / \log^3 n$, within a logarithmic factor of the upper bound. Second, we will describe a number of conjectures as to what the true behaviour of the variation in $\chi(G_{n,1/2})$ is, and why. The first form of this conjecture arises from recent work of Bollobás, Heckel, Morris, Panagiotou, Riordan and Smith. We will also give much more detailed conjectures, suggesting that the true width, for the worst case n , matches our lower bound up to a constant factor. These conjectures also predict a Gaussian limiting distribution.

MSC 2020
05C15, 05C80 (primary)

1 | INTRODUCTION

Given a graph G , a *colouring* of G is an assignment of colours to the vertices of G so that no two adjacent vertices are coloured the same. The smallest number of colours for which this is possible is called the *chromatic number* of G , and is denoted by $\chi(G)$. This graph parameter plays a very important role in applications, in particular in assignment problems. Here, however, we focus on $\chi(G)$ from a theoretical point of view, simply as a natural and fundamental parameter of a graph.

As with any important graph parameter, an interesting question is: what is its typical value, if we choose G uniformly at random from all graphs on n (labelled) vertices? Also, how much does the chromatic number fluctuate around this critical value? Given $n \in \mathbb{N}$ and $p \in [0, 1]$, the *binomial random graph* $G_{n,p}$ is the graph on n labelled vertices where each possible edge is included independently with probability p , so a uniformly random graph on n vertices is simply $G_{n,1/2}$. The question just described was raised (in the sparse setting) by Erdős and Rényi [11], in one of their seminal papers which initiated the study of random graphs. Erdős later posed this question for the dense case, see Bollobás [7]. In this section we first outline the history of this problem, concentrating on the most relevant results. Then we state our new results. Finally, we present a number of conjectures as to the true behaviour of $\chi(G_{n,1/2})$, in various levels of detail. The basic conjecture is due to Bollobás, Morris, Panagiotou and Smith together with the present authors; the finer conjectures are new.

1.1 | Past results and questions

In 1975, Grimmett and McDiarmid [14] found the likely order of magnitude of $\chi(G_{n,p})$ for $0 < p < 1$ constant. In a landmark contribution in 1987, Bollobás [5] determined the asymptotic behaviour

of $\chi(G_{n,p})$ in this case. In stating this result we follow a standard convention, writing q for $1 - p$ and b for $1/q = 1/(1 - p)$ to make the formulae more compact.

Theorem 1 [5]. *Let $0 < p < 1$ be constant, and let $b = 1/(1 - p)$. With high probability,[†]*

$$\chi(G_{n,p}) \sim \frac{n}{2 \log_b n}.$$

Formally, this means that for any constant $\epsilon > 0$, with high probability $\chi(G_{n,p})$ is between $1 - \epsilon$ and $1 + \epsilon$ times the bound on the right-hand side.

Theorem 1 has been sharpened several times [12, 23, 24], most recently in [15].

Theorem 2 [15]. *Fix $p \leq 1 - 1/e^2$. Then, with high probability (whp),*

$$\chi(G_{n,p}) = \frac{n}{2 \log_b n - 2 \log_b \log_b n - 2 \log_b 2} + o\left(\frac{n}{\log^2 n}\right) \tag{1}$$

where $b = 1/(1 - p)$.

For constant $p > 1 - 1/e^2$, there is a slightly more complicated expression which also determines $\chi(G_{n,p})$ whp up to accuracy $o(\frac{n}{\log^2 n})$ [15].

Łuczak [19] extended Theorem 1 to the case $p \rightarrow 0$, giving a similar expression for $\chi(G_{n,p})$ whenever $p > C/n$ for some large enough constant C .

All the results we have mentioned so far examine the likely *value* of the chromatic number — they give increasingly sharp upper or lower bounds for $\chi(G_{n,p})$ which hold with high probability. A separate line of enquiry asks for the *concentration* of the chromatic number: even if we cannot pin down $\chi(G_{n,p})$ exactly, can we say something about how much it varies?

The starting point for these questions is the classic result of Shamir and Spencer from their 1987 paper [26], in which they pioneered the use of martingale concentration inequalities in probabilistic combinatorics, something which has now become a standard tool in the area. They proved that for any function $p = p(n)$, the chromatic number of $G_{n,p}$ takes one of at most about \sqrt{n} consecutive values whp.

Theorem 3 [26]. *Let $p = p(n) \in (0, 1)$ and $\omega(n) \rightarrow \infty$ be arbitrary functions. Then there is a sequence of intervals $[s_n, t_n]$ of length*

$$\ell_n := t_n - s_n \leq \sqrt{n\omega(n)}$$

such that, whp,

$$\chi(G_{n,p}) \in [s_n, t_n].$$

It is not hard to show that for certain extreme values of $p = p(n)$, Theorem 3 is tight: Alon and Krivelevich [2] note that $\chi(G_{n,p})$ is not concentrated on fewer than $\Theta(\sqrt{n})$ values for $p = 1 - 1/(10n)$.

[†] As usual, we say that a sequence $(E_n)_{n \in \mathbb{N}}$ of events holds whp if $\mathbb{P}(E_n) \rightarrow 1$ as $n \rightarrow \infty$.

For the dense case, where p is constant, Alon gave a slight improvement to intervals of length about $\frac{\sqrt{n}}{\log n}$ ([4], Section 7.9, Exercise 3; see also [25]). If p tends to 0 quickly enough, however, Theorem 3 can be improved considerably.

Shamir and Spencer [26] showed that if $p = n^{-c}$ for $c \in (0, \frac{1}{2})$, then $\chi(G_{n,p})$ is concentrated on at most about $n^{\frac{1}{2}-c} \log n$ values. For $c > \frac{1}{2}$, they proved concentration on *constantly* many values. Łuczak [20] showed that if $c > \frac{5}{6}$, then $\chi(G_{n,p})$ is maximally concentrated: whp it takes one of at most two consecutive values. Finally, Alon and Krivelevich [2] proved two-point concentration whenever $p < n^{-c}$ with $c > \frac{1}{2}$ constant.

It should be noted that none of these concentration results gives any information about the *location* of the concentration intervals. In a breakthrough contribution, Achlioptas and Naor [1] found two *explicit* values for $\chi(G_{n,p})$ with $p = d/n$ where d is constant. Later, Coja-Oghlan, Panagiotou and Steger [10] extended this result to $p < n^{-\frac{3}{4}-\varepsilon}$, giving three explicit values in this case.

In view of strong results asserting sharp concentration of the chromatic number, starting in the late 1980s Bollobás raised, and he and Erdős disseminated, the *opposite* question: can we find any examples where the chromatic number of $G_{n,p}$ is *not* very sharply concentrated? Of course there are cases where this is trivially true, such as when $p = 1 - 1/(10n)$ as mentioned above. But what about interesting examples, and what about the most natural special case, $p = 1/2$?

It took quite a while for this question to appear in print. In an open problems appendix to the first edition of *The Probabilistic Method* [3], Erdős asked: can we prove that $\chi(G_{n,1/2})$ is not concentrated on constantly many values? Bollobás reiterated this question in [7], asking for *any* non-trivial results asserting a lack of concentration. The problem is also discussed in [2, 6, 9, 13, 18].

The first result of this type was recently given by the first author in [16]: it turns out that, at least for some values of n , the chromatic number of $G_{n,1/2}$ is not concentrated on fewer than about $n^{\frac{1}{4}}$ values.

Theorem 4 [16]. *Let $c < \frac{1}{4}$ be a constant, and let $([s_n, t_n])_{n \geq 1}$ be a (deterministic) sequence of intervals such that $\mathbb{P}(\chi(G_{n,1/2}) \in [s_n, t_n]) \rightarrow 1$ as $n \rightarrow \infty$. Then there are infinitely many n such that $t_n - s_n > n^c$.*

In other words, slightly informally, for $c < 1/4$ there is no sequence of intervals of length n^c which contain $\chi(G_{n,1/2})$ with high probability.

1.2 | Main results

In this paper, we improve the lower bound in Theorem 4 to an almost optimal one, giving a lower bound on the concentration interval length which nearly matches the upper bound from Theorem 3.

Theorem 5. *Fix $p \in (0, 1)$ and $c < \frac{1}{2}$, and let $([s_n, t_n])_{n \geq 1}$ be a (deterministic) sequence of intervals such that $\mathbb{P}(\chi(G_{n,p}) \in [s_n, t_n]) \rightarrow 1$ as $n \rightarrow \infty$. Then there are infinitely many n such that $t_n - s_n > n^c$.*

It is clear from the form of the result that Theorem 5 also holds if we replace c with $\frac{1}{2} - o(1)$ for some function $o(1)$ which tends to 0 sufficiently slowly. Up to this vanishing term, the exponent

matches the classic upper bound of \sqrt{n} for the concentration interval length given by Shamir and Spencer, and Alon’s improved upper bound of $\frac{\sqrt{n}}{\log n}$.

Considering intervals centred on the expectation of $\chi(G_{n,p})$, Theorem 5 implies (but is not implied by) a corresponding bound on the variance of $\chi(G_{n,p})$. Concretely, for any $c < 1$, we do not have $\text{Var}(\chi(G_{n,p})) = O(n^c)$.

Note that neither Theorem 4 nor Theorem 5 tells us anything about the concentration of the chromatic number of $G_{n,p}$ for any particular n , let alone every n . They only state that whenever $[s_n, t_n]$ is a sequence of intervals which contain $\chi(G_{n,p})$ whp, there must be a subsequence of long intervals. Thus, these results do not rule out the unlikely scenario that the chromatic number of $G_{n,p}$ is spread out over about \sqrt{n} values on some sparse subsequence of the integers, and is one-point concentrated everywhere else.

We will prove a stronger result than Theorems 5 and 6. To state this we introduce some notation, and review some classic results, concerning the independence number of $G_{n,p}$.

A set of vertices is *independent* in a graph G if there are no edges of G between them; the *independence number* of G , denoted by $\alpha(G)$, is the maximum size of such a set in G . As before, let $q = 1 - p$ and $b = 1/q$. For p constant, $\alpha(G_{n,p})$ can be determined precisely as follows: let

$$\alpha_0 = \alpha_0(n) := 2 \log_b n - 2 \log_b \log_b n + 2 \log_b (e/2) + 1; \tag{2}$$

then Matula [21, 22] and independently Bollobás and Erdős [8] proved that

$$\alpha(G_{n,p}) = \lfloor \alpha_0 + o(1) \rfloor \text{ whp,}$$

pinning down $\alpha(G_{n,p})$ to at most two consecutive values. If we let

$$\alpha = \alpha(n) := \lfloor \alpha_0(n) \rfloor, \tag{3}$$

then in fact for most n , whp $\alpha(G_{n,p}) = \alpha$.

Given $t \geq 1$, we call an independent set of size t a *t-set*. Let X_t count the number of t -sets in $G_{n,p}$, and let

$$\mu_t = \mu_t(n) := \mathbb{E}[X_t] = \binom{n}{t} q^{\binom{t}{2}}. \tag{4}$$

If we interpret the formula above suitably for non-integer t , then $\alpha_0(n)$ is, to a good approximation, the value of t at which $\mu_t = 1$. In particular, unless α_0 is very close to an integer, we expect many α -sets and no $(\alpha + 1)$ -sets, so it is no surprise that $\alpha(G_{n,p}) = \alpha$ whp.

With this notation, we can now state our next, more precise, result.

Theorem 6. Fix $p \leq 1 - 1/e^2$ and $\varepsilon > 0$, and let $[s_n, t_n]$ be a sequence of intervals such that $\mathbb{P}(\chi(G_{n,p}) \in [s_n, t_n]) \geq 0.9$. Then, for each n such that $\mu_{\alpha(n)}(n) < n^{1-\varepsilon}$, there is an integer $n^* = (1 + o(1))n$ such that

$$t_{n^*} - s_{n^*} > C \frac{\sqrt{\mu_{\alpha(n^*)}(n^*)}}{\log n^*},$$

where

$$C = C(p, \varepsilon) = \frac{\varepsilon \log b}{9}$$

and, as usual, $b = 1/(1 - p)$.

Theorem 6 readily implies the case $p \leq 1 - 1/e^2$ of Theorem 5: we simply pick a sequence of n where μ_α is close to n , which is certainly possible; see, for example, [16, Lemma 4]. The case $p > 1 - 1/e^2$ of Theorem 5 will also follow easily from the proof of Theorem 6 (see the final part of Section 2.6). We have replaced the assumption that $\chi(G_{n,p})$ is in a certain interval whp with a weaker concrete assumption, since this is what the proof allows. The specific constant 0.9 is not optimised.

Theorem 6 still does not imply non-concentration for any particular n — this is a feature of the method — but for every n it will find some nearby n^* where the concentration interval is long. In many cases we believe that the bound above is tight up to the constant factor, including the dependence on ε ; see Section 1.3, and in particular Remark 17.

Even stronger bounds

Theorem 5 implies that there are *some* values of n such that $\chi(G_{n,p})$ is not concentrated on fewer than $n^{\frac{1}{2}-o(1)}$ values for some unspecified function $o(1)$. Can this be pushed any further towards Alon’s upper bound of $\sqrt{n}/\log n$? We focus on the case $p = \frac{1}{2}$.

The main bottleneck is the form of the error term in the estimate (1) in Theorem 2, which we make essential use of in the proof of Theorem 5. Specifically, we use that we have an explicit estimate for $\chi(G_{n,1/2})$, and that the derivative (w.r.t. n) of this estimate is sufficiently larger than $1/\alpha(n)$; see Remark 24 for how this affects the final bound.

Konstantinos Panagiotou and the first author [17] recently announced a sharper explicit estimate for $\chi(G_{n,1/2})$. To state this we need some definitions.

Definition 7. A vertex colouring of G is t -bounded if all colour classes have size at most t ; the t -bounded chromatic number of G , denoted $\chi_t(G)$, is the minimum number of colours in such a colouring. By an *unordered* (t -bounded) k -colouring of a graph G , we mean a partition of $V(G)$ into k non-empty independent sets (of size at most t). We may think of this as an equivalence class of k -colourings under permuting colours. Let $E_{n,k,t}$ denote the expected number of unordered t -bounded k -colourings of $G_{n,1/2}$. Then the t -bounded first moment threshold of $G_{n,1/2}$ is defined to be

$$k_t(n) := \min\{k : E_{n,k,t} \geq 1\}. \tag{5}$$

Note that $E_{n,n,t} = 1$, so this definition makes sense.

In [17] it is shown that if $a = a(n)$ is such that $n^{0.1} < \mu_a(n) < n^{1.9}$ (where $\mu_a(n)$ is defined in (4)), then whp

$$\chi_{a-1}(G_{n,1/2}) = k_{a-1}(n) + O(n^{0.99}). \tag{6}$$

Unsurprisingly, when μ_α is not too large, then $\chi_{\alpha-1}(G_{n,1/2})$ and $\chi(G_{n,1/2})$ are close, and then (6) (applied with $a = \alpha$) provides a good bound on the latter. For example, we trivially have that the expectation of the difference is at most μ_α , though we will need a much tighter bound (see Lemma 28). Assuming a much weaker form of a special case of (6), we can prove a stronger lower bound on the non-concentration interval.

Theorem 8. *Suppose that, for any integers n and $a = a(n)$ such that $\mu_a(n) = \Theta(n/\log^2 n)$, we have*

$$\chi_{a-1}(G_{n,1/2}) = k_{a-1}(n) + o(n \log \log n / \log^4 n) \text{ whp,} \tag{7}$$

where $k_t(n)$ is defined in (5). Then there is a constant $c > 0$ so that for any sequence of intervals $[s_n, t_n]$ such that $\mathbb{P}(\chi(G_{n,1/2}) \in [s_n, t_n]) \geq 0.9$, there is a sequence of integers n^* such that

$$t_{n^*} - s_{n^*} \geq c \frac{\sqrt{n^*} \log \log n^*}{\log^3 n^*}.$$

Remark 9. Theorem 8 immediately implies (assuming (7)) a corresponding lower bound on the variance of $Y_n = \chi(G_{n,1/2})$: writing w_n for $n^{1/2} \log \log n / \log^3 n$, if we take intervals I_n of length $cw_n/2$ centred on the mean of Y_n , then there are infinitely many n such that $\mathbb{P}(Y_n \notin I_n) > 0.1$, which implies $\text{Var}(Y_n) > 0.1(cw_n/4)^2 = \Omega(w_n^2)$, so $\limsup \text{Var}(Y_n)/w_n^2 > 0$.

As we shall describe in the next section, we believe that the bound given by Theorem 8 is optimal up to the constant factor.

1.3 | Conjectured behaviour

The behaviour of the chromatic number of $G_{n,p}$ is closely linked to that of the number of large independent sets, and specifically to X_α and $X_{\alpha-1}$ (where X_t is the number of independent t -sets), so we take a closer look at the distributions of these random variables.

First consider X_α . Let $\theta = \theta(n) = \log \mu_\alpha / \log n$, so that

$$\mu_\alpha = n^\theta. \tag{8}$$

Standard calculations (see [23, Section 3.c]) give

$$\theta = \alpha_0 - \alpha + o(1) \in [-o(1), 1 + o(1)]. \tag{9}$$

Thus θ behaves as shown in Figure 1: when α_0 is close to an integer, θ is close to 0. As we increase n , θ increases to near 1 (roughly linearly in $\log n$), until α_0 gets close to the next integer. At this point $\alpha(n)$ increases by 1 and θ drops back to near 0.

As for $X_{\alpha-1}$, note that

$$\mu_{\alpha-1} = \Theta\left(\frac{n}{\log n} \mu_\alpha\right) = n^{1+\theta+o(1)}. \tag{10}$$




FIGURE 1 The exponent $\theta = \theta(n)$ so that $\mu_\alpha = n^\theta$. When $\alpha_0(n)$ is close to an integer, $\alpha(n) = \lfloor \alpha_0 \rfloor$ increases by 1 and θ drops from close to 1 to close to 0. Note that one can think of each line segment as graphing the expected number of t -sets of some particular size t (or rather, the log of this divided by $\log n$). These lines extend above and below the strip shown in the figure, but when we are considering the largest independent set, we jump from one size to the next as n increases.

It turns out that both X_α and $X_{\alpha-1}$ are approximately Poisson for almost all n (see [6, Theorem 11.9]). In particular, X_α and $X_{\alpha-1}$ are not whp contained in any sequences of intervals shorter than $\sqrt{\mu_\alpha} = n^{\theta/2}$ and $\sqrt{\mu_{\alpha-1}} = n^{(1+\theta)/2+o(1)}$, respectively.

1.3.1 | The Zigzag Conjecture

We are now ready to state a conjecture on the correct length of the concentration interval made by Bollobás, Heckel, Panagiotou, Morris, Riordan and Smith. The conjecture states that the concentration interval length for $\chi(G_{n,1/2})$ is essentially the maximum of two proposed lower bounds, one which comes from fluctuations in X_α and one which comes from fluctuations in $X_{\alpha-1}$, which we will describe below.

We shall consider only the case $p = \frac{1}{2}$, for a number of reasons. Firstly, this is the original question; secondly, this simplifies the formulae somewhat; finally, and most importantly, for some constant p — in particular when $p > 1 - 1/e^2$ — the chromatic number of $G_{n,p}$ behaves differently to the case $p = \frac{1}{2}$ (see [15]), so its concentration may well behave differently too.

The chromatic number of $G_{n,1/2}$ is closely linked to its independence number. Every colour class in a colouring is an independent set, and so for any graph G on n vertices, $\chi(G) \geq n/\alpha(G)$. In $G_{n,1/2}$, this simple bound for the value of the chromatic number is asymptotically correct: Bollobás’ classic result implies that whp $\chi(G_{n,1/2}) \sim n/\alpha(G_{n,1/2})$, and Theorem 2 states that whp,

$$\chi(G_{n,1/2}) = \frac{n}{\alpha_0 - 1 - \frac{2}{\log 2} + o(1)} \approx \frac{n}{\alpha_0 - 3.89}.$$

It is plausible that an optimal colouring of $G_{n,1/2}$ contains all or almost all α -sets as colour classes. To see this heuristically, fix a number $k \approx \frac{n}{2 \log_2 n}$ of colours. Each essentially different colouring of the vertex set of $G_{n,1/2}$ with k colours corresponds to a *profile*, that is, a sequence of sizes for the colour classes. Among all profiles, it turns out that the expected number of




FIGURE 2 Exponent of the concentration interval length (in n). The dashed line is the conjectured lower bound $\theta/2$. The dotted line is the conjectured lower bound $(1 - \theta)/2$. The thicker ‘zigzag’ line is the maximum of these two lower bounds. The Zigzag Conjecture proposes that the concentration interval length of $\chi(G_{n,1/2})$ fluctuates between $n^{1/4+o(1)}$ and $n^{1/2+o(1)}$ along this line.

colourings[†] is maximised if all or almost all α -sets are included as colour classes. More precisely, the expectation is maximised by *unrealisable* profiles containing even more α -sets (order $n / \log n$). Although the expected number of colourings with such a profile is large, whp no such colouring exists, as there are not enough α -sets.

We saw above that X_α is approximately Poisson with mean n^θ . In particular, X_α varies by about $\sqrt{\mu_\alpha} \approx n^{\theta/2}$. If the number of available α -sets for our colouring varies by $\sqrt{\mu_\alpha}$, intuitively the total number of colours we need should vary by at least about

$$\frac{\sqrt{\mu_\alpha}}{\log n} = \frac{n^{\theta/2}}{\log n}. \tag{11}$$

Perhaps it is not immediately clear where the factor $\log n$ comes from. One heuristic way to see this is the following: if there are $n^{\theta/2}$ fewer α -sets, we can cover $n^{\theta/2}\alpha$ fewer vertices with α -sets and need to colour them in colour classes of size $\alpha - 1$ or less. On average we colour with classes of size $\approx \alpha_0 - 3.89$. So each α -set that we use covers $\Theta(1)$ extra vertices compared to a typical colour class, and hence saves $\Theta(1/\alpha_0) = \Theta(1/\log n)$ colours. This argument is an oversimplification; see Section 1.3.2 for a detailed discussion.

The first part of the Zigzag Conjecture states that (11) is indeed a lower bound for the concentration interval length of $\chi(G_{n,1/2})$ (see Figure 2).

The second part is another conjectured lower bound which comes from the variations of $X_{\alpha-1}$, and is slightly trickier to understand. Again fix a number $k \sim \frac{n}{2 \log_2 n}$ of colours, and consider the optimal colouring profile, that is, choose the number of colour classes of each possible size so that the expected number of such colourings is maximised. A reasonable guess is that $\chi(G_{n,1/2})$ is close to the smallest k such that, for the optimal colouring profile with k colours, the expected number of colourings is at least 1; it can be shown that the expected total number of (equivalence classes under permuting colours of) k -colourings is then not much more than 1.

It turns out that the optimal profile contains $l = \Theta(n / \log n)$ colour classes of size $\alpha - 1$, that is, some constant proportion of colour classes have this size. We now make some extremely rough

[†] As before, we actually count partitions into independent sets (with a given profile), rather than colourings.

estimates on how much the expected number of k -colourings with this profile changes as $X_{\alpha-1}$ varies, at least in the highest order terms.

Since we pick l colour classes from the $X_{\alpha-1}$ available $(\alpha - 1)$ -sets, the expected number of k -colourings with the optimal profile should be roughly proportional to $\binom{X_{\alpha-1}}{l} \approx X_{\alpha-1}^l$. Of course, in reality, not every choice of l colour classes is possible because not all $(\alpha - 1)$ -sets are disjoint, but the highest order term should match, or rather, should change in the same way as $X_{\alpha-1}$ varies.

Consider $G_{n,1/2}$ conditioned on some typical values for $X_{\alpha-1}$ which are $r \approx \sqrt{\mu_{\alpha-1}}$ apart, first on $X_{\alpha-1} = m \approx \mu_{\alpha-1}$ and then on $X_{\alpha-1} = m - r$. In the second case, where we have r fewer $(\alpha - 1)$ -sets, the expected number of k -colourings with optimum profile decreases by a factor of roughly

$$(m - r)^l / m^l \approx \exp(-rl/m) = \exp\left(-\Theta\left(\frac{n}{\sqrt{\mu_{\alpha-1}} \log n}\right)\right). \tag{12}$$

So how much does the chromatic number increase when $X_{\alpha-1} = m - r$ compared to the case $X_{\alpha-1} = m$? It can be shown (see Corollary 39) that adding one colour increases the expected number of colourings by a factor of size $\exp(\Theta(\log^2 n))$. So in order to make up for the decrease in the expectation in (12), we need to introduce order

$$\frac{n}{\sqrt{\mu_{\alpha-1}} \log^3 n}$$

additional colours. By (10), note that

$$\frac{n}{\sqrt{\mu_{\alpha-1}} \log^3 n} = \Theta\left(\frac{\sqrt{n}}{\sqrt{\mu_{\alpha}} \log^{5/2} n}\right) = \Theta\left(\frac{n^{(1-\theta)/2}}{\log^{5/2} n}\right). \tag{13}$$

The second part of the Zigzag Conjecture states that (13) is another lower bound for the concentration interval length of $\chi(G_{n,1/2})$ (see Figure 2).

Are counts of α -sets and $(\alpha - 1)$ -sets the only significant sources of non-concentration of the chromatic number? A recently announced result by the first author and Konstantinos Panagiotou [17] strongly suggests this (at least for $p = \frac{1}{2}$). Recall that the t -bounded chromatic number $\chi_t(G)$ is defined like the normal chromatic number except that we only allow colourings in which all colour classes have size at most t . The announced result is that the $(\alpha - 2)$ -bounded chromatic number of $G_{n,m}$ with $m = \frac{1}{2} \binom{n}{2}$ is two-point concentrated. In other words, once α -sets and $(\alpha - 1)$ -sets are banned as colour classes, and the number of edges is fixed, the required number of colours is extremely narrowly concentrated. It is easy to see that, in $G_{n,1/2}$, the variation in the number of edges only has a very small effect on the chromatic number, accounting for fluctuations of order at most $\log n$; see [16, Section 3] for a simple coupling argument showing this.

The full conjecture, therefore, states that the maximum of the lower bounds (11) and (13) is indeed the correct concentration interval length for $\chi(G_{n,1/2})$ — at least whenever $\theta(n)$ is bounded away from 0 and 1.

Ignoring terms of size $n^{o(1)}$, a simplified statement is the following.

Conjecture 10 (Zigzag Conjecture; Bollobás, Heckel, Morris, Panagiotou, Riordan and Smith). *Set $p = \frac{1}{2}$ and define $\theta = \theta(n)$ as in (8). Let*

$$\lambda = \lambda(n) := \max\left(\frac{\theta}{2}, \frac{1-\theta}{2}\right). \tag{14}$$

Then there is a sequence of intervals of length $n^{\lambda+o(1)}$ which contains $\chi(G_{n, \frac{1}{2}})$ whp. However, for any fixed $\varepsilon > 0$ and any sequence $(I_n)_{n \in \mathbb{N}}$ of intervals of length $n^{\lambda-\varepsilon}$, we have

$$\mathbb{P}(\chi(G_{n,1/2}) \in I_n) = o(1).$$

An analogous statement presumably holds for any constant $p \in (0, 1 - 1/e^2]$, or perhaps $p \in (0, 1 - 1/e^2)$.

Conjecture 10 would imply that the concentration interval length of $\chi(G_{n,1/2})$ fluctuates between $n^{\frac{1}{4}+o(1)}$ and $n^{\frac{1}{2}+o(1)}$ as shown in Figure 2.

Theorem 6 almost proves the first lower bound (11) coming from fluctuations in X_α : we show that, for any integer n with $\theta(n)$ bounded away from 1, there is another integer n^* nearby such that (11) holds. It is of course extremely unlikely that the width of the distribution of $\chi(G_{n,1/2})$ is significantly different between n and n^* , so our result presumably holds for all n , but we cannot prove this.

1.3.2 | Further conjectures

In this section we state a number of further conjectures refining Conjecture 10. We will explain the intuition behind these conjectures in the Appendix. So far we have focussed on the width of the distribution as measured by concentration in an interval; here it will often be more convenient to work with the variance. Of course we expect these to be equivalent: if $Y_n := \chi(G_{n,1/2})$ then we expect that Y_n is concentrated on some sequence of intervals of length ℓ_n if and only if $\ell_n/\sigma_n \rightarrow \infty$, where $\sigma_n^2 = \text{Var}(Y_n)$. However, we do not know this, only the one-way implication that small variance implies tight concentration.

We start with a conjecture on the worst case concentration width: we believe that, up to a constant factor, the lower bound given in Theorem 8 is optimal.

Conjecture 11. Let $p \in (0, 1 - 1/e^2]$ be constant, let $Y_n = \chi(G_{n,p})$, let $\sigma_n^2 = \text{Var}(Y_n)$, and set

$$w_n := \frac{\sqrt{n} \log \log n}{\log^3 n}. \tag{15}$$

Then

$$0 < \limsup \frac{\sigma_n}{w_n} < \infty. \tag{16}$$

Moreover, for any constant $c > 0$ there is a constant $d > 0$ such that along any sequence of integers n with $\mu_{\alpha(n)}(n) \sim cn / \log^2 n$ we have $\sigma_n \sim dw_n$.

As noted in Remark 9, Theorem 8 implies the first inequality in (16), subject to (7).

We have a corresponding conjecture for the best case, although we are less confident of this, so we state only the basic form.

Conjecture 12. Let $p \in (0, 1 - 1/e^2]$ be constant, let $Y_n = \chi(G_{n,p})$, let $\sigma_n^2 = \text{Var}(Y_n)$, and set

$$\tilde{w}_n := \frac{n^{1/4}}{\log^{7/4} n}. \tag{17}$$

Then

$$0 < \liminf \frac{\sigma_n}{\tilde{w}_n} < \infty.$$

In fact, we believe that for most (probably all) n , the chromatic number is asymptotically normally distributed, with a variance that follows (a refined version of) the graph suggested by the Zigzag Conjecture. We are least confident about points close to the minima in this graph, which we call ‘bad’.

Fix a constant $\delta > 0$, and call n ‘bad’ if $n^{1/2-\delta} \leq \mu_{\alpha(n)}(n) \leq n^{1/2+\delta}$, and ‘good’ otherwise.

Conjecture 13. Let $p \in (0, 1 - 1/e^2]$ be constant, and let $Y_n = \chi(G_{n,p})$. There are functions $f(n)$ and $g(n)$ such that, at least for ‘good’ n ,

$$\frac{Y_n - f(n)}{g(n)} \xrightarrow{d} N(0, 1),$$

where $N(0, 1)$ is a standard Gaussian distribution. Moreover, $g(n) = n^{\lambda(n)+o(1)}$, where $\lambda(n)$ is defined in (14).

For good n , the dominant source of the variation should be (as described earlier) the variation in the number of independent sets of a certain size $a = \alpha - 1$ or $a = \alpha$, depending on the parameters. Specifically, let

$$a(n) := \lfloor \alpha_0(n) - 1/2 \rfloor, \tag{18}$$

so, for good n , we have

$$n^{1/2+\delta+o(1)} \leq \mu_{a(n)}(n) \leq n^{3/2-\delta+o(1)}.$$

One can alternatively take this to define $a(n)$. For bad n , at least at a certain transition point that we do not identify precisely, it is not clear how to define $a(n)$. Indeed, in a certain range two sizes should contribute. However, the distribution should still be asymptotically normal, since a linear combination of two Gaussians is Gaussian.

Conjecture 14. Let $p \in (0, 1 - 1/e^2]$ be constant, let $Y_n = \chi(G_{n,p})$, and let Z_n be the number of independent sets of size $a(n)$ in $G_{n,p}$, where $a(n)$ is defined in (18). Then there are functions $f(n)$, $g(n)$ such that, for good n ,

$$\left(\frac{Y_n - f(n)}{g(n)}, \frac{Z_n - \mathbb{E}Z_n}{\sqrt{\text{Var}(Z_n)}} \right) \xrightarrow{d} (Z, Z),$$

where $Z \sim N(0, 1)$.

In other words, knowing $f(n)$ and $g(n)$ (which we do not), the value of Z_n is enough to predict $Y_n = \chi(G_{n,p})$ up to an error that is $o(g(n))$, that is, smaller order than the standard deviation. In fact, for good n , this $o(\cdot)$ term should be $n^{-\Omega(1)}$. We would expect the conclusion of Conjecture 14 to hold outside a much smaller ‘bad’ set, perhaps only having to exclude n such that $\mu_{\alpha(n)}(n)$ is $\Theta(h(n))$ for some $h(n)$ close to $n^{1/2}$.

Finally, we believe that, except for ‘bad’ n , we can describe the width of the distribution up to a constant factor, and in a significant fraction of cases, up to a $1 + o(1)$ factor. Defining $a = a(n)$ as above, define $x = x(n)$ by

$$\mu_a(n) = \frac{2xn}{a^2} = \Theta\left(\frac{xn}{\log^2 n}\right). \tag{19}$$

The precise normalisation here is not so important; the second formula is the key one.

Conjecture 15. Define $a(n)$ and $x(n)$ as in (18) and (19). For good n , the function $g(n)$ in Conjecture 13, or equivalently $\sigma_n = \sqrt{\text{Var}(Y_n)}$, satisfies the following bounds, with $c_0 = 2/\log 2$:

(i) if $x \rightarrow 0$, then

$$g(n) \sim \sqrt{\mu_a} \frac{\log \log n + \log(1/x)}{c_0 \log^2 n};$$

(ii) if $x = \Theta(1)$, then, defining w_n as in (15),

$$g(n) = \Theta\left(\sqrt{\mu_a} \frac{\log \log n}{\log^2 n}\right) = \Theta(w_n);$$

(iii) if $x \rightarrow \infty$ with $x = n^{o(1)}$, then

$$g(n) \sim \frac{\sqrt{\mu_a}}{x} \cdot \frac{\log \log n + \log x}{c_0 \log^2 n} \sim c_1 n^{1/2} \frac{\log \log n + \log x}{\sqrt{x} \log^3 n},$$

and (iv) if $x \geq (\log n)^C$ for some constant $C > 0$, then

$$g(n) = \Theta\left(\frac{\sqrt{\mu_a} \log x}{x \log^2 n}\right) = \Theta\left(\frac{\sqrt{n} \log x}{\sqrt{x} \log^3 n}\right).$$

Note that the four ranges above cover all good n , with some overlap between (iii) and (iv). (The formula for (iii) applies in case (iv) too, but simplifies to (iv) in that case.) We can give a single formula applicable in all cases, but it is not clear that this is informative – the transition from case (i) to cases (iii)/(iv) is rather arbitrary, since in case (ii) we do not even have a guess as to what the implicit constant should be (as a function of x). Still, defining

$$g_0(n) = \sqrt{\mu_a} \frac{\log \log n + |\log x|}{c_0(1+x) \log^2 n},$$

in all cases we conjecture that $g(n) = \Theta(g_0(n))$, with \sim in cases (i) and (iii).

Remark 16. If $\log \mu_\alpha / \log n$ is bounded away from $1/2$, 1 and $3/2$, then the formulae in (i) and (iv) match our earlier heuristics (11) and (13) up to constant factors. Thus, cases (i) and (iv) of Conjecture 15 refine the ‘zig’ and ‘zag’ parts of the Zigzag Conjecture. Case (ii), and also case (iii), interpolate between these parts, describing the conjectured shape of the top of the zigzag curve. For the bottom, we have not stated a very detailed conjecture, but extrapolating the formulae in (i) and (iv) suggests that when $\mu_{\alpha(n)}(n) = \Theta(\sqrt{n} / \log^{3/2} n)$ and $\mu_{\alpha(n)-1}(n) = \Theta(n^{3/2} / \log^{5/2} n)$ (which is within, and indeed in some sense the centre of, the ‘bad n ’ case), then the contributions from $(\alpha - 1)$ -sets and α -sets to $g(n)$ should both be of order $n^{1/4} / \log^{7/4} n$, and this is how Conjecture 12 arises.

Remark 17. The agreement between the lower bound in Theorem 6 and case (i) of Conjecture 15 is in some sense surprisingly strong. The formula for $t_{n^*} - s_{n^*}$ in the former matches $g(n^*)$ up to a constant factor, noting that $1/x$ is at least approximately n^ε . Since we may let ε tend to zero at some rate, and the dependence on ε matches, this shows that $\sqrt{\text{Var}(Y_n)} = \Omega(g_0(n))$ not for every n , but at least for some n^* near any good n with $\mu_\alpha(n) \leq n^{1-\gamma(n)}$, where $\gamma(n)$ is a function tending to zero at a rate that we have not determined. Similarly, it was quite a surprise to us (and not the case when we first formulated the conjectures) that we can prove a (conditional) lower bound (Theorem 8) that (for a subsequence) matches the upper bound in Conjecture 11.

A completely satisfactory understanding of the asymptotic distribution of $\chi(G_{n,1/2})$ would involve two further ingredients: we would like to know $g(n)$ (or σ_n) asymptotically, not just up to constant factors. It is quite possible that one could read out such a formula from our intuitive justification of the conjectures above (see the Appendix), though of course we are nowhere near a proof. The second is that we would of course like to know $f(n)$ up to an additive error of $o(g(n))$. This seems to be a much harder problem, for which we do not even have a conjecture. See the discussion in the Appendix.

The rest of the paper is organised as follows. First, in Section 2.1, we outline the general strategy of the proofs. In Section 2.2 we state and prove a concrete ‘framework lemma’ that formalises this strategy, essentially giving a conditional result subject to two ingredients. In Section 2.3 we provide the first ingredient, a simple coupling lemma. The details of the other ingredient vary from case to case; after some preliminaries in Section 2.4 we provide these, and so prove Theorems 6 and 5, in Sections 2.5 and 2.6, respectively. The (very much more involved) argument for Theorem 8 is given in Section 3, with the proof of the key lemmas in Sections 3.1 and 3.2. Finally, we discuss the intuition behind our more precise conjectures in the Appendix.

2 | PROOFS

2.1 | Proof outline

Before turning to the details, we outline the method, which is in principle simple but involves significant calculation. Throughout we fix $0 < p < 1$. There will be two key ingredients which are illustrated in Figure 3. First, we take as an input a suitable result establishing whp concentration of $\chi(G_{n,p})$ on some explicit interval $f(n) \pm \Delta(n)$. Here the interval length $2\Delta(n)$ will be much larger than the scale on which we are aiming to establish non-concentration. It will be essential that, interpolating $f(n)$ to non-integer values, over the range of n that we consider we have $\frac{df}{dn} > 1/\alpha$,




FIGURE 3 Illustration of the basic strategy. We know $\chi(G_{n,p})$ is concentrated in the (wide) grey band around a function $f(n)$ with slope more than $1/\alpha$. A coupling argument shows that for suitable r it is likely that $\chi(G_{n',p}) \leq \chi(G_{n,p}) + r$, where $n' = n + \alpha r$ (dotted lines with slope $1/\alpha$). If the concentration intervals $[s_n, t_n]$ are too short, a contradiction results.

where $\alpha = \alpha(n) = \lfloor \alpha_0(n) \rfloor$ as before, and we consider a range of values for n such that α does not change in this ‘window’. More specifically, we will suppose that

$$\frac{df}{dn} \geq \frac{1}{\alpha} + \delta$$

for some $\delta > 0$.

The second key ingredient is a simple coupling result, Lemma 19 below, and in particular its consequence, Corollary 21, which states, slightly informally, that for r not too large we may couple the random graphs $G_n = G_{n,p}$ and $G_{n'} = G_{n+\alpha r,p}$ so that with significant probability (say > 0.4 , though we could write > 0.99 by changing the constants) we have $\chi(G_{n'}) \leq \chi(G_n) + r$. Here we can take r up to roughly $\sqrt{\mu_\alpha(n)}$, the standard deviation of the number X_α of α -sets. The intuition behind this is, roughly speaking, that because the number of α -sets varies by at least r , planting r extra ones does not affect the distribution of our graph too much. Planting these sets in $G_{n'}$, $n' = n + \alpha r$, we can view the graph on the remaining vertices as G_n , giving the coupling.

Suppose for the moment that $\chi(G_n) = \chi(G_{n,p})$ were in fact deterministic, equal to some function $f_0(n)$. Then the coupling just described would show that $f_0(n') \leq f_0(n) + r$, that is, (essentially) that the function $f_0(n)$ has slope $f'_0(n)$ at most $1/\alpha$. This would lead to a

contradiction, considering a suitably large range of values of n . Indeed, by our first (concentration) assumption, $|f_0(n) - f(n)| \leq \Delta(n)$. But (ignoring the variation of $\Delta(n)$ over the relevant window), a line $f_0(n)$ with slope $1/\alpha$ cannot stay this close to a curve $f(n)$ with slope at least $1/\alpha + \delta$ for more than roughly $2\Delta/\delta$ consecutive values of n .

Of course, $\chi(G_n)$ is not deterministic. But in proving our result, we may assume that it is *almost* deterministic: for each n , we may assume that $\chi(G_n)$ is concentrated on some interval $[s_n, t_n]$ of length ℓ_n . With $n' = n + \alpha r$ for r not too large, as before, our coupling implies that $s_{n'} \leq t_n + r$, since it is reasonably likely that all inequalities in the chain $s_{n'} \leq \chi(G_{n'}) \leq \chi(G_n) + r \leq t_n + r$ hold. In turn, this gives

$$s_{n'} \leq s_n + \ell_n + r.$$

Can we still get a contradiction? Yes, if the numbers work out correctly. Defining $f_0(n) = s_n$, we see that between n and $n' = n + \alpha r$ this function has slope

$$\frac{s_{n'} - s_n}{\alpha r} \leq \frac{1}{\alpha} + \frac{\ell_n}{\alpha r},$$

so if $\ell_n/(\alpha r) \leq \delta/2$, say, we will get a contradiction much as before. Hence there must be some n such that $\ell_n > \alpha r \delta/2$. Note that to obtain a strong non-concentration result, we wish to take r as large as possible.

2.2 | The framework lemma

In this subsection we formalise the outline above in the following lemma. We have replaced various $o(1)$ bounds here by concrete bounds for definiteness, though in the application we mostly start with $o(1)$ bounds and take n large enough. In the application we will take $a = \alpha(n^-) = \alpha(n^+)$, where $\alpha(n)$ is defined in (3) and $[n^-, n^+]$ is the range of values of n we consider. Thus a will be the typical independence number of the relevant graphs $G_{n,p}$.

Lemma 18. *Let p, δ and Δ be positive real numbers with $p < 1$, and let $n^- < n^+$ and a be positive integers. Let $I = [n^-, n^+]$. Suppose that the following hold. Firstly, there is some function $f(n)$ such that for each (integer) $n \in I$ we have*

$$\mathbb{P}(\chi(G_{n,p}) \in [f(n) - \Delta, f(n) + \Delta]) \geq 0.99. \tag{20}$$

Secondly, for all (real) $n \in I$ we have

$$\frac{df}{dn} \geq \frac{1}{a} + \delta. \tag{21}$$

Thirdly, for each $n \in I$ we have

$$\mathbb{P}(\chi(G_{n,p}) \in [s_n, t_n]) \geq 0.9 \tag{22}$$

for some integers s_n, t_n . Fourthly, there is an increasing integer-valued function $r(n)$ such that for each $n \in I$ we have a coupling of $G_{n,p}$ and $G_{n+ar(n),p}$ such that

$$\mathbb{P}(\chi(G_{n+ar(n),p}) \leq \chi(G_{n,p}) + r(n)) \geq 0.4. \tag{23}$$

Finally, suppose also that

$$n^+ - n^- \geq 5\Delta/\delta, \quad \text{and} \quad n^+ - n^- \geq 5ar(n^+). \tag{24}$$

Then there is some integer $n \in I$ for which $t_n - s_n > \frac{a\delta r(n)}{2}$.

Proof. We follow the plan described in the previous section, with the minor complication that we allow r to vary with n ; this is not essential, but gives stronger results in some applications. Throughout we write G_n for $G_{n,p}$.

Firstly, for $n \in I$ define $\tilde{s}_n = \max\{s_n, f(n) - \Delta\}$ and $\tilde{t}_n = \min\{t_n, f(n) + \Delta\}$. Then by (20) and (22) we have

$$\mathbb{P}(\chi(G_n) \in [\tilde{s}_n, \tilde{t}_n]) \geq 0.89, \tag{25}$$

and in particular this interval is non-empty, which implies that

$$|\tilde{s}_n - f(n)| \leq \Delta. \tag{26}$$

Let us suppose for a contradiction that for every $n \in I$ we have

$$\ell_n := t_n - s_n \leq \frac{a\delta r(n)}{2},$$

and note for later that $\tilde{t}_n - \tilde{s}_n \leq \ell_n$. Now, for any $n \in I$ such that $n' = n + ar(n) \in I$, by (23) and (25) (applied twice), with probability at least $0.4 - 2 \times 0.11 > 0$ all three inequalities $\chi(G_{n'}) \geq \tilde{s}_{n'}$, $\chi(G_n) \leq \tilde{t}_n$ and $\chi(G_{n'}) \leq \chi(G_n) + r(n)$ hold. Hence, with positive probability

$$\tilde{s}_{n'} \leq \chi(G_{n'}) \leq \chi(G_n) + r(n) \leq \tilde{t}_n + r(n),$$

and in particular $\tilde{s}_{n'} \leq \tilde{t}_n + r(n)$. Since this is a deterministic statement, it always holds. Thus, recalling that $\tilde{t}_n - \tilde{s}_n \leq \ell_n$, we have

$$\tilde{s}_{n'} \leq \tilde{s}_n + \ell_n + r(n) \leq \tilde{s}_n + r(n) + \frac{a\delta r(n)}{2} = \tilde{s}_n + (n' - n) \left(\frac{1}{a} + \frac{\delta}{2} \right). \tag{27}$$

Finally, define a sequence $(n_i)_{0 \leq i \leq j}$ as follows: let $n_0 = n^-$ and, given n_i , let $n_{i+1} = n_i + ar(n_i)$ unless this value exceeds n^+ , in which case we set $j = i$ and stop. Note that by the stopping condition and the monotonicity of r ,

$$n_j > n^+ - ar(n_j) \geq n^+ - ar(n^+). \tag{28}$$

Applying (27) with $n = n_i$ (and so $n' = n_{i+1}$) for $0 \leq i < j$ and telescoping, we see that

$$\tilde{s}_{n_j} - \tilde{s}_{n_0} \leq (n_j - n_0) \left(\frac{1}{a} + \frac{\delta}{2} \right).$$

On the other hand, from (21), $f(n_j) - f(n_0) \geq (n_j - n_0) \left(\frac{1}{a} + \delta \right)$, so writing $h(n) = f(n) - \tilde{s}_n$ we have

$$h(n_j) - h(n_0) \geq (n_j - n_0) \delta / 2.$$

From (28) we have $n_j - n_0 > n^+ - n^- - ar(n^+)$. Hence, by (24), we have $n_j - n_0 > (4/5)(n^+ - n^-) \geq 4\Delta/\delta$. Thus $h(n_j) - h(n_0) > 2\Delta$, which contradicts (26). \square

2.3 | The coupling argument

In this section we present the coupling lemma we shall use. We state it somewhat more generally than needed here; in the application we will take $a = \alpha(n)$ (the typical independence number of $G_{n,p}$).

Lemma 19. *Let $p \in (0, 1)$ be constant, let $b = 1/(1 - p)$, and let $a = a(n)$ satisfy $1.01 \log_b n \leq a \leq 100 \log_b n$ and $\mu \leq n^{1.99}$, where $\mu = \mu_a(n) = \binom{n}{a}(1 - p)^{\binom{a}{2}}$. Then there is a coupling of the random graphs $G_n = G_{n,p}$ and $G_{n-a} = G_{n-a,p}$ with the property that*

$$\mathbb{P}(\chi(G_n) \leq \chi(G_{n-a}) + 1) \geq 1 - \frac{1 + o(1)}{2\sqrt{\mu}}.$$

Proof. Let U be a uniform random subset of $V = [n]$ of size a . Given U , let P_n be the random graph on V with no edges inside U , in which each of the other $\binom{n}{2} - \binom{a}{2}$ possible edges is present independently with probability p . Thus P_n is G_n with a random independent a -set ‘planted’. From the definition, we may realise G_{n-a} as $P_n[V \setminus U]$. Furthermore, since U is an independent set in P_n , we have

$$\chi(P_n) \leq \chi(G_{n-a}) + 1.$$

It remains only to show that we can couple the distributions of P_n and G_n to agree with sufficiently high probability.

The key observation is that P_n has the distribution of G_n ‘size-biased’ by the number X_a of independent a -sets. To see this, let H be any graph on $[n]$, let $q_H = \mathbb{P}(P_n = H)$, and let $X_a(H)$ be the number of independent a -sets in H . For $P_n = H$ to hold, our random set U must be independent in H , which has probability $X_a(H) / \binom{n}{a}$. Given such a choice of U , exactly the right edges outside U must be present. Hence

$$q_H = X_a(H) \binom{n}{a}^{-1} p^{e(H)} (1 - p)^{\binom{n}{2} - \binom{a}{2} - e(H)} = \frac{X_a(H)}{\mu} p^{e(H)} (1 - p)^{e(H^c)} = \frac{X_a(H)}{\mu} p_H, \tag{29}$$

where $p_H = \mathbb{P}(G_n = H)$.

Let τ be the total variation distance between the distributions of P_n and of G_n . Then

$$2\tau := \sum_H |q_H - p_H| = \sum_H \left| \frac{X_a}{\mu} - 1 \right| p_H = \mathbb{E} \left[\left| \frac{X_a - \mu}{\mu} \right| \right],$$

where the expectation refers to the random graph G_n . Thus by Jensen’s inequality (or by Cauchy–Schwarz),

$$4\tau^2 \leq \mu^{-2} \mathbb{E}[(X_a(G_n) - \mu)^2] = \mu^{-2} \text{Var}[X_a(G_n)].$$

Writing $\text{Var}[X_a(G_n)]$ as a sum (of covariances of indicator functions) over pairs U_1, U_2 of a -sets in V , the contribution from $U_1 = U_2$ is at most μ , while by a standard exercise the contribution from the remaining terms is $O(\mu^2 a^4/n^2 + \mu a n(1-p)^{a-1})$, with the two terms corresponding to U_1 and U_2 intersecting in 2 or $a - 1$ vertices, respectively. Under our assumptions $\text{Var}[X_a(G_n)] \sim \mu$, so $\tau \lesssim 1/(2\sqrt{\mu})$. Since G_n and P_n can be coupled to agree with probability $1 - \tau$, this completes the proof. \square

Remark 20. It is perhaps interesting that the proof of our coupling lemma relies on a variance bound, that is, an upper bound on how much $X_a(G_n)$ varies. In the end, we use the lemma to show, roughly speaking, that $\chi(G_n)$ varies *at least* a certain amount, because $X_a(G_n)$ does.

Corollary 21. *Let $p \in (0, 1)$ be constant, let $b = 1/(1 - p)$ and let $1 \leq a = a(n) \leq n$ satisfy $1.02 \log_b n \leq a \leq 99 \log_b n$ and $\mu \leq n^{1.98}$, where $\mu = \mu_a(n) = \binom{n}{a}(1 - p)^{\binom{a}{2}}$. Let $r \leq \sqrt{\mu}$ be an integer. Then if n is large enough, there is a coupling of the random graphs $G_n = G_{n,p}$ and $G_{n+ar} = G_{n+ar,p}$ with the property that*

$$\mathbb{P}(\chi(G_{n+ar}) \leq \chi(G_n) + r) > 0.4.$$

Proof. For $i = 0, \dots, r$, let $n_i = n + ai$, and let $\mu_i = \binom{n_i}{a}(1 - p)^{\binom{a}{2}}$. In this notation, $n_0 = n$, $n_r = n + ar$ and $\mu_0 = \mu$. Since $r \leq \sqrt{\mu} \leq n^{0.99}$ and $a \leq 99 \log_b n$, if n is large enough, for all $0 \leq i \leq r$,

$$n \leq n_i \leq n + n^{0.999}.$$

In particular, $\log n_i = \log n + O(n^{-0.001})$, so $1.01 \log_b n_i \leq a \leq 100 \log_b n_i$ if n is large enough. Furthermore, if n is large enough,

$$\mu_i = \mu_0 \left(1 + O\left(\frac{ra}{n}\right) \right)^a = \mu_0 \left(1 + O\left(\frac{ra^2}{n}\right) \right) \leq n^{1.99} \leq n_i^{1.99}.$$

So we may apply Lemma 19 to show that, for every $i \in \{1, \dots, r\}$, there is a coupling of the random graphs G_{n_i} and $G_{n_{i-1}}$ such that

$$\mathbb{P}(\chi(G_{n_i}) \leq \chi(G_{n_{i-1}}) + 1) \geq 1 - \frac{1 + o(1)}{2\sqrt{\mu_i}} \geq 1 - \frac{1 + o(1)}{2\sqrt{\mu}}. \tag{30}$$

The Gluing Lemma (which is trivial in this finite setting[†]) implies that there is a joint coupling of the random graphs G_{n_0}, \dots, G_{n_r} so that (30) holds for every $1 \leq i \leq r$. In this coupling, with probability at least

$$1 - (1 + o(1)) \frac{r}{2\sqrt{\mu}} \geq 1 - \frac{1 + o(1)}{2} > 0.4$$

we have $\chi(G_{n+ar}) \leq \chi(G_n) + r$. □

2.4 | Preliminaries for Theorem 6

Fix $p \leq 1 - 1/e^2$, and let

$$f(n) = f_p(n) := \frac{n}{2 \log_b n - 2 \log_b \log_b n - 2 \log_b 2} \tag{31}$$

be the estimate for $\chi(G_{n,p})$ given in Theorem 2. Note that

$$f(n) = \frac{n}{\alpha_0(n) - 1 - \frac{2}{\log b}},$$

where $\alpha_0(n)$ was defined in (2).

Lemma 22. *Treating α_0 and f as functions of a real-valued input n , we have*

$$\frac{df}{dn} = \frac{1}{\alpha_0(n)} + \frac{1}{\alpha_0(n)^2} + O\left(\frac{1}{\log^3 n}\right).$$

Proof. Elementary calculus! □

Let us note some simple properties of $\alpha_0(n)$, $\alpha(n) = \lfloor \alpha_0(n) \rfloor$, and $\theta(n)$, defined in (8). Firstly, as noted in the introduction (see (9) and Figure 1),

$$\theta(n) = \alpha_0(n) - \alpha(n) + o(1). \tag{32}$$

In other words, θ is essentially the fractional part of α_0 . Secondly, it is immediate from the definition that $\alpha_0(n)$ is an increasing function of n (for n at least some constant n_0), and that

$$n' \sim n \Rightarrow \alpha_0(n') = \alpha_0(n) + o(1). \tag{33}$$

As outlined in Section 2.2, we will want to compare f' to $1/\alpha$. The following lemma is a convenient form of the statement, allowing us to conveniently consider all n in a suitable range.

[†] Given couplings of X and Y and of Y and Z , that is, desired distributions for (X, Y) and for (Y, Z) , construct (X, Y, Z) by starting with Y and, given the value of Y , taking the appropriate conditional distributions for X and for Z — for example, with conditional independence.

Lemma 23. *If $n \sim n'$, then*

$$\frac{df}{dn} \Big|_{n'} = \frac{1}{\alpha(n)} + \frac{1 - \theta(n)}{\alpha(n)^2} + o\left(\frac{1}{\log^2 n}\right).$$

Proof. The case $n' = n$ is immediate from Lemma 22 and (32). To see the result for $n' \sim n$, note that when we change n by a factor of $1 + o(1)$, from (33) the expression for f' given in Lemma 22 changes by $o(1/\alpha_0^2) = o(1/\log^2 n)$. □

2.5 | Proof of Theorem 6

Proof of Theorem 6. Throughout we fix $p \leq 1 - 1/e^2$ and $\varepsilon > 0$, and consider a positive integer (or rather a sequence) n such that $\mu(n) := \mu_{\alpha(n)}(n) \leq n^{1-\varepsilon}$, or, equivalently,

$$\theta(n) \leq 1 - \varepsilon. \tag{34}$$

We will find the required n^* if n is large enough.

We will apply Lemma 18 with $n^- = n$. Thus from now on we write n^- for our ‘input’ value of n . We choose $\gamma = \gamma(n^-)$ tending to zero sufficiently slowly for various estimates below to hold, and will choose n^+ so that $n^+ \leq n^-(1 + \gamma)$. Thus the desired condition $n^* \sim n^-$ will follow from the conclusion $n^* \in I = [n^-, n^+]$ of Lemma 18.

As noted above, it is immediate from the definition (2) of $\alpha_0(n)$ that (i) α_0 is an increasing function, and (ii) $n^+ \sim n^-$ implies $\alpha_0(n^+) = \alpha_0(n^-) + o(1)$. From (32), the fractional part of $\alpha_0(n^-)$ is $\theta(n^-) + o(1)$, which is at most $1 - \varepsilon/2$, say, if n^- is large enough, which we assume from now on. Thus $\alpha(n^+) = \alpha(n^-)$. In other words, the condition (34) ensures that $\alpha_0(n)$ is not just about to pass through an integer value as we increase n from n^- . Let us write

$$a = \alpha(n^-),$$

noting that in fact $\alpha(n) = a$ for all $n \in I$.

Let $f(n)$ and $\Delta(n)$ be as in Theorem 2. In particular, the error function $\Delta(n)$ is $o(n/\log^2 n)$. We will take

$$\Delta = \max_{n \in I} \Delta(n) = o(n^- / \log^2 n^-). \tag{35}$$

By Lemma 23 and (34), if n^- is large enough we have $f'(n) \geq 1/a + \delta$ for all $n \in I$, where

$$\delta = \frac{\varepsilon}{2a^2} \sim \frac{\varepsilon}{8 \log_b^2 n^-}. \tag{36}$$

So far, we have verified the first two conditions of Lemma 18. For the third, by assumption we have $\mathbb{P}(\chi(G_{n,p}) \in [s_n, t_n]) \geq 0.9$, and our aim is to prove a lower bound on some $\ell_n := t_n - s_n$.

For (23), we take $r(n) = \lfloor \sqrt{\mu(n)} \rfloor$ where $\mu(n) = \mu_a(n) = \binom{n}{a}(1-p)^{\binom{a}{2}}$ as usual; note that here $a = \alpha(n)$. This is clearly an increasing function of n . Moreover, since $a \sim 2 \log_b n$ the first condition of Corollary 21 holds with room to spare. For the second, for any $n \in I$ we have $\mu(n) = n^{\theta(n)}$

by definition, and from (32) and (33) we have $\theta(n) = \theta(n^-) + o(1) \leq 1$, so $\mu(n) \leq n \leq n^{1.98}$. Hence Corollary 21 applies, establishing (23).

Finally, from (35) and (36) we have $\Delta/\delta = o(n^-)$. Also, $r(n^+) \leq \sqrt{\mu(n^+)} \leq \sqrt{n^+}$ as above, so both lower bounds on $n^+ - n^-$ in (24) are $o(n^-)$, and we can choose n^+ to satisfy these bounds as long as $\gamma(n) \rightarrow 0$ slowly enough.

Thus, all conditions of Lemma 18 are met, and we conclude that there is some $n^* \in I$ such that

$$t_{n^*} - s_{n^*} \geq \frac{a\delta}{2} r(n^*) \sim \frac{\varepsilon}{4\alpha(n^*)} \sqrt{\mu(n^*)} \sim \frac{\varepsilon}{8 \log_b n^*} \sqrt{\mu(n^*)} = \frac{\varepsilon \log b}{8 \log n^*} \sqrt{\mu(n^*)}.$$

This establishes the conclusion of Theorem 6 if n is large enough. □

Remark 24. Let us comment briefly on how the error bound in Theorem 2 affects the final bounds we obtain. At first sight, it appears to play little role: the interval length we obtain depends on δ (the gradient difference) and $r = \sqrt{\mu(n)}$. However, via (24), if Δ is large we need to consider a large range I of possible values of n . This not only weakens the conclusion (finding n^* far from n) but can cause a more serious problem: over the interval I both $\mu(n)$ and $\delta(n) = f'(n) - 1/\alpha(n)$ vary, so if our bound on Δ is too weak, we will not obtain a useful lower bound on δ and the argument will fail. Conversely, to obtain a final non-concentration length very close to $n^{1/2}$, we need to consider values of n such that $\mu(n)$ is very close to n , which will only be true over a relatively short interval. So for this we need a better bound on Δ . We revisit this in Section 3.

2.6 | Proof of Theorem 5

Fix $c < 1/2$, and suppose that $[s_n, t_n]$ is a sequence of intervals which contains $\chi(G_{n,p})$ whp, with interval lengths $\ell_n := t_n - s_n$. We will show that there is an integer n^* such that $\ell_{n^*} \geq (n^*)^c$, which suffices to prove Theorem 5.

Let

$$\varepsilon = \frac{1 - 2c}{3} \in \left(0, \frac{1}{3}\right). \tag{37}$$

It is very easy to see that we can find an arbitrarily large integer n such that

$$\theta(n) \in (1 - 2\varepsilon, 1 - \varepsilon). \tag{38}$$

(Recall from (32) that θ is essentially the fractional part of $\alpha_0(n)$, which increases smoothly with n .) By definition of $\theta(n)$, this implies that

$$\mu(n) = n^{\theta(n)} \in (n^{1-2\varepsilon}, n^{1-\varepsilon}). \tag{39}$$

We first consider the case $p \leq 1 - 1/e^2$; the case $p > 1 - 1/e^2$ will follow by some straightforward modifications which we describe at the end of the proof. By Theorem 6, there is an integer $n^* \sim n$ such that

$$\ell_{n^*} \geq C(\varepsilon, p) \frac{\sqrt{\mu(n^*)}}{\log n^*}. \tag{40}$$

As $n^* \sim n$, it follows that $\alpha_0(n^*) = \alpha_0(n) + o(1)$, and so (by (32) and (38)) $\alpha(n^*) = \alpha(n)$. Therefore,

$$\mu(n^*) \sim \mu(n)(n^*/n)^{\alpha(n)} = \mu(n)(1 + o(1))^{O(\log n)} = \mu(n)n^{o(1)} \geq n^{1-2\varepsilon+o(1)} = (n^*)^{1-2\varepsilon+o(1)}.$$

From (40) it follows that if n is large enough, then

$$\ell_{n^*} \geq (n^*)^{\frac{1-2\varepsilon+o(1)}{2}} > (n^*)^{\frac{1-3\varepsilon}{2}} = (n^*)^c,$$

as required.

Now suppose that $p > 1 - 1/e^2$. So far in this paper, whenever we assumed $p \leq 1 - 1/e^2$, it was only to be able to use the estimate for the chromatic number from Theorem 2. More specifically, we only used that we have some estimate $\chi(G_{n,p}) = f(n) + o(\frac{n}{\log^2 n})$ so that the derivative $f'(n)$ is sufficiently larger than $\frac{1}{\alpha}$; namely that

$$f'(n) \geq \frac{1}{\alpha(n)} + \frac{1-\theta}{\alpha(n)^2} + o\left(\frac{1}{\log^2 n}\right). \tag{41}$$

If $p > 1 - 1/e^2$, [15] gives a more complicated expression which also determines $\chi(G_{n,p})$ up to an error term of size $o(\frac{n}{\log^2 n})$. Fortunately, if $\theta = \alpha_0 - \alpha + o(1)$ is close to 1, this estimate takes a simple form which is given in the following lemma.

Lemma 25. Fix $p > 1 - 1/e^2$, and let $u = \frac{2}{\log b} < 1$. For all n such that $\alpha_0(n) - \alpha(n) \geq u$, whp

$$\chi(G_{n,p}) = \frac{n}{\alpha(n) - 1} + o\left(\frac{n}{\log^2 n}\right).$$

Proof. By [15, Theorem 1], whp

$$\chi(G_{n,p}) = \frac{n}{\gamma(n) - x_0} + o\left(\frac{n}{\log^2 n}\right), \tag{42}$$

where $\gamma(n) = \alpha_0(n) - 1 - u$, and, letting $d = \gamma - \lfloor \gamma \rfloor$, x_0 is the smallest non-negative solution to

$$\varphi(x) := (1 - d + x) \log(1 - d + x) + (d - x)(1 - d)/u \leq 0.$$

Suppose that $\alpha_0(n) - \alpha(n) \geq u$. Then $\lfloor \gamma \rfloor = \lfloor \alpha_0 - 1 - u \rfloor = \alpha - 1$ and $d = \gamma - \lfloor \gamma \rfloor = \alpha_0 - \alpha - u$. In particular, $u < 1 - d$. Then for $0 \leq x \leq d$, note that

$$\varphi'(x) = \log(1 - d + x) + 1 - (1 - d)/u \leq 1 - (1 - d)/u < 0.$$

As $\varphi(d) = 0$, this implies that d is the smallest non-negative solution to $\varphi(x) \leq 0$, and so $x_0 = d$. By (42), whp

$$\chi(G_{n,p}) = \frac{n}{\gamma - d} + o\left(\frac{n}{\log^2 n}\right) = \frac{n}{\lfloor \gamma \rfloor} + o\left(\frac{n}{\log^2 n}\right) = \frac{n}{\alpha - 1} + o\left(\frac{n}{\log^2 n}\right). \quad \square$$

Let $u = \frac{2}{\log b}$, and fix $\varepsilon > 0$. For n large enough, if $\theta > u + \varepsilon$ then $\alpha_0 - \alpha = \theta + o(1) \geq u$ and Lemma 25 above applies. Let $f(n) = \frac{n}{\alpha(n)-1}$. If we only consider n in an interval where $\alpha(n)$ is constant — as we did the proof of Theorem 6 — we have

$$f'(n) = \frac{1}{\alpha - 1} = \frac{1}{\alpha} + \frac{1}{\alpha(\alpha - 1)} > \frac{1}{\alpha} + \frac{1}{\alpha^2} \geq \frac{1}{\alpha} + \frac{1 - \theta}{\alpha^2}.$$

Comparing this to (41), all our conclusions from the case $p \leq 1 - 1/e^2$ remain valid as long as $\theta \in (u + \varepsilon, 1)$. To prove the statement of Theorem 5, we can assume c is arbitrarily close to $\frac{1}{2}$, so by (37) we can make ε arbitrarily small. By (38), we can assume that $\theta \in (u + \varepsilon, 1)$. The rest of the proof of Theorem 5 is unchanged from the case $p \leq 1 - 1/e^2$.

3 | PROOF OF THEOREM 8

In this section we prove our final result, Theorem 8. Throughout, we fix $p = \frac{1}{2}$. When we use our assumption (7), we shall state this explicitly. This happens only at one point in the proof of Theorem 8; the assumption is not needed for any of our lemmas. The overall proof strategy is very similar to that we used for Theorem 6, based on our Framework Lemma, Lemma 18. Before turning to the details, let us outline roughly why we choose the parameters that we do, as motivation for the arguments that follow.

We use the same coupling lemma as before which, in terms of the parameters of Lemma 18, leads to choosing $r \approx \sqrt{\mu_\alpha(n)}$, where $\alpha = \alpha(n)$. As Lemma 18 produces an interval of length at least $\alpha\delta r/2$, for a given value of α we want δr to be as large as possible, so we try to choose n so that $\delta(n)\sqrt{\mu_\alpha(n)}$ is as large as possible, which turns out to be when $\mu_\alpha(n) \approx n/\log^2 n$. For why this is optimal, see the Appendix.

As we shall see below, in this range the difference δ between the slope of the chromatic number and $1/\alpha$ is quite small, of order $\Theta(\log \log n / \log^3 n)$. We will consider an interval of values of n differing by at most a factor $1 + x$ where $x \approx 1/\log n$, so that, over the range of n , μ_α , which is roughly proportional to n^α , varies by a constant factor. This means that we need our Δ to be at most roughly $xn\delta \approx n \log \log n / \log^4 n$, to satisfy the first condition in (24).

The error bound from the concentration result (7) is much smaller than this. The trouble is that it applies to the β -bounded chromatic number χ_β , where $\beta = \alpha - 1$, not the chromatic number itself. However, it turns out that, by a first moment argument, we can bound χ from below by $k_\beta - O(\mu_\alpha(n) \log \log n / \log^2 n)$, where k_β is the first moment threshold for β -bounded colourings; see Definition 7. Since (7) gives $\chi_\beta \leq k_\beta + o(n \log \log n / \log^4 n)$ and $\chi \leq \chi_\beta$ by definition, we thus have that $\chi(G_{n,1/2})$ is (just) close enough to k_β for our argument to work.

Throughout the section we consider (sometimes only integer, sometimes real) values of n in a set $W \subset \mathbb{R}$ with the following property: W is a disjoint union of intervals, on each of which $\alpha(n)$ is constant, where $\alpha(n)$ is defined in (2) and (3). In short, $\alpha(n)$ is *locally constant* on W , formally meaning that it has derivative zero.

In the following arguments, there are two relevant ways that n varies: within an interval, and between intervals. When we differentiate with respect to n , we are (by definition) working locally within an interval, and then $\alpha = \alpha(n)$ is constant. On the other hand, for asymptotics (such as the bound $\alpha(n) = O(\log n)$), the variation between intervals is relevant.

Intuitively, one can think of n as very large (so that various asymptotic estimates hold), and in the analysis, in particular the application of the framework lemma, it is only the variation within

an interval that matters. So one should think of $\alpha(n)$ as a constant (derivative zero) that happens to be of logarithmic order. Formally, of course, there is no issue: $\frac{d\alpha}{dn} = 0$ on the set W .

The hardest part of the proof turns out to be understanding the behaviour of (a suitable approximation to) $k_\beta(n)$, where $\beta = \beta(n) = \alpha(n) - 1$. The following lemma, proved in the next section, provides this. Note that we work almost all the time with $\alpha - 1$ rather than α , so to keep the formulae compact we write β for $\alpha - 1$. In fact, although we do not need it here, the same method works with no difficulty for $\alpha - 2$ also. We prove the more general case since it may be useful elsewhere, but the reader may wish to simply consider $\beta = \alpha - 1$.

Lemma 26. *Define $\beta = \beta(n) = \alpha(n) - i$, where $\alpha(n)$ is defined in (3) and $i \in \{1, 2\}$ is constant, and let $W \subset \mathbb{R}$ consist of a disjoint union of intervals on each of which $\alpha(n)$ is constant. For integer $n \in W$, define $k_\beta(n)$ as in (5). Then there is a real-valued function $k^*(n) = k_\beta^*(n)$, defined for all $n \in W$, with the following three properties:*

$$k_\beta(n) = k^*(n) + O(\log^2 n) \text{ for all integer } n \in W,$$

while for all real $n \in W$ we have

$$\frac{n}{k^*(n)} = \alpha(n) + \frac{\log(\mu_{\alpha(n)}(n))}{\log n - \log \log n} - \frac{2}{\log 2} - 1 + O(1/\log n), \tag{43}$$

and

$$\left(\frac{dk^*(n)}{dn}\right)^{-1} = \frac{n}{k^*(n)} + \frac{2}{\log 2} + O(1/\log n).$$

The (somewhat lengthy) proof of Lemma 26 is given in Section 3.1.

Remark 27. The formula (43) may seem slightly mysterious; we make two observations. Firstly, $\alpha(n)$ here can be replaced by any integer $a(n)$ such that $a(n) = \alpha_0(n) + O(1)$, provided we use the same a in both places. This follows from the fact that $\mu_a/\mu_{a-1} = \Theta(\log n/n)$ for $a = \alpha_0(n) + O(1)$. Secondly, a straightforward but rather tedious calculation shows that, for $a = \alpha_0(n) + O(1)$, we have

$$a(n) + \frac{\log(\mu_{a(n)}(n))}{\log n - \log \log n} = \alpha_0(n) + \left(\frac{2}{\log 2} - \frac{1}{2}\right) \frac{\log \log n}{\log n} + O(1/\log n). \tag{44}$$

To make sense of this note that one can interpolate the definition of $\mu_a(n)$ to non-integer values of a in a natural way. As noted above, the left-hand side is then (roughly) constant for a near α_0 . Substituting in $a = \alpha_0$, we expect μ_{α_0} to be close to 1. This explains (44) apart from the $c \log \log n / \log n$ term. This term is only there because we have taken a simple definition of α_0 , rather than solve $\mu_{\alpha_0} = 1$ very precisely; if we were to do so, we would simply have $\alpha_0(n) + O(1/\log n)$ here, but there would be minor additional complications in other formulae. Finally, we do not use this expression in (43) because both in the proof and in the application, it is easier to work with α and μ_α than with α_0 .

Our next lemma, proved in Section 3.2, is the promised lower bound on $\chi(G_{n,1/2})$ in terms of k_β .

Lemma 28. *Let $\varepsilon > 0$ be constant. Suppose that $\mu_{\alpha(n)}(n) = \Theta(n/\log^2 n)$. Then, whp,*

$$\chi(G_{n,1/2}) \geq k_\beta(n) - (1 + \varepsilon)\mu_{\alpha(n)}(n) \frac{\log \log n}{c_0 \log^2 n},$$

where $c_0 = 2/\log 2$, and $k_\beta(n)$ is defined in (5), with $\beta = \beta(n) = \alpha(n) - 1$.

At this point we are ready to prove Theorem 8, subject to the (in the first case lengthy) proofs of Lemmas 26 and 28, given in next two sections.

Proof of Theorem 8. Set $c_0 = 2/\log 2$, and let c_1 be a positive constant with $c_1 < \frac{1}{5c_0^2}$. We consider the set

$$W := \left\{ n \in \mathbb{R} : \frac{c_1 n}{e \log^2 n} \leq \mu_{\alpha(n)}(n) \leq \frac{c_1 n}{\log^2 n} \right\}. \tag{45}$$

This set is easily seen to be a disjoint union of intervals, one for each value of $\alpha(n)$. Our aim is to show the existence of at least one n in each interval (apart perhaps from the first few) such that $\chi(G_{n,1/2})$ is not too concentrated.

First, we consider the length of a single interval $I = [n^-, n^+] \subset W$. With $a = \alpha(n)$ constant (as it is over I), $\mu_a(n)$ is proportional to $\binom{n}{a}$, which is asymptotically $n^a/a!$. It follows easily that $n^+ = (1 + \varepsilon)n^-$ for some ε such that $(1 + \varepsilon)^a \sim e$. This gives $\varepsilon \sim 1/a \sim 1/(c_0 \log n^-)$, say. Thus

$$n^+ - n^- \sim n^- / (c_0 \log n^-).$$

We will apply Lemma 18 to each interval, with $f(n) = k^*(n)$, where $k^*(n)$ is as in Lemma 26. Let $\beta = \beta(n) = \alpha(n) - 1$, which is constant on each interval. By Lemma 28 and the definition of W , whp we have

$$\chi(G_{n,1/2}) \geq k_\beta(n) - (1 + o(1))\mu_{\alpha(n)}(n) \frac{\log \log n}{c_0 \log^2 n} \geq k_\beta(n) - (1 + o(1)) \frac{c_1 n \log \log n}{c_0 \log^4 n}.$$

We have $\chi_\beta(G_{n,1/2}) = k_\beta(n) + o(n \log \log n / \log^4 n)$ whp by our assumption (7).[†] Thus, whp

$$k_\beta(n) - (1 + o(1)) \frac{c_1 n \log \log n}{c_0 \log^4 n} \leq \chi(G_{n,1/2}) \leq \chi_\beta(G_{n,1/2}) \leq k_\beta(n) + \frac{c_1 n \log \log n}{c_0 \log^4 n}.$$

From Lemma 26 we have $k^*(n) - k_\beta(n) = O(\log^2 n)$ for integer $n \in W$, so it follows that for $n \in I$ we have $\chi(G_{n,1/2}) \in [f(n) - \Delta, f(n) + \Delta]$ whp, for some Δ satisfying

$$\Delta \sim \frac{c_1 n^- \log \log n^-}{c_0 \log^4 n^-}.$$

This establishes the first condition (20) of Lemma 18.[‡]

[†] This is the only place in the proof where we use (7); the lemmas stated in this section do not rely on it.

[‡] The reader may wonder why we take $f(n) = k^*(n)$ rather than $f(n) = k_\beta(n)$. The reason is that we do not know precisely enough how the latter varies.

We set $a = \alpha(n)$, which is constant over the interval we are considering. For $n \in W$, by Lemma 26 we have

$$\begin{aligned} \frac{1}{f'(n)} &= \alpha(n) + \frac{\log(\mu_{\alpha(n)}(n))}{\log n - \log \log n} - 1 + O(1/\log n) \\ &= a + \frac{\log n - 2 \log \log n + O(1)}{\log n - \log \log n} - 1 + O(1/\log n) \\ &= a - \frac{\log \log n}{\log n} + O(1/\log n). \end{aligned}$$

Since $a \sim c_0 \log n$, it follows (using $(a - \varepsilon)^{-1} = a^{-1}(1 - \varepsilon/a)^{-1} = a^{-1} + \varepsilon a^{-2} + \dots$) that

$$f'(n) = \frac{1}{a} + (1 + o(1)) \frac{\log \log n}{c_0^2 \log^3 n},$$

so (21) holds for all $n \in I$ for some δ satisfying

$$\delta \sim \frac{\log \log n^-}{c_0^2 \log^3 n^-}.$$

As usual (22) is part of our assumption; we assume $G_{n,1/2}$ is concentrated like this and our aim is to give a lower bound on $t_n - s_n$ for some n .

As before, condition (23) follows from our coupling result, Corollary 21, taking $r(n) = \lfloor \sqrt{\mu_{\alpha(n)}} \rfloor$, say.

Now $\Delta/\delta \sim c_0 c_1 n^- / \log n^-$, while $ar(n^+)$ is $O(\sqrt{n^-})$. By choice of c_1 we have $1/c_0 > 5c_0 c_1$, so it follows that the inequalities in (24) hold for large enough n .

Thus Lemma 18 implies that for some n in each interval (except perhaps for the first $O(1)$), we have

$$t_n - s_n \geq \frac{a\delta r(n)}{2} \geq (1 + o(1))c_0 \log n \frac{\log \log n}{c_0^2 \log^3 n} \frac{\sqrt{c_1 n}}{2\sqrt{e} \log n} \sim c_2 \frac{\sqrt{n} \log \log n}{\log^3 n}$$

where $c_2 = \sqrt{c_1}/(2c_0\sqrt{e})$, and we have replaced n^- by n since $n \sim n^-$. □

3.1 | Proof of Lemma 26

In this section we prove Lemma 26. This will take some time. In principle, this is a matter of calculation, but it seems to require considerable work, and several tricks, to get the calculations to come out to the required accuracy. For the reader to refer back to later, we collect in Table 1 some notation used in this and the next section.

The t -bounded first moment threshold $k_t(n)$ is defined in terms of $E_{n,k,t}$, the expected number of unordered t -bounded k -colourings of $G_{n,1/2}$. One key idea of the proof is to replace $E_{n,k,t}$ by a simpler quantity, and to define k^* as the threshold for this simpler estimate to cross 1. We will simplify in three simple steps, proved together in one lemma (Lemma 29): (i) we replace the expected number of colourings with a given profile (see below) by a simpler formula, (ii) we replace the sum over profiles by a maximum, and (iii) we replace the maximum over integer-valued profiles (a complicated set) by the maximum over a certain region in \mathbb{R}^t .

TABLE 1 The various functions involved in defining and approximating the t -bounded expectation threshold. While L_0 is the key approximation to $\log(E_{n,k,t})$, in different parts of the analysis it turns out to be much simpler to consider the transformed functions \widehat{L}_0 and \widetilde{L}_0 .

$E_{n,k,t}$	Expected number of unordered t -bounded k -colourings
$k_t(n)$	Threshold where $E_{n,k,t}$ reaches 1
$L_0(n, k, t)$	Approximation to $\log(E_{n,k,t})$ defined in (48)
$\widehat{L}_0(n, k, t)$	$\frac{1}{k}L_0(n, k, t)$
$\widetilde{L}_0(\rho, k, t)$	Defined by $\widehat{L}_0(n, k, t) = \widetilde{L}_0(n/k, k, t)$
$k^*(n)$	Defined by solving $L_0 = 0$ (or $\widehat{L}_0 = 0$)

To state and prove the lemma we need some notation. Let $\pi = (n_i)_{i=1}^t$ denote a t -bounded profile, where n_i represents the number of colour classes with i vertices. Let $P_{n,k,t}$ denote the set of all profiles π satisfying

$$n_i \geq 0, \quad \sum_{i=1}^t n_i = k \quad \text{and} \quad \sum_{i=1}^t in_i = n. \tag{46}$$

Thus $P_{n,k,t}$ consists of all profiles corresponding to t -bounded k -colourings. Extending to real values, given positive reals $k < n$ and a positive integer t , let

$$P_{n,k,t}^0 = \{ (n_i)_{i=1}^t \in \mathbb{R}^t : (46) \text{ holds} \}.$$

Two key quantities appearing in many places in our calculation will be

$$d_i := 2^{\binom{i}{2}} i! \tag{47}$$

and

$$L_0(n, k, t) := \sup_{\pi \in P_{n,k,t}^0} \left\{ n \log n - n + k - \sum_{i=1}^t n_i \log(n_i d_i) \right\}. \tag{48}$$

As we now show, the latter is a good approximation to $\log(E_{n,k,t})$, where $E_{n,k,t}$ is the expected number of unordered t -bounded k -colourings of $G_{n,1/2}$ (see Definition 7).

Lemma 29. *Suppose that $t = t(n) = O(\log n)$. For all (large enough) n and for all k with $1 < n/k < t$ we have*

$$\log(E_{n,k,t}) = L_0(n, k, t) + O(\log^4 n).$$

Proof. For a given profile π , let E_π be the expected number of unordered colourings with this profile, so by definition

$$E_{n,k,t} = \sum_{\pi \in P_{n,k,t}} E_\pi.$$

Since the order of the parts does not matter, there are

$$\frac{1}{\prod_i n_i!} \frac{n!}{\prod_i i!^{n_i}}$$

ways to partition $[n]$ into k parts with n_i of size i for each i (the second fraction is the relevant multinomial coefficient). Such a partition is indeed an unordered k -colouring if and only if there are no edges of $G_{n,1/2}$ within the parts. Hence

$$E_\pi = \frac{n!}{\prod_i n_i!} \frac{1}{\prod_i i!^{n_i}} 2^{-\sum_i n_i \binom{i}{2}} = \frac{n!}{\prod_i n_i!} \prod_i d_i^{-n_i}.$$

Let

$$L_\pi := n \log n - \sum_i n_i \log(n_i d_i) - n + k.$$

Then using Stirling's formula it is easy to see that for any $\pi \in P_{n,k,t}$ we have

$$\log(E_\pi) = L_\pi + O(\log^2 n).$$

Indeed, this follows by absorbing the (logarithm of) all $\sqrt{2\pi m}$ factors into the error term.

There are at most $(n + 1)^t = \exp(O(\log^2 n))$ possible profiles, so $E_{n,k,t}$ is within this factor of $\max_\pi E_\pi$. Hence

$$\log(E_{n,k,t}) = \max_{\pi \in P_{n,k,t}} L_\pi + O(\log^2 n). \tag{49}$$

Now $P_{n,k,t} \subset P_{n,k,t}^0$, so the inequality $\max_{\pi \in P_{n,k,t}} L_\pi \leq L_0(n, k, t)$ holds trivially. It remains to show the reverse inequality, up to a small error term. For this, let $\pi = (n_i)_{i=1}^t \in P_{n,k,t}^0$ be arbitrary. Our aim is to find a profile $\pi' \in P_{n,k,t}$ with $L_{\pi'}$ not too far from L_π . To do so, we modify π in a series of small steps. Firstly, round each (non-integer) n_i either up or down to the nearest integer, choosing whether to round up or down in such a way that after all such roundings $\sum n_i$ is unchanged. At this point, $\sum i n_i$ has changed by no more than $\sum_{i=1}^t i \leq t^2$. We obtain π' by making a number of further changes, each of which consists of altering the size of one class by 1, that is, decreasing some n_i by 1 and increasing either n_{i-1} or n_{i+1} by 1; clearly we can fix the error in $\sum i n_i$ by at most t^2 such changes. In total, we have made $O(t^2)$ small changes, each of which consists of altering a single value n_i by at most 1.

Now each d_i is at most $2^{t^2} t! = \exp(O(t^2))$. Also

$$\frac{d}{dn_i} n_i \log(n_i d_i) = \log(n_i d_i) + 1,$$

which is thus $O(\log n + t^2) = O(\log^2 n)$ for $1 \leq n_i \leq n$. It is easy to check that $n_i \log(n_i d_i)$ is $O(\log^2 n)$ for $0 \leq n_i \leq 1$. It follows that each of the changes above (changing a single n_i by at most 1) changes $\sum n_i \log(n_i d_i)$ by at most $O(\log^2 n)$. The remaining terms in L_π are the same for π' as

for π , so we conclude that

$$|L_\pi - L_{\pi'}| = O(t^2 \log^2 n) = O(\log^4 n).$$

Hence L_0 is within $O(\log^4 n)$ of the maximum over (integer) profiles π , which, combined with (49), gives the result. \square

At this point it will be convenient to rescale in two ways: we replace each n_i by $p_i = n_i/k$, the fraction of colour classes having size i (at least, this is the interpretation when n_i is an integer). We will also divide the logarithm we are considering by k . To formalise this, for t a positive integer and ρ a real number with $1 < \rho < t$ define

$$\tilde{P}_{\rho,t} = \left\{ (p_i)_{i=1}^t \in [0, 1]^t : \sum_i p_i = 1 \quad \text{and} \quad \sum_i i p_i = \rho \right\}.$$

When $\rho = n/k$, this is exactly the set $P_{n,k,t}^0$ rescaled by replacing each n_i by $p_i = n_i/k$. Note that $\tilde{P}_{\rho,t}$ is simply the set of probability distributions (or probability mass functions) on $[t]$ with expectation ρ .

Let t be a positive integer, and ρ and k positive reals with $1 < \rho < t$. For $\mathbf{p} = (p_i)_{i=1}^t \in \tilde{P}_{\rho,t}$, let

$$\tilde{L}(\rho, k, \mathbf{p}) := \rho \log(\rho k) - \log k - \rho + 1 - \sum_i p_i \log(p_i d_i), \tag{50}$$

and define

$$\tilde{L}_0(\rho, k, t) := \sup_{\mathbf{p} \in \tilde{P}_{\rho,t}} \tilde{L}(\rho, k, \mathbf{p}). \tag{51}$$

Lemma 30. *If t is a positive integer and n and k are positive reals with $1 < n/k < t$ then*

$$L_0(n, k, t) = k \tilde{L}_0(\rho, k, t),$$

where $\rho = n/k$.

Proof. This is simply a matter of rescaling: for $\pi \in P_{n,k,t}^0$, letting $p_i = n_i/k$ we have

$$\begin{aligned} \frac{L_\pi}{k} &= \frac{n}{k} \log n - \sum_i \frac{n_i}{k} \log(n_i d_i) - \frac{n}{k} + 1 = \rho \log(\rho k) - \sum_i p_i \log(k p_i d_i) - \rho + 1 \\ &= \rho \log(\rho k) - \sum_i p_i \log(p_i d_i) - \log k - \rho + 1, \end{aligned}$$

since $\sum p_i = 1$. The result follows from the bijection between $P_{n,k,t}^0$ and $\tilde{P}_{\rho,t}$ given by $p_i = n_i/k$. \square

Corollary 31. *Suppose that $t = t(n) = O(\log n)$. For all (large enough) n and for all k with $1 < k < n/t$ we have*

$$\log(E_{n,k,t}) = k\tilde{L}_0(\rho, k, t) + O(\log^4 n),$$

where $\rho = n/k$.

Proof. Immediate from Lemmas 29 and 30. □

In the next few lemmas our aim is to study the functions $\tilde{L}(\rho, k, \mathbf{p})$ and $\tilde{L}_0(\rho, k, t)$ defined in (50) and (51). Although, as in Corollary 31, we will eventually evaluate \tilde{L}_0 at $(n/k, k, t)$, where n and k are integers, for the moment this is irrelevant. We are simply studying the functions defined in (50) and (51), where ρ and k are real inputs, t is an integer, and \mathbf{p} is a vector in $\tilde{P}_{\rho,t}$. In particular, n appears nowhere in these definitions.

We start by studying the location and value of the maximum of $\tilde{L}(\rho, k, \mathbf{p})$ over \mathbf{p} .

Lemma 32. *Let $1 < \rho < t$, where t is an integer. Then, for any real $k > 1$, there is a unique $\mathbf{p} = \mathbf{p}_{\rho,t} \in \tilde{P}_{\rho,t}$ maximising $\tilde{L}(\rho, k, \mathbf{p})$. This maximising \mathbf{p} is independent of k , and is given by*

$$p_i = e^{x+iy} d_i^{-1} \tag{52}$$

for $1 \leq i \leq t$, where $x = x_t(\rho)$ and $y = y_t(\rho)$ satisfy

$$\sum_{i=1}^t e^{x+iy} d_i^{-1} = 1 \tag{53}$$

and

$$\sum_{i=1}^t i e^{x+iy} d_i^{-1} = \rho. \tag{54}$$

Furthermore,

$$\tilde{L}_0(\rho, k, t) = \rho \log(\rho k) - \log k - \rho + 1 - x - \rho y. \tag{55}$$

Proof. Throughout the proof k, t and ρ are fixed, and we are maximising only over $\mathbf{p} \in \tilde{P}_{\rho,t}$. Thus, the only term in $\tilde{L}(\rho, k, \mathbf{p})$ that varies is the term

$$f(\mathbf{p}) = \sum_{i=1}^t -p_i \log(p_i d_i).$$

Note that k does not appear in this expression. In contrast, ρ appears implicitly via the constraint $\sum i p_i = \rho$. Hence the location of the maximum will depend on ρ and t , but not on k .

Now $-x \log(xd)$ is strictly concave as a function of x , so viewed as a function on $[0, 1]^t$, $f(\mathbf{p})$ is a sum of concave functions and hence concave. It is thus concave also on the domain $\tilde{P}_{\rho,t}$. Thus

$f(\mathbf{p})$, and hence $\tilde{L}(\rho, k, \mathbf{p})$, has a unique maximiser \mathbf{p} . This maximiser lies in the interior of $\tilde{P}_{\rho,t}$, since the derivative of $-x \log(xd)$, namely $-\log(xd) - 1$, approaches infinity as x approaches 0.[†]

The second statement now follows easily by the method of Lagrange multipliers, viewing $f(\mathbf{p})$ as a function on $[0, 1]^t$, which we wish to maximise subject to the constraints

$$\sum p_i = 1 \quad \text{and} \quad \sum i p_i = \rho. \tag{56}$$

Indeed, we have

$$\frac{\partial f}{\partial p_i} = -\log(p_i d_i) - 1,$$

so at the maximum there are λ and μ such that

$$-\log(p_i d_i) - 1 = \lambda + \mu i$$

for $1 \leq i \leq t$. Rearranging and setting $y = -\mu$ and $x = -\lambda - 1$ gives (52). The relations (53) and (54) follow immediately from the constraints (56).

Finally, to obtain (55) we substitute (52) into the definition of \tilde{L}_0 , noting that for this specific \mathbf{p} we have

$$\sum p_i \log(p_i d_i) = \sum p_i (x + iy) = x + \rho y,$$

again using (56). □

It is easy to see that, for a given integer t , (53) and (54) define x and y uniquely as functions of ρ (where $1 < \rho < t$), and furthermore that these functions $x(\rho) = x_t(\rho)$ and $y(\rho) = y_t(\rho)$ are (infinitely) differentiable. Indeed, dividing (54) by (53) gives

$$\frac{\sum_{i=1}^t i e^{iy} d_i^{-1}}{\sum_{i=1}^t e^{iy} d_i^{-1}} = \rho.$$

The left-hand side is strictly increasing and (infinitely) differentiable as a function of y , and tends to 1 or to t as y tends to $-\infty$ or $+\infty$, respectively. Having solved this equation to determine $y_t(\rho)$, we may use (53), say, to find $x_t(\rho)$.

We next investigate the derivatives of L_0 .

Lemma 33. *For t fixed the 2-variable function $\tilde{L}_0(\rho, k, t)$ has partial derivatives*

$$\frac{\partial}{\partial k} \tilde{L}_0(\rho, k, t) = \frac{\rho - 1}{k} \quad \text{and} \quad \frac{\partial}{\partial \rho} \tilde{L}_0(\rho, k, t) = \log(\rho k) - y_t(\rho).$$

[†] To spell this out completely, suppose that at the maximum some $p_i = 0$. To obtain a contradiction it suffices to find a direction that we can move within $\tilde{P}_{\rho,t}$ in which p_i increases. Then for a small enough change in this direction, the increase in the term $-p_i \log(p_i d_i)$ will outweigh the decrease in any other terms. Such a direction exists because $\tilde{P}_{\rho,t}$ certainly contains a point \mathbf{p}' with $p'_i > 0$, so we may choose the direction from \mathbf{p} to \mathbf{p}' .

Proof. We use (55), recalling that with t fixed $x = x_t(\rho)$ and $y = y_t(\rho)$ depend only on ρ , not on k . The formula for the k -derivative is immediate (since then x , y and ρ are constants). For the ρ -derivative by elementary calculus we have

$$\frac{\partial}{\partial \rho} \tilde{L}_0(\rho, k, t) = \log(\rho k) - \frac{d}{d\rho}(x + \rho y) = \log(\rho k) - \frac{dx_t(\rho)}{d\rho} - \rho \frac{dy_t(\rho)}{d\rho} - y_t(\rho).$$

At this point something miraculous-seeming happens: if we differentiate the constraint (53) with respect to ρ we obtain

$$\sum_{i=1}^t \left(\frac{dx_t(\rho)}{d\rho} + i \frac{dy_t(\rho)}{d\rho} \right) e^{x_t(\rho) + iy_t(\rho)} d_i^{-1} = 0,$$

which, using (53) and (54), simplifies to

$$\frac{dx_t(\rho)}{d\rho} + \rho \frac{dy_t(\rho)}{d\rho} = 0.$$

Combined with the formula above, this gives the result. □

So far, it was convenient to work in terms of ρ and k rather than n and k , because certain key functions then depended only on ρ . However, in the end we wish to find a threshold k^* as a function of n , so we now undo this change of variables. Noting/recalling that the definitions (51) and (48) of $\tilde{L}_0(\rho, k, t)$ and $L_0(n, k, t)$ do not require n and k to be integers, for t a positive integer and n and k positive reals with $1 < n/k < t$, define

$$\hat{L}_0(n, k, t) := \tilde{L}_0(n/k, k, t), \tag{57}$$

so, by Lemma 30,

$$L_0(n, k, t) = k\tilde{L}_0(n/k, k, t) = k\hat{L}_0(n, k, t). \tag{58}$$

Lemma 34. *For t fixed the 2-variable function $\hat{L}_0(n, k, t)$ has partial derivatives*

$$\frac{\partial}{\partial k} \hat{L}_0(n, k, t) = -\frac{n}{k^2}(\log n - y_t(n/k)) + \frac{n}{k^2} - \frac{1}{k} \quad \text{and} \quad \frac{\partial}{\partial n} \hat{L}_0(n, k, t) = \frac{\log n - y_t(n/k)}{k}.$$

Proof. This is straightforward calculus: using (57) and the Chain Rule we have

$$\frac{\partial}{\partial k} \hat{L}_0 = -\frac{n}{k^2} \frac{\partial}{\partial \rho} \tilde{L}_0 + \frac{\partial}{\partial k} \tilde{L}_0,$$

and

$$\frac{\partial}{\partial n} \hat{L}_0 = \frac{1}{k} \frac{\partial}{\partial \rho} \tilde{L}_0.$$

The result thus follows from Lemma 33. □

Our next aim is to find the value of y ; it turns out that a fairly crude bound is enough, and for this we can use a ‘soft’ argument, rather than trying to exactly solve the constraints (53) and (54).

Lemma 35. *Suppose that $\rho = t - \Theta(1)$ and $\rho \geq 2$.[†] Then*

$$y_t(\rho) = \log(t2^t) + O(1).$$

Proof. Note that $y_t(\rho)$ is defined for any positive integer t and any real ρ with $1 < \rho < t$. The statement is that if we restrict the parameter space to (ρ, t) such that $\rho \geq 2$ and $c < t - \rho < C$ for some constants $C > c > 0$, then the difference between $y_t(\rho)$ and $\log(t2^t)$ is bounded.

Fix, for the moment, ρ and t with $1 < \rho < t$, and let $y = y_t(\rho)$. Recall that $(p_i)_{i=1}^t$ with p_i defined by (52) is a probability distribution on $[t]$ with mean ρ .

For $2 \leq i \leq t$, from (52) we have

$$r_i := \frac{p_i}{p_{i-1}} = e^y \frac{d_{i-1}}{d_i} = \frac{e^y}{i2^{i-1}}, \tag{59}$$

recalling the definition (47) of d_i . In particular, (r_i) is a decreasing function of i , so the sequence (p_i) is unimodal. Furthermore, for $i = t - O(1)$ we have

$$r_i = e^y \Theta\left(\frac{1}{t2^t}\right), \tag{60}$$

where the implicit constants do not depend on t or ρ . We claim that, uniformly over (ρ, t) with $t - \rho = \Theta(1)$, we have $r_t = r_t(\rho, t) = \Theta(1)$; then (60) gives the result.

To establish the claim suppose first (for a contradiction) that for fixed c, C there exist (ρ, t) with $c < t - \rho < C$ such that $r_t = r_t(\rho, t)$ is arbitrarily large. If $r_t \geq D$ then $r_i \geq D$ for all $2 \leq i \leq t$, so (for large D) the sequence p_i is rapidly increasing and the mean ρ of this probability distribution is very close to t . We thus obtain a contradiction for some $D = D(c)$.

Next suppose that, with $c < t - \rho < C$, we may choose ρ and t such that r_t is arbitrarily small. Since $r_i = \Theta(r_t)$ for $i \geq t - 2C$, say, r_i is also small (say $< 1/2$) for $i \geq t - 2C$. Thus p_i decreases rapidly on $[t - 2C, t]$. If $t \geq 2C + 1$ then it follows that the mean of this probability distribution is less than $t - C$, a contradiction. If $t \leq 2C + 1$ then we conclude that p_i decreases rapidly on the whole domain $[1, t]$, which implies that the mean is less than 2, again contradicting our assumptions. □

We also give a useful bound on $x + ty$, in a slightly more general form.

Lemma 36. *Suppose that $\rho \geq 2$, that $\rho = t - \Theta(1)$, and that $a = t + O(1)$ is a positive integer. Then*

$$x_t(\rho) + ay_t(\rho) = \log(d_a) + O(1),$$

where d_a is defined in (47).

[†] This condition will be irrelevant in the end; ρ and t will be order $\log n$. It is needed only to rule out values of ρ very close to 1.

Proof. We continue the argument in the proof of the previous lemma. As shown there, defining r_i as in (59), we have $r_i = \Theta(r_t) = \Theta(1)$ for $i = t - O(1)$. Since (p_i) is a probability distribution on $[t]$ with mean ρ , it follows that $p_t = \Theta(1)$. Indeed, if $p_t = o(1)$ then we would have $p_i = o(1)$ for $i = t - O(1)$, contradicting that the mean ρ is within $O(1)$ of t .

Now (purely as a notational convenience) extend the definition of p_i to $i > t$ also, taking $p_i = e^{x+iy} d_i^{-1}$ as in (52), with $x = x_t(\rho)$ and $y = y_t(\rho)$. Then (59) holds for $i > t$ too, and (from this equation) we have $r_i = \Theta(r_t)$ for $i = t + O(1)$. Hence $p_a = \Theta(p_t) = \Theta(1)$. Taking logs,

$$\log(p_a) = x_t(\rho) + ay_t(\rho) - \log(d_a) = O(1),$$

giving the result. □

We will be interested in the β -bounded chromatic number where $\beta = \alpha(n) - i = \alpha_0(n) + O(1)$, for $i = 1$ (the important case for us) or $i = 2$. It will turn out that the relevant values of ρ (the average colour-class size) are of the form $\beta - \Theta(1)$. The next corollary gives the value of y in this key case.

Corollary 37. *Suppose that $t = t(n) = \alpha_0(n) + O(1)$ is an integer. Uniformly over all n and all real $\rho \geq 2$ such that $t - \rho = \Theta(1)$ we have*

$$y_t(\rho) = 2 \log n - \log \log n + O(1). \tag{61}$$

Proof. We apply Lemma 35, noting that, recalling (2), for $t = \alpha_0(n) + O(1)$ we have $2^t = \Theta(n^2 / \log^2 n)$. □

Using this value of $y_t(\rho)$, and Lemma 36, we can estimate \widehat{L}_0 (or \widetilde{L}_0 , which is the same function reparametrised). Recall that \widehat{L}_0 is defined by dividing L_0 (a good approximation to the logarithm of the expected number of t -bounded k -colourings) by k , so the $+O(1)$ error below corresponds in the end to a factor $\exp(O(k)) = \exp(O(n / \log n))$.

Lemma 38. *Suppose that $k < n$ are positive reals, and $t \leq a$ are positive integers, such that $a, t = \alpha_0(n) + O(1)$ and $2 \leq n/k = t - \Theta(1)$. Then*

$$\widehat{L}_0(n, k, t) = \left(a - \rho - 1 - \frac{2}{\log 2} \right) (\log n - \log \log n) + \log(\mu_a(n)) + O(1), \tag{62}$$

where $\rho = n/k$ and, as usual, $\mu_a(n) = \binom{n}{a} 2^{-\binom{a}{2}}$ is the expected number of independent a -sets in $G_{n,1/2}$.

Proof. Let $\rho = n/k$, so by assumption $\rho \geq 2$ and $t - \rho = \Theta(1)$. By the formula (55) from Lemma 32 we have

$$L := \widehat{L}_0(n, k, t) = \widetilde{L}_0(\rho, k, t) = \rho \log n - \log k - \rho + 1 - (x + ay) + (a - \rho)y,$$

where $x = x_t(\rho)$ and $y = y_t(\rho)$. By Lemma 36 we have

$$x + ay = \log(d_a) + O(1).$$

Since

$$\mu_a(n) = \binom{n}{a} 2^{-\binom{a}{2}} \sim \frac{n^a}{a! 2^{\binom{a}{2}}} = \frac{n^a}{d_a},$$

we have $\log(d_a) = a \log n - \log(\mu_a(n)) + o(1)$. Thus

$$\begin{aligned} L &= \rho \log n - \log k - \rho - a \log n + \log(\mu_a(n)) + (a - \rho)y + O(1) \\ &= (a - \rho)(y - \log n) + \log(\mu_a(n)) - \log k - \rho + O(1). \end{aligned} \tag{63}$$

Now by assumption

$$\rho = n/k = t - \Theta(1) = \alpha_0(n) + O(1) = 2 \log_2 n - 2 \log_2 \log n + O(1),$$

since $\log_2 n = \Theta(\log n)$. Thus

$$\rho = \frac{2}{\log 2}(\log n - \log \log n) + O(1).$$

Also, crudely, $k = n/\rho = \Theta(n/\log n)$, so

$$\log k = \log n - \log \log n + O(1).$$

Substituting the last two formulae into (63), we have

$$L = (a - \rho)(y - \log n) + \log(\mu_a(n)) - \left(1 + \frac{2}{\log 2}\right)(\log n - \log \log n) + O(1).$$

Finally, note that $a - \rho = O(1)$ and that, from (61), $y = 2 \log n - \log \log n + O(1)$. Thus

$$L = \left(a - \rho - 1 - \frac{2}{\log 2}\right)(\log n - \log \log n) + \log(\mu_a(n)) + O(1),$$

as claimed. □

We can also use the value of y from Corollary 37 to give approximate bounds on the partial derivatives of L_0 and $\hat{L}_0 = L_0/k$.

Corollary 39. *Suppose that $t = t(n) = \alpha_0(n) + O(1)$ is an integer. Uniformly over all $k \leq n/2$ such that $k = n/(t - \Theta(1))$ we have*

$$\frac{\partial}{\partial k} \hat{L}_0(n, k, t) = \Theta\left(\frac{\log^3 n}{n}\right), \quad \frac{\partial}{\partial n} \hat{L}_0(n, k, t) = -\Theta\left(\frac{\log^2 n}{n}\right)$$

and

$$\frac{\partial}{\partial k} L_0(n, k, t) = \frac{2}{\log 2} \log^2 n + O(\log n \log \log n).$$

Proof. Note that the dependence of t on n is only relevant for the asymptotics; by definition of partial derivative, we hold t constant when differentiating. Also, in the end $t = \alpha(n) - 1$ or $\alpha(n) - 2$ will be locally constant. The bounds on the partial derivatives of \widehat{L}_0 follow by substituting the value $y = 2 \log n + O(\log \log n) \sim 2 \log n$ from (61) into the conclusion of Lemma 34, noting that $n/k = \alpha_0(n) + O(1) = 2 \log_2 n + O(\log \log n) \sim 2 \log_2 n$.

For $L_0(n, k, t) = k\widehat{L}_0(n, k, t)$, calculating slightly more precisely,

$$\begin{aligned} \frac{\partial}{\partial k} L_0 &= \frac{\partial}{\partial k} (k\widehat{L}_0) = \widehat{L}_0 + k \frac{\partial}{\partial k} \widehat{L}_0 = \widehat{L}_0 + \frac{n}{k} (y_t(n/k) - \log n) + \frac{n}{k} - 1 \\ &= \widehat{L}_0 + \frac{n}{k} \log n + O(\log n \log \log n) = \widehat{L}_0 + \frac{2}{\log 2} \log^2 n + O(\log n \log \log n). \end{aligned}$$

The result follows since $\widehat{L}_0(n, k, t) = O(\log n)$ by Lemma 38. □

For the rest of the section we consider a function $\beta(n)$ satisfying the following assumptions; the upper bound on β is of no particular significance.

Assumption 40. The function β is defined on a subset W of \mathbb{R} which is a union of intervals, and is constant on each interval. Furthermore, for some constant $\varepsilon > 0$ we have

$$\alpha_0(n) - 1 - \frac{2}{\log 2} + \varepsilon \leq \beta(n) \leq \alpha_0(n) + 100$$

for all large enough n .

Note in the assumptions of Lemma 26, we specified $\beta(n) = \alpha(n) - 1$ or $\beta(n) = \alpha(n) - 2$. These both satisfy Assumption 40, since $\beta \geq \alpha_0 - 3$ and $2/\log 2 > 2$. For $n \in W$ let

$$I_n = \left[\frac{n}{\beta - \varepsilon/4}, \frac{n}{\alpha_0 - 100} \right].$$

Recall that $L_0(n, k, t) = k\widehat{L}_0(n, k, t)$, so one is zero if and only if the other is.

Lemma 41. For each large enough (real) $n \in W$ there is a unique $k^* = k^*(n) \in I_n$ such that

$$\widehat{L}_0(n, k^*(n), \beta(n)) = 0 = L_0(n, k^*(n), \beta(n)). \tag{64}$$

Furthermore,

$$\frac{n}{k^*(n)} = \alpha(n) - 1 - \frac{2}{\log 2} + \frac{\log(\mu_{\alpha(n)}(n))}{\log n - \log \log n} + O(1/\log n),$$

and if n is an integer then $k^*(n) - k_\beta(n) = O(\log^2 n)$.

Proof. Keeping n fixed, from Corollary 39, if n is large enough, then $\widehat{L}_0(n, k, \beta(n))$ is strictly increasing as a function of $k \in I_n$, with derivative $\Theta(\log^3 n/n)$. This implies uniqueness of $k^*(n)$ once we

show existence. Define $k_0 = k_0(n)$ by

$$\frac{n}{k_0} = \alpha(n) - 1 - \frac{2}{\log 2} + \frac{\log(\mu_{\alpha(n)}(n))}{\log n - \log \log n}.$$

Then, recalling that $\mu_{\alpha(n)}(n) = n^{\alpha(n) - \alpha_0(n) + o(1)}$, we have $n/k_0 = \alpha_0(n) - 1 - 2/\log 2 + o(1) \leq \beta(n) - \varepsilon/2$, so $k_0 \in I_n$ with $\varepsilon/4$ room to spare.

By Lemma 38 we have $\widehat{L}_0(n, k_0, \beta(n)) = O(1)$; we chose k_0 so that the main term in (62) vanishes, leaving only the error term. Since, as a function of k , \widehat{L}_0 has derivative $\Theta(\log^3 n/n)$, it follows immediately that $k^*(n)$ exists, and that $k^*(n) - k_0(n) = O(n/\log^3 n)$. Since k^* and k_0 are of order $n/\log n$, this translates to $n/k^* = n/k_0 + O(1/\log n)$, proving the first statement.

For the second statement, recall the bound

$$\log(E_{n,k,\beta}) = k\widehat{L}_0(n, k, \beta) + O(\log^4 n) \tag{65}$$

given by Corollary 31 and (58). Consider $k = k^*(n) + x$, where x will be of larger order than $\log^2 n$ but not too large (say $o(n/\log^2 n)$). Then from the derivative bound, $\widehat{L}_0(n, k, \beta) = \Theta(x \log^3 n/n)$, so $k\widehat{L}_0(n, k, \beta) = \Theta(x \log^2 n)$. For x of the magnitude indicated this quantity is $\omega(\log^4 n)$. Choosing such an x so that k is an integer, from (65) we conclude that $\log(E_{n,k,\beta}) > 0$, so $k_\beta(n) \leq k = k^*(n) + x$. A similar argument with x negative shows that $k_\beta(n) = k^*(n) + O(\log^2 n)$. \square

Lemma 42. *The function $k^*(n)$ is differentiable on W , and its derivative satisfies*

$$\left(\frac{dk^*(n)}{dn}\right)^{-1} = \frac{n}{k^*(n)} + \frac{2}{\log 2} + O(1/\log n).$$

Proof. The Implicit Function Theorem, applied to the continuously (in fact, infinitely) differentiable function $\widehat{L}_0(n, k, t)$ with t fixed tells us that $k^*(n)$, defined by $\widehat{L}_0(n, k^*, \beta(n)) = 0$, is differentiable, and that its derivative is $-\frac{\partial \widehat{L}_0}{\partial n} / \frac{\partial \widehat{L}_0}{\partial k}$. Writing ρ for n/k , by Lemma 34 and Corollary 37 the reciprocal of the derivative is thus

$$-\frac{\partial \widehat{L}_0}{\partial k} / \frac{\partial \widehat{L}_0}{\partial n} = \frac{n}{k} - \frac{\rho - 1}{\log n - y_\beta(\rho)} = \rho + \frac{\rho + O(1)}{\log n - \log \log n + O(1)}.$$

Now $\rho = \alpha_0(n) + O(1) = 2 \log_2 n - 2 \log_2 \log_2 n + O(1)$, so the last fraction above is

$$\frac{2 \log_2 n - 2 \log_2 \log_2 n + O(1)}{\log n - \log \log n + O(1)} = \frac{2 \log_2 n - 2 \log_2 \log_2 n + O(1)}{(\log 2)(\log_2 n - \log_2 \log_2 n) + O(1)}$$

since $\log x = (\log 2) \log_2 x$, and hence $\log \log n = \log(\log_2 n) + O(1) = (\log 2) \log_2 \log_2 n + O(1)$. The result follows. \square

Together, Lemmas 41 and 42 imply Lemma 26, so the proof of Lemma 26 is complete.

3.2 | Proof of Lemma 28

We shall prove the following sharper form of Lemma 28, since it seems that the lower bound here is perhaps quite close to the truth (see the discussion in the Appendix), so this might be useful elsewhere.

Lemma 43. *Suppose that $\log^5 n \leq \mu_{\alpha(n)}(n) = O(n/\log^2 n)$. Then, whp,*

$$\chi(G_{n,1/2}) \geq k^*(n) - (1 + \varepsilon) \frac{\mu \log \nu}{\alpha(\log n - \log \log n)},$$

where $\alpha = \alpha(n)$, $\mu = \mu(n) = \mu_{\alpha(n)}(n)$, $\nu = (n/\log n)/\mu$, $\varepsilon = \varepsilon(n) = O(1/\log \nu) \rightarrow 0$, and $k^*(n) = k_{\beta}^*(n)$ is defined in Lemma 41.

Before giving the proof, we note that the result we need, Lemma 28, follows.

Proof of Lemma 28. This is immediate from Lemma 43, noting that by assumption $\nu(n)$ as defined there is $\Theta(\log n)$, so $\log \nu \sim \log \log n$, recalling that $\alpha(n) \sim c_0 \log n$, and noting that by Lemma 41, $k_{\beta} - k^*(n) = O(\log^2 n)$, which is much smaller than the error term we are aiming for. \square

Proof of Lemma 43. Let

$$\delta = C/\log \nu,$$

where $C \geq 3$ is a constant that we will specify later. Let $k_0 = \lfloor k^*(n) - d \rfloor$, where

$$d = (1 + 5\delta) \frac{\mu \log \nu}{\alpha(\log n - \log \log n)},$$

so our aim is to show that whp $\chi(G_{n,1/2}) \geq k_0$. To do this, it suffices to show that whp $G_{n,1/2}$ has no proper k_0 -colouring. Note for later that $k^*(n) = \Theta(n/\log n)$ while, recalling our assumptions on μ , we have $d = O(\mu/\log n) = O(n/\log^3 n)$. Thus, crudely, $d = o(k^*/\log n)$ and it follows easily that

$$n/k_0 = n/k^*(n) + o(1). \tag{66}$$

Let $\alpha = \alpha(n)$. By assumption, $\mu := \mu_{\alpha}(n) = O(n/\log^2 n)$, so $\mu_{\alpha+1}(n) = O(\mu \log n/n) = O(1/\log n) \rightarrow 0$, and whp $G_{n,1/2}$ contains no independent sets of size $\alpha + 1$. Thus it suffices to show that whp $G_{n,1/2}$ has no α -bounded k_0 -colouring.

We will group the potential colourings (or, more precisely, partitions into independent sets), according to the number m of α -sets included. Let

$$m^+ = \mu(1 + \delta).$$

Recalling that X_{α} , the number of independent α -sets, has mean μ and variance $O(\mu)$, we know from Chebyshev's inequality that whp $X_{\alpha} \leq m^+$. Thus it suffices to show that whp $G_{n,1/2}$ has no α -bounded k_0 -colouring using at most m^+ α -sets.

Let C_m denote the number of partitions of $[n] = V(G_{n,1/2})$ into exactly k_0 independent sets of which exactly m have size α and none has size larger than α . We claim that, if n is large enough, for each $m \leq m^+$ we have

$$\mathbb{E}[C_m] \leq 1/n. \tag{67}$$

Assuming this, then summing over the $m^+ + 1 = O(\mu) = o(n)$ values of m and applying Markov's inequality, the proof is complete. Thus it suffices to prove (67).

Now a potential colouring/partition of the type counted by C_m may be described as follows: we pick an unordered m -tuple of disjoint α -vertex subsets of $[n]$, and then we pick a partition P of the remaining $n - \alpha m$ vertices into $k_0 - m$ parts of size at most $\alpha - 1$. The partition gives a legal colouring if and only if the m α -sets are independent, and P induces a legal colouring of the corresponding subgraph of G . Hence,

$$\mathbb{E}[C_m] = \frac{1}{m!} \binom{n}{\alpha} \binom{n - \alpha}{\alpha} \dots \binom{n - (m - 1)\alpha}{\alpha} 2^{-m \binom{\alpha}{2}} E_{n - \alpha m, k_0 - m, \alpha - 1},$$

where $E_{n', k', t}$ is the expected number of t -bounded unordered k' -colourings of $G_{n', 1/2}$. Hence, bounding each binomial coefficient above by $\binom{n}{\alpha}$, we have

$$\mathbb{E}[C_m] \leq \frac{\mu^m}{m!} E_{n - \alpha m, k_0 - m, \alpha - 1}.$$

Taking logs, and using the standard bound $\mu^m / m! \leq e^\mu$ (the former is one term in the expansion of the latter), we see that

$$\log \mathbb{E}[C_m] \leq \mu + \log E_m$$

where $E_m := E_{n - \alpha m, k_0 - m, \alpha - 1}$.

Fortunately, we have a good approximation for $\log E_m$. Recalling (66) and noting from Lemma 41 that $n/k^*(n) \leq \alpha(n) - 2/\log 2 < \alpha(n) - 2$, we have $n/k_0 < \alpha - 1$ for n large enough, and it follows that

$$\frac{n - \alpha m}{k_0 - m} \leq \frac{n}{k_0} < \alpha - 1.$$

Thus we can apply Lemma 29 to conclude that

$$\log E_m = L_0(n - \alpha m, k_0 - m, \alpha - 1) + O(\log^4 n),$$

where L_0 is defined in (48).

Unfortunately we do not have a direct formula for L_0 sufficiently accurate for our present purpose. Fortunately, however, we do have indirect bounds, expressed in terms of $k^*(n)$, defined in Lemma 41. Note that we will consider a range of values n' satisfying $n' \in I$, where

$$I = [n - \alpha m^+, n].$$

Since $\alpha m^+ = O(\alpha\mu) = O(n/\log n)$, it follows easily that $\mu_\alpha(n') = \Theta(\mu)$ for all such n' . In particular, $\alpha(n') = \alpha$ does not vary over this range of n' , and it makes sense to consider $k^*(n')$ as in Lemma 41, defined with $\beta = \alpha - 1$, as a function of n' .

By definition $L_0(n - \alpha m, k^*(n - \alpha m), \alpha - 1) = 0$ (see (64)). From the last part of Corollary 39 we thus have

$$L_0(n - \alpha m, k_0 - m, \alpha - 1) \sim c_0 \log^2 n(k_0 - m - k^*(n - \alpha m)), \tag{68}$$

where $c_0 = 2/\log 2$. Since k_0 is defined in terms of $k^*(n)$, the next step is to consider how $k^*(n')$ varies as n' varies between n and $n - \alpha m$.

Now by Lemma 26, for $n' \in I$ we have

$$\left(\frac{dk^*(n')}{dn'}\right)^{-1} = \alpha + \frac{\log(\mu_\alpha(n'))}{\log n' - \log \log n'} - 1 + O(1/\log n'),$$

recalling that $\alpha(n') = \alpha$ for all $n' \in I$. For $n' \in I$ we have $\log n' = \log n + o(1)$ and, as noted above, $\mu_\alpha(n') = \Theta(\mu)$. It follows that

$$\left(\frac{dk^*(n')}{dn'}\right)^{-1} = \alpha + \frac{\log \mu}{\log n - \log \log n} - 1 + O(1/\log n) = \alpha - \frac{\log \nu}{\log n - \log \log n} + O(1/\log n),$$

recalling that $\nu = (n/\log n)/\mu$. Thus,

$$\left(\frac{dk^*(n')}{dn'}\right)^{-1} \geq \alpha - (1 + \delta) \frac{\log \nu}{\log n - \log \log n},$$

provided the constant C appearing in the definition of δ is chosen large enough.

We now take the reciprocal. Using the expansion $(\alpha - x)^{-1} = \alpha^{-1}(1 - x/\alpha)^{-1} = \alpha^{-1} + x\alpha^{-2} + \dots$ we see that for $n' \in I$ we have

$$\frac{dk^*(n')}{dn'} \leq \frac{1}{\alpha} + (1 + 2\delta) \frac{\log \nu}{\alpha^2(\log n - \log \log n)}.$$

For any $m \leq m^+$ this estimate applies for all n' in the interval $(n - \alpha m, n) \subset I$, so it follows immediately that

$$k^*(n) - k^*(n - \alpha m) \leq m + (1 + 2\delta) \frac{m \log \nu}{\alpha(\log n - \log \log n)}.$$

Hence

$$\begin{aligned} k_0 - m - k^*(n - \alpha m) &= k_0 - k^*(n) - m + (k^*(n) - k^*(n - \alpha m)) \\ &\leq -d - m + (k^*(n) - k^*(n - \alpha m)) \\ &\leq -d + (1 + 2\delta) \frac{m \log \nu}{\alpha(\log n - \log \log n)} \end{aligned}$$

$$\begin{aligned} &\leq -d + (1 + 4\delta) \frac{\mu \log \nu}{\alpha(\log n - \log \log n)} \\ &\leq -\frac{\delta \mu \log \nu}{\alpha(\log n - \log \log n)} \\ &\leq -\frac{C\mu}{\alpha \log n}, \end{aligned}$$

where in the last three steps we used the fact that $m \leq m^+ = (1 + \delta)\mu$, then the definition of d , and finally the definition of δ .

Hence, from (68), if n is large enough

$$L_0(n - \alpha m, k_0 - m, \alpha - 1) \leq -0.99c_0 \log^2 n \frac{C\mu}{\alpha \log n} \leq -0.98C\mu \leq -2\mu,$$

recalling that $C \geq 3$.

Putting the pieces together, we have

$$\log \mathbb{E}[C_m] \leq \mu - 2\mu + O(\log^4 n) \sim -\mu,$$

recalling that $\mu \geq \log^5 n$ by assumption. Thus, if n is large enough, $\mathbb{E}[C_m] \leq 1/n$ with plenty of room to spare, giving (67). Thus the proof of Lemma 28, and hence of Theorem 8, is complete. \square

APPENDIX: INTUITION BEHIND CONJECTURES

In this section we motivate the more refined conjectures in Section 1.3.2. There are two basic starting points, both described previously, so we only recall them briefly. Firstly, the very first guess at the chromatic number is from the ‘expectation threshold’, the least k such that the expected number of partitions into k independent sets is larger than 1. In calculating this, since there are rather few profiles (a list specifying how many independent sets have each possible size) to consider, one can consider only the optimal profile.

This intuition fails immediately when we look at independent sets of size α : the naive ‘optimal profile’ is ‘unachievable’ because it would like us to use $\Theta(n/\log n)$ independent sets of size α , but the actual number X_α will be close to μ_α which (for most n) will be much smaller than this. So the first approximation is to consider α -sets separately, expecting (since the naive optimum is to use many more than there are) that we will use as many as we can, and then considering the expectation threshold for colourings without α -sets.

This same ‘unachievability’ phenomenon can also arise with $(\alpha - 1)$ -sets; again, the optimal profile would like to use $\Theta(n/\log n)$ of them. There are certainly enough present, but not necessarily enough disjoint ones. Numerical calculations carried out by the first author suggest that this is an issue for $\mu_{\alpha-1}(n)$ up to around n^{1+x_0} for some small positive constant x_0 .

As in Section 1.3.2, to avoid a discontinuity when α changes, from now on we work in terms of $a = a(n)$, chosen so that $\mu_a(n)$ is between $n^{1/2+\delta}$ and $n^{3/2-\delta}$ for some positive δ . We only consider the ‘good’ n , for which such an a exists. Then $a = \alpha$ or $\alpha - 1$. In the latter case $\mu_\alpha(n)$ is at most $n^{1/2-\delta}$. For us, the independent sets of size $\alpha = a + 1$ can be ignored in this case: there may be enough of them to affect the chromatic number significantly, but the standard deviation of X_α is at most around $n^{1/4}$, which is smaller than any of our predictions for $g(n)$. Heuristically, we include all α -sets in our colouring, but do not need to consider them any further.

As outlined above, our main heuristic (we discuss another below) is as follows: to colour we choose as large as possible a collection C of disjoint independent sets of size a . Then we assume that the rest of the graph can be coloured with colour classes of size $a - 1$ as predicted by the relevant expectation threshold. Let us write m for X_a , the number of independent sets of size a , which will typically be μ_a plus or minus order $\sqrt{\mu_a}$, recalling that the distribution of X_a is approximately Poisson, and hence asymptotically Gaussian when $\mu_a \rightarrow \infty$. We write t for the size of C . Somewhat informally, we need to understand: (I) roughly how big t is, and (hence) roughly how much t varies as m varies, and (II) how much a given change in t affects the $(a - 1)$ -bounded chromatic number of the remaining graph $G_{n',p}$, where $n' = n - ta$.

Let us rescale by writing

$$m = \frac{2xn}{a^2} \quad \text{and} \quad t = \frac{2yn}{a^2}.$$

Rather than consider the actual distribution of independent sets of size a , we work heuristically in the random hypergraph model $H_a(m)$, or rather the essentially equivalent variant where the m hyperedges are chosen independently and uniformly from all a -sets. Since two a -sets intersect with probability $\sim a^2/n$, we see that on average one a -set intersects $\sim 2x$ others.

Case 1: $x = o(1)$, that is, $m = o(n/\log^2 n)$. Then almost all a -sets intersect no others, so we have $t \sim m$. Moreover, if we add an extra a -set, it is very likely to be disjoint from the current maximum matching, so (somewhat informally)

$$\frac{dy}{dx} = \frac{dt}{dm} \sim 1.$$

Case 2: $x = \Theta(1)$. Here it is hard to say anything very precise, but it is nevertheless clear that $t = \Theta(m)$, since we still have a constant fraction of a -sets that intersect no others. Certainly we expect that for some[†] well-behaved increasing function $h : (0, \infty) \rightarrow (0, \infty)$ we have $y \sim h(x)$, and hence

$$\frac{dy}{dx} \sim h'(x) = \Theta(1).$$

Case 3: $x \rightarrow \infty$. This case is more difficult, but for our heuristic we assume that the maximum matching is at least approximately given by the first moment threshold in the random hypergraph, that is, by solving

$$\binom{m}{t} \frac{(n)_{at}}{(n)_a^t} \approx 1,$$

where $(n)_k$ is the falling factorial $n!/(n - k)!$, and the ratio above is the probability that t randomly chosen a -sets are disjoint. In turn this gives

$$\log x - \log y + 1 + o(1) = y(1 + O(at/n)) = \Theta(y), \tag{A.1}$$

[†] One can probably describe h in terms of the size of the largest independent set in a suitable random graph $G_{n,2x/n}$, but it is not clear that this adds much. In any case, we believe we understand the asymptotic behaviour as $x \rightarrow 0$ or $x \rightarrow \infty$ from cases 1 and 3.

and we arrive at

$$y = \Theta(\log x) \quad \text{and} \quad \frac{dy}{dx} = \Theta(1/x),$$

with the implicit constants being $1 + o(1)$ when $x = n^{o(1)}$.

Let us now turn to (II), considering how the $(a - 1)$ -bounded chromatic number of the rest of the graph, which we treat simply as $G_{n',p}$, $n' = n - at$, varies as t , and hence n' , varies. Heuristically, we assume the actual number of colours needed will be essentially the relevant first moment threshold, or rather the approximation k^* from Lemma 26.

If there are n' vertices left, then for each extra vertex covered by a -sets we expect to need $\frac{dk^*}{dn} |_{n=n'}$ fewer colours. We temporarily write γ for the reciprocal of this quantity. From Lemma 26 and Remark 27 (which tells us that we can replace $\alpha(n)$ by a in (43)) we have

$$\begin{aligned} \gamma &= a - 1 + \frac{\log \mu_a(n')}{\log n' - \log \log n'} + O(1/\log n') \\ &= a - 1 + \frac{\log \mu_a(n')}{\log n - \log \log n} + O(1/\log n) \end{aligned}$$

since we will always have $n' = \Theta(n)$.

It is convenient to work in terms of $\mu_{a+1} = \Theta(\mu_a \log n/n)$. Let $\mu' = \mu_{a+1}(n')$, then

$$\begin{aligned} \gamma &= a + \frac{\log \mu'}{\log n - \log \log n} + O(1/\log n) \\ &= a + (1 + o(1)) \frac{\log \mu'}{\log n}, \end{aligned}$$

since we will see later that $-\log \mu'$ is at minimum at least $\omega(1)$ (in fact at least order $\log \log n$). This gives

$$\frac{dk^*}{dn} = \gamma^{-1} = a^{-1} - (1 + o(1)) \frac{\log \mu'}{a^2 \log n}.$$

So each extra a -set in the matching should save a times this many colours, minus the one used for the set itself, giving ‘benefit’ (per a -set used)

$$B \sim \frac{-\log \mu'}{a \log n} \sim \frac{-\log \mu'}{c_0 \log^2 n}$$

where $c_0 = 2/\log 2$.

Now

$$\mu_{a+1}(n) = \Theta(\mu_a(n) \log n/n) = \Theta(x/\log n).$$

In all cases, writing \approx for agreement up to constant factors,

$$\frac{\mu'}{\mu_{a+1}(n)} \sim (n'/n)^{a+1} \approx (n'/n)^a = (1 - at/n)^a = (1 - 2y/a)^a.$$

In cases 1 and 2, where x and hence y are $O(1)$, this is $\Theta(1)$ and hence irrelevant. In these cases we thus have $\mu' = \Theta(x/\log n)$, so $-\log \mu' \sim \log \log n + |\log x|$. Thus

$$B \sim \frac{\log \log n + |\log x|}{c_0 \log^2 n}.$$

In case 3, when x grows but not too quickly, say $x = n^{o(1)}$, then $y = o(\log n)$ and hence, from (A.1), $y \sim \log x$. Then

$$\frac{\mu'}{\mu_{a+1}(n)} \approx (1 - 2y/a)^a = \exp(-(2 + o(1)) \log x),$$

so μ' is roughly $1/(x \log n)$, with asymptotic agreement in the logarithms. In this case we thus obtain

$$B \sim \frac{\log \log n + \log x}{c_0 \log^2 n}.$$

Finally, if x is at least $n^{\Omega(1)}$, then cruder estimates give $\mu' = n^{-\Omega(1)}$, so $-\log \mu' = \Theta(\log x)$. In this case

$$B = \Theta\left(\frac{\log x}{\log^2 n}\right).$$

In all cases, multiplying $\sqrt{\mu_a}$, our estimate for how much the number m of independent a -sets varies, by $\frac{dy}{dx} = \frac{dt}{dm}$, and then by B , gives our estimate for $g(n)$, the standard deviation of $\chi(G_{n,p})$.

A.1 | Complications

In this subsection we discuss a number of issues that arise when attempting to understand the behaviour of $\chi(G_{n,p})$ even more precisely. First, we should note that in any attempt at proving Conjecture 15, there are major problems with the heuristic above. The key one is that, having removed some collection of independent a -sets, the graph that remains certainly does not have the same distribution as $G_{n',p}$ for appropriate n' . But even at the intuitive level, there are additional complications.

For one thing, the alert reader may have noticed that our heuristic above does not make sense in case 3 when x is too large, in particular when $\mu_a \geq n^{1+x_0}$, the point up to which the naive optimum profile wants us to use more disjoint a -sets than can be found. Here we justify our prediction rather by the heuristic in Section 1.3.1. With t fixed, then as $m = X_a$ varies, the number of ways of choosing t (disjoint) a -sets varies, and this translates into variation in the chromatic number. Fortunately, for $\mu_a = n^{1+\Theta(1)}$ the two predictions agree within a constant factor, so we do not need to resolve exactly how they interact.

This same effect arises in other cases, however. Suppose we have a strategy for partially colouring with a -sets where we use a slightly smaller than maximum matching, so there are $N \gg 1$ choices for this matching. Then we might expect to find a colouring if the expected number of $(a - 1)$ -bounded colourings of the remaining n' vertices is roughly $1/N$. As noted earlier, from

Corollary 39, for given n' we should expect the extra N choices to lead to a reduction in k^* of around $\log N / (c_0 \log^2 n') = \Theta(\log N / \log^2 n)$.

Considering the simpler case in which almost all a -sets are disjoint, we have $N \approx \binom{m}{t}$, so there is a large increase in the number of choices for leaving out the first few a -sets. Our calculations suggest that in this range we will leave out order $\Theta(\mu_a^2 \log n / n)$ a -sets from a maximum matching. This will affect the chromatic number significantly, but we do not expect it to lead to a significant change in the variance of the chromatic number.

A further issue is that in our case $x \rightarrow \infty$, the first moment threshold is not a terribly good estimate of the size of a maximum matching of a -sets. In the case where x does not grow too quickly, a heuristic explanation is the following. Since two a -sets intersect with probability $\pi_0 \sim a^2/n$, we expect t a -sets to be disjoint with probability around $\exp(-\pi_0 \binom{t}{2})$. However, there is some variability in the number M of overlapping pairs of a -sets. This quantity, which is of order $M_0 = \binom{m}{2} a^2/n$, varies by around $\Delta = \sqrt{M_0}$. If we condition on this number, then our new heuristic for the probability t a -sets are disjoint is $\exp(-\pi \binom{t}{2})$ where $\pi = M / \binom{m}{2} = \pi_0(1 + \Delta/M_0)$. This variation may well be significant, and it leads to a situation where the overall expectation of the number of t -matchings (collections of t disjoint a -sets) is dominated by the contribution from the case where M is atypically small. Hence the first moment will not be an accurate guide to the existence of a t -matching. We do not explore this further here since it does not seem to affect $g(n)$. However, this, and more complicated such effects, would (at least in some cases) alter $f(n)$ by a significant amount. Thus the problem of predicting, let alone proving, a ‘full result’ $(\chi(G_{n,p}) - f(n)) / \sqrt{g(n)} \xrightarrow{d} N(0, 1)$ with explicit functions f and g seems extremely difficult.

ACKNOWLEDGEMENTS

This author’s research was funded by ERC Grants 676632-RanDM and 772606-PTRCSP.

JOURNAL INFORMATION

The *Journal of the London Mathematical Society* is wholly owned and managed by the London Mathematical Society, a not-for-profit Charity registered with the UK Charity Commission. All surplus income from its publishing programme is used to support mathematicians and mathematics research in the form of research grants, conference grants, prizes, initiatives for early career researchers and the promotion of mathematics.

ORCID

Annika Heckel  <https://orcid.org/0000-0001-8961-0574>

REFERENCES

1. D. Achlioptas and A. Naor, *The two possible values of the chromatic number of a random graph*, *Ann. of Math.* **162** (2005), 1335–1351.
2. N. Alon and M. Krivelevich, *The concentration of the chromatic number of random graphs*, *Combinatorica* **17** (1997), no. 3, 303–313.
3. N. Alon and J. Spencer, *The probabilistic method (with an open problems appendix by Paul Erdős)*, 1st ed., Wiley, New York, 1992.
4. N. Alon and J. Spencer, *The probabilistic method*, 4th ed., Wiley, Hoboken, NJ, 2015.
5. B. Bollobás, *The chromatic number of random graphs*, *Combinatorica* **8** (1988), no. 1, 49–55.
6. B. Bollobás, *Random graphs*, 2nd ed., Cambridge University Press, 2001.
7. B. Bollobás, *How sharp is the concentration of the chromatic number?* *Combin. Probab. Comput.* **13** (2004), no. 01, 115–117.

8. B. Bollobás and P. Erdős, *Cliques in random graphs*, Mathematical Proceedings of the Cambridge Philosophical Society, vol. 80, Cambridge University Press, 1976, pp. 419–427.
9. F. Chung and R. Graham, *Erdős on Graphs: his legacy of unsolved problems*, AK Peters/CRC Press, Boca Raton, FL, 1998.
10. A. Coja-Oghlan, K. Panagiotou, and A. Steger, *On the chromatic number of random graphs*, J. Combin. Theory Ser. B **98** (2008), no. 5, 980–993.
11. P. Erdős and A. Rényi, *On the evolution of random graphs*, Publ. Math. Institute Hung. Acad. Sci. **5** (1960), 17–61.
12. N. Fountoulakis, R. Kang, and C. McDiarmid, *The t -stability number of a random graph*, Electron. J. Combin. **17** (2010), no. 1, R59.
13. R. Glebov, A. Liebenau, and T. Szabó, *On the concentration of the domination number of the random graph*, SIAM J. Discrete Math. **29** (2015), no. 3, 1186–1206.
14. G. R. Grimmett and C. McDiarmid, *On colouring random graphs*, Mathematical Proceedings of the Cambridge Philosophical Society, vol. 77, Cambridge University Press, 1975, pp. 313–324.
15. A. Heckel, *The chromatic number of dense random graphs*, Random Structures Algorithms **53** (2018), no. 1, 140–182.
16. A. Heckel, *Non-concentration of the chromatic number of a random graph*, J. Amer. Math. Soc. **34** (2021), 245–260.
17. A. Heckel and K. Panagiotou, *Colouring random graphs: tame colourings*, arxiv.org/abs/2306.07253, 2023.
18. R. Kang and C. McDiarmid, *Colouring random graphs*, Topics in chromatic graph theory, Encyclopedia of Mathematics and Its Applications, vol. 156, Cambridge University Press, 2015, pp. 199–229.
19. T. Łuczak, *The chromatic number of random graphs*, Combinatorica **11** (1991), no. 1, 45–54.
20. T. Łuczak, *A note on the sharp concentration of the chromatic number of random graphs*, Combinatorica **11** (1991), no. 3, 295–297.
21. D. Matula, *On the complete subgraphs of a random graph*, Proceedings of the 2nd Chapel Hill Conference on Combinatorial Mathematics and its Applications (Chapel Hill, NC, 1970), 1970, pp. 356–369.
22. D. Matula, *The employee party problem*, Notices Amer. Math. Soc. **19** (1972), no. 2, A–382.
23. C. McDiarmid, *On the method of bounded differences*, Surv. Comb. **141** (1989), no. 1, 148–188.
24. K. Panagiotou and A. Steger, *A note on the chromatic number of a dense random graph*, Discrete Math. **309** (2009), no. 10, 3420–3423.
25. A. Scott, *On the concentration of the chromatic number of random graphs*, arxiv.org/abs/0806.0178, 2008.
26. E. Shamir and J. Spencer, *Sharp concentration of the chromatic number on random graphs $G_{n,p}$* , Combinatorica **7** (1987), no. 1, 121–129.